

# Table of Contents

Table of Contents	1
pgRouting Manual (3.8)	14
pgRouting Manual (3.8)	14
Table of Contents	14
General	14
Introduction	14
Licensing	14
Contributors	14
This Release Contributors	14
Individuals in this release v3.8.x (in alphabetical order)	14
Translators (in alphabetical order)	14
Corporate Sponsors in this release (in alphabetical order)	14
Contributors Past & Present:	14
Individuals (in alphabetical order)	14
Corporate Sponsors (in alphabetical order)	14
More Information	15
Installation	15
Short Version	15
Get the sources	15
Enabling and upgrading in the database	15
Dependencies	15
Configuring	16
Configurable variables	17
Building	17
Testing	17
See Also	17
Support	17
Reporting Problems	18
Mailing List and GIS StackExchange	18
Commercial Support	18
Sample Data	18
Main graph	18
Edges	19
Edges data	19
Vertices	20
Vertices data	20
The topology	20
Topology data	20
Points outside the graph	20
Points of interest	20
Points of interest fill up	21
Support tables	21
Combinations	21
Combinations data	21
Restrictions	22
Restrictions data	22
Images	22
Directed graph with cost and reverse_cost	22
Undirected graph with cost and reverse_cost	22
Directed graph with cost	22
Undirected graph with cost	23
Pick & Deliver Data	23
The vehicles	23
The original orders	23
The orders	24
Pgrouting Concepts	24
pgRouting Concepts	24
Graphs	25
Graph definition	25
Graph with cost	25
Graph with cost and reverse_cost	26
Graphs without geometries	27
Wiki example	27
Prepare the database	27
Create a table	27
Insert the data	28
Find the shortest path	28
Vertex information	28
Graphs with geometries	28
Create a routing Database	28
Load Data	28
Build a routing topology	29
Adjust costs	29
Update costs to length of geometry	29
Update costs based on codes	29
Check the Routing Topology	30
Crossing edges	30
Fixing an intersection	31
Touching edges	31
Fixing a gap	32
Connecting components	33
Contraction of a graph	34
Dead ends	34
Linear edges	34
Function's structure	34
Function's overloads	35
One to One	35
One to Many	35
Many to One	35
Many to Many	35
Combinations	35
Inner Queries	36
Edges SQL	36
General	36
General without id	36
General with (X,Y)	37
Flow	37
Combinations SQL	37
Restrictions SQL	38
Points SQL	38
Parameters	39
Parameters for the Via functions	39
For the TRSP functions	39
Result columns	39
Result columns for a path	40
Multiple paths	41

Selective for multiple paths¶	41
Non selective for multiple paths¶	41
Result columns for cost functions¶	42
Result columns for flow functions¶	42
Result columns for spanning tree functions¶	42
Performance Tips¶	43
For the Routing functions¶	43
How to contribute¶	43
Function Families¶	43
Function Families¶	43
Functions by categories¶	44
All Pairs - Family of Functions¶	45
pgr_floydWarshall¶	45
Description¶	45
Signatures¶	45
Parameters¶	46
Optional parameters¶	46
Inner Queries¶	46
Edges SQL¶	46
Result columns¶	46
See Also¶	46
pgr_johnson¶	46
Description¶	47
Signatures¶	47
Parameters¶	47
Optional parameters¶	47
Inner Queries¶	47
Edges SQL¶	47
Result columns¶	48
See Also¶	48
Introduction¶	48
Parameters¶	48
Optional parameters¶	48
Inner Queries¶	48
Edges SQL¶	48
Result columns¶	49
Performance¶	49
Data¶	49
Results¶	49
See Also¶	51
A* - Family of functions¶	51
pgr_aStar¶	51
Description¶	51
Signatures¶	52
One to One¶	52
One to Many¶	52
Many to One¶	52
Many to Many¶	53
Combinations¶	53
Parameters¶	53
Optional parameters¶	54
aStar optional parameters¶	54
Inner Queries¶	54
Edges SQL¶	54
Combinations SQL¶	54
Result columns¶	55
Additional Examples¶	55
See Also¶	56
pgr_aStarCost¶	56
Description¶	56
Signatures¶	57
One to One¶	57
One to Many¶	57
Many to One¶	57
Many to Many¶	57
Combinations¶	57
Parameters¶	58
Optional parameters¶	58
aStar optional parameters¶	58
Inner Queries¶	58
Edges SQL¶	58
Combinations SQL¶	59
Result columns¶	59
Additional Examples¶	59
See Also¶	60
pgr_aStarCostMatrix¶	60
Description¶	60
Signatures¶	60
Parameters¶	60
Optional parameters¶	61
aStar optional parameters¶	61
Inner Queries¶	61
Edges SQL¶	61
Result columns¶	61
Additional Examples¶	62
See Also¶	62
Description¶	62
aStar optional parameters¶	62
Advanced documentation¶	63
Heuristic¶	63
Factor¶	63
See Also¶	63
Bidirectional A* - Family of functions¶	63
pgr_bdAStar¶	63
Description¶	64
Signatures¶	64
One to One¶	64
One to Many¶	64
Many to One¶	65
Many to Many¶	65
Combinations¶	65
Parameters¶	66
Optional parameters¶	66
aStar optional parameters¶	66
Inner Queries¶	66
Edges SQL¶	66
Combinations SQL¶	67
Result columns¶	67
Additional Examples¶	67
See Also¶	68
pgr_bdAStarCost¶	68
Description¶	68
Signatures¶	69
One to One¶	69
One to Many¶	69
Many to One¶	69
Many to Many¶	70
Combinations¶	70
Parameters¶	70
Optional parameters¶	70
aStar optional parameters¶	70
Inner Queries¶	71
Edges SQL¶	71
Combinations SQL¶	71
Result columns¶	71
Additional Examples¶	72
See Also¶	72
pgr_bdAStarCostMatrix¶	72
Description¶	72
Signatures¶	73
Parameters¶	73
Optional parameters¶	73
aStar optional parameters¶	73
Inner Queries¶	73
Edges SQL¶	73
Result columns¶	74
Additional Examples¶	74
See Also¶	74
Description¶	74
See Also¶	75

Bidirectional Dijkstra - Family of functions¶	75
pgr_bdDijkstra¶	75
Description¶	75
Signatures¶	75
One to One¶	76
One to Many¶	76
Many to One¶	76
Many to Many¶	76
Combinations¶	77
Parameters¶	77
Optional parameters¶	77
Inner Queries¶	77
Edges SQL¶	77
Combinations SQL¶	78
Result columns¶	78
Additional Examples¶	78
See Also¶	79
pgr_bdDijkstraCost¶	79
Description¶	79
Signatures¶	80
One to One¶	80
One to Many¶	80
Many to One¶	80
Many to Many¶	80
Combinations¶	80
Parameters¶	81
Optional parameters¶	81
Inner Queries¶	81
Edges SQL¶	81
Combinations SQL¶	81
Result columns¶	81
Additional Examples¶	82
See Also¶	82
pgr_bdDijkstraCostMatrix¶	82
Description¶	82
Signatures¶	83
Parameters¶	83
Optional parameters¶	83
Inner Queries¶	84
Edges SQL¶	84
Result columns¶	84
Additional Examples¶	84
See Also¶	84
Synopsis¶	84
Characteristics¶	84
See Also¶	85
Components - Family of functions¶	85
pgr_connectedComponents¶	85
Description¶	85
Signatures¶	86
Parameters¶	86
Inner Queries¶	86
Edges SQL¶	86
Result columns¶	86
Additional Examples¶	87
Connecting disconnected components¶	87
See Also¶	88
pgr_strongComponents¶	88
Description¶	88
Signatures¶	88
Parameters¶	89
Inner Queries¶	89
Edges SQL¶	89
Result columns¶	89
See Also¶	89
pgr_biconnectedComponents¶	89
Description¶	90
Signatures¶	90
Parameters¶	90
Inner Queries¶	90
Edges SQL¶	90
Result columns¶	91
See Also¶	91
pgr_articulationPoints¶	91
Description¶	91
Signatures¶	91
Parameters¶	92
Inner Queries¶	92
Edges SQL¶	92
Result columns¶	92
See Also¶	92
pgr_bridges¶	92
Description¶	92
Signatures¶	93
Parameters¶	93
Inner Queries¶	93
Edges SQL¶	93
Result columns¶	93
See Also¶	93
pgr_makeConnected - Experimental¶	94
Description¶	94
Signatures¶	94
Parameters¶	94
Inner Queries¶	94
Edges SQL¶	94
Result columns¶	95
See Also¶	95
See Also¶	95
Contraction - Family of functions¶	95
pgr_contraction¶	96
Description¶	96
Signatures¶	96
Parameters¶	96
Optional parameters¶	97
Contraction optional parameters¶	97
Inner Queries¶	97
Edges SQL¶	97
Result columns¶	97
Additional Examples¶	98
Only dead end contraction¶	98
Only linear contraction¶	98
The cycle¶	98
Contracting sample data¶	98
Construction of the graph in the database¶	98
The process to create the contraction graph on the database¶	99
Add additional columns¶	99
Store contraction information¶	99
Update the edges and vertices tables¶	99
The contracted graph¶	99
Using the contracted graph¶	100
See Also¶	101
pgr_contractionDeadEnd - Proposed¶	101
Description¶	101
Signatures¶	101
Parameters¶	101
Optional parameters¶	102
Contraction optional parameters¶	102
Inner Queries¶	102
Edges SQL¶	102
Result columns¶	102
Additional Examples¶	102
Dead end vertex on undirected graph¶	103
Dead end vertex on directed graph¶	103
Step by step dead end contraction¶	103
Creating the contracted graph¶	104
Steps for the creation of the contracted graph¶	104
The contracted graph¶	104
Using when departure and destination are in the contracted graph¶	104
Using when departure/destination is not in the contracted graph¶	104
Using when departure and destination are not in the contracted graph¶	104
See Also¶	105
pgr_contractionLinear - Proposed¶	105
Description¶	105
Signatures¶	105

Parameters¶	105
Optional parameters¶	105
Contraction optional parameters¶	106
Inner Queries¶	106
Edges SQL¶	106
Result columns¶	106
Additional Examples¶	106
Linear edges¶	106
Linearity is not symmetrical¶	107
Linearity is symmetrical¶	107
Step by step linear contraction¶	108
Creating the contracted graph¶	108
Steps for the creation of the contracted graph¶	108
The contracted graph¶	108
Using when departure and destination are in the contracted graph¶	109
Using when departure/destination is not in the contracted graph¶	109
See Also¶	109
pgr_contractionHierarchies - Experimental¶	109
Description¶	109
Signatures¶	109
Parameters¶	110
Optional parameters¶	110
Contraction hierarchies optional parameters¶	110
Inner Queries¶	110
Edges SQL¶	110
Result columns¶	110
Examples¶	111
On an undirected graph¶	111
On an undirected graph with forbidden vertices¶	112
Contraction process steps details¶	112
Shortcut building process¶	112
Initialize the queue with a first vertices order¶	112
Build the final vertex order¶	113
Add shortcuts to the initial graph¶	113
Use the contraction¶	113
Build the contraction¶	113
Add shortcuts and hierarchy in the existing tables¶	113
Use a Dijkstra shortest path algorithm on it¶	113
See Also¶	113
Introduction¶	113
See Also¶	113
Dijkstra - Family of functions¶	113
pgr_dijkstra¶	114
Description¶	114
Signatures¶	114
One to One¶	115
One to Many¶	115
Many to One¶	115
Many to Many¶	115
Combinations¶	116
Parameters¶	116
Optional parameters¶	116
Inner Queries¶	116
Edges SQL¶	116
Combinations SQL¶	117
Result columns¶	117
Additional Examples¶	117
For directed graphs with cost and reverse_cost columns¶	119
1) Path from \6) to \10)¶	119
2) Path from \6) to \7)¶	119
3) Path from \12) to \10)¶	119
4) Path from \12) to \7)¶	119
5) Using One to Many to get the solution of examples 1 and 2¶	119
6) Using Many to One to get the solution of examples 2 and 4¶	119
7) Using Many to Many to get the solution of examples 1 to 4¶	120
8) Using Combinations to get the solution of examples 1 to 3¶	120
For undirected graphs with cost and reverse_cost columns¶	120
9) Path from \6) to \10)¶	120
10) Path from \6) to \7)¶	120
11) Path from \12) to \10)¶	121
12) Path from \12) to \7)¶	121
13) Using One to Many to get the solution of examples 9 and 10¶	121
14) Using Many to One to get the solution of examples 10 and 12¶	121
15) Using Many to Many to get the solution of examples 9 to 12¶	121
16) Using Combinations to get the solution of examples 9 to 11¶	121
For directed graphs only with cost column¶	121
17) Path from \6) to \10)¶	122
18) Path from \6) to \7)¶	122
19) Path from \12) to \10)¶	122
20) Path from \12) to \7)¶	122
21) Using One to Many to get the solution of examples 17 and 18¶	122
22) Using Many to One to get the solution of examples 18 and 20¶	122
23) Using Many to Many to get the solution of examples 17 to 20¶	122
24) Using Combinations to get the solution of examples 17 to 19¶	122
For undirected graphs only with cost column¶	123
25) Path from \6) to \10)¶	123
26) Path from \6) to \7)¶	123
27) Path from \12) to \10)¶	123
28) Path from \12) to \7)¶	123
29) Using One to Many to get the solution of examples 25 and 28¶	123
30) Using Many to One to get the solution of examples 26 and 28¶	124
31) Using Many to Many to get the solution of examples 25 to 28¶	124
32) Using Combinations to get the solution of examples 25 to 27¶	124
Equivalences between signatures¶	124
33) Using One to One¶	124
34) Using One to Many¶	124
35) Using Many to One¶	124
36) Using Many to Many¶	124
37) Using Combinations¶	125
See Also¶	125
pgr_dijkstraCost¶	125
Description¶	125
Signatures¶	125
One to One¶	126
One to Many¶	126
Many to One¶	126
Many to Many¶	126
Combinations¶	126
Parameters¶	127
Optional parameters¶	127
Inner Queries¶	127
Edges SQL¶	127
Combinations SQL¶	127
Result columns¶	127
Additional Examples¶	128
See Also¶	128
pgr_dijkstraCostMatrix¶	128
Description¶	128
Signatures¶	129
Parameters¶	129
Optional parameters¶	129
Inner Queries¶	129
Edges SQL¶	129
Result columns¶	130
Additional Examples¶	130
See Also¶	130
pgr_drivingDistance¶	130
Description¶	131
Signatures¶	131
Single Vertex¶	131
Multiple Vertices¶	131
Parameters¶	131
Optional parameters¶	131
Driving distance optional parameters¶	132
Inner Queries¶	132
Edges SQL¶	132
Result columns¶	132
Additional Examples¶	132
See Also¶	133
pgr_KSP¶	133
Description¶	133
Signatures¶	133
One to One¶	133
One to Many¶	133
Many to One¶	134

Many to Many¶	134
Combinations¶	135
Parameters¶	135
Optional parameters¶	135
KSP Optional parameters¶	136
Inner Queries¶	136
Edges SQL¶	136
Combinations SQL¶	136
Result columns¶	136
Additional Examples¶	137
See Also¶	138
pgr_dijkstraVia - Proposed¶	138
Description¶	139
Signatures¶	139
One Via¶	139
Parameters¶	139
Optional parameters¶	139
Via optional parameters¶	139
Inner Queries¶	139
Edges SQL¶	139
Result columns¶	140
Additional Examples¶	140
The main query¶	140
Aggregate cost of the third path.¶	140
Route's aggregate cost of the route at the end of the third path.¶	141
Nodes visited in the route.¶	141
The aggregate costs of the route when the visited vertices are reached.¶	141
Status of "passes in front" or "visits" of the nodes.¶	141
See Also¶	141
pgr_dijkstraNear - Proposed¶	141
Description¶	142
Characteristics¶	142
Signatures¶	142
One to Many¶	142
Many to One¶	142
Many to Many¶	143
Combinations¶	143
Parameters¶	144
Dijkstra optional parameters¶	144
Near optional parameters¶	144
Inner Queries¶	144
Edges SQL¶	144
Combinations SQL¶	145
Result columns¶	145
See Also¶	145
pgr_dijkstraNearCost - Proposed¶	145
Description¶	145
Characteristics¶	146
Signatures¶	146
One to Many¶	146
Many to One¶	146
Many to Many¶	146
Combinations¶	147
Parameters¶	148
Dijkstra optional parameters¶	148
Near optional parameters¶	148
Inner Queries¶	148
Edges SQL¶	148
Combinations SQL¶	148
Result columns¶	148
See Also¶	149
Introduction¶	149
Parameters¶	149
Optional parameters¶	149
Inner Queries¶	149
Edges SQL¶	149
Combinations SQL¶	150
Advanced documentation¶	150
The problem definition (Advanced documentation)¶	150
See Also¶	151
Flow - Family of functions¶	151
pgr_maxFlow¶	151
Description¶	151
Signatures¶	152
One to One¶	152
One to Many¶	152
Many to One¶	152
Many to Many¶	152
Combinations¶	152
Parameters¶	153
Inner Queries¶	153
Edges SQL¶	153
Combinations SQL¶	153
Result columns¶	153
Additional Examples¶	153
See Also¶	154
pgr_boykovKolmogorov¶	154
Description¶	154
Signatures¶	154
One to One¶	154
One to Many¶	155
Many to One¶	155
Many to Many¶	155
Combinations¶	155
Parameters¶	155
Inner Queries¶	156
Edges SQL¶	156
Combinations SQL¶	156
Result columns¶	156
Additional Examples¶	156
See Also¶	157
pgr_edmondsKarp¶	157
Description¶	157
Signatures¶	157
One to One¶	157
One to Many¶	158
Many to One¶	158
Many to Many¶	158
Combinations¶	158
Parameters¶	159
Inner Queries¶	159
Edges SQL¶	159
Combinations SQL¶	159
Result columns¶	159
Additional Examples¶	159
See Also¶	160
pgr_pushRelabel¶	160
Description¶	160
Signatures¶	160
One to One¶	160
One to Many¶	161
Many to One¶	161
Many to Many¶	161
Combinations¶	161
Parameters¶	162
Inner Queries¶	162
Edges SQL¶	162
Combinations SQL¶	162
Result columns¶	162
Additional Examples¶	163
See Also¶	163
pgr_edgeDisjointPaths¶	163
Description¶	163
Signatures¶	163
One to One¶	163
One to Many¶	164
Many to One¶	164
Many to Many¶	164
Combinations¶	165
Parameters¶	165
Optional parameters¶	165
Inner Queries¶	165
Edges SQL¶	165
Combinations SQL¶	166
Result columns¶	166

Additional Examples¶	166
See Also¶	167
pgr_maxCardinalityMatch¶	167
Description¶	167
Signatures¶	167
Parameters¶	167
Inner Queries¶	167
Edges SQL¶	167
Result columns¶	168
See Also¶	168
pgr_maxFlowMinCost - Experimental¶	168
Description¶	168
Signatures¶	169
One to One¶	169
One to Many¶	169
Many to One¶	169
Many to Many¶	170
Combinations¶	170
Parameters¶	170
Inner Queries¶	170
Edges SQL¶	170
Combinations SQL¶	171
Result columns¶	171
Additional Examples¶	171
See Also¶	171
pgr_maxFlowMinCost_Cost - Experimental¶	172
Description¶	172
Signatures¶	172
One to One¶	173
One to Many¶	173
Many to One¶	173
Many to Many¶	173
Combinations¶	173
Parameters¶	173
Inner Queries¶	174
Edges SQL¶	174
Combinations SQL¶	174
Return columns¶	174
Additional Examples¶	174
See Also¶	175
Flow Functions General Information¶	175
Inner Queries¶	175
Edges SQL¶	175
Combinations SQL¶	176
Result columns¶	176
Advanced Documentation¶	177
See Also¶	177
Kruskal - Family of functions¶	177
pgr_kruskal¶	177
Description¶	178
Signatures¶	178
Parameters¶	178
Inner Queries¶	178
Edges SQL¶	178
Result columns¶	178
See Also¶	179
pgr_kruskalBFS¶	179
Description¶	179
Signatures¶	179
Single vertex¶	179
Multiple vertices¶	179
Parameters¶	180
BFS optional parameters¶	180
Inner Queries¶	180
Edges SQL¶	180
Result columns¶	180
See Also¶	181
pgr_kruskalDD¶	181
Description¶	181
Signatures¶	181
Single vertex¶	182
Multiple vertices¶	182
Parameters¶	182
Inner Queries¶	182
Edges SQL¶	182
Result columns¶	183
See Also¶	183
pgr_kruskalDFS¶	183
Description¶	183
Signatures¶	183
Single vertex¶	184
Multiple vertices¶	184
Parameters¶	184
DFS optional parameters¶	184
Inner Queries¶	184
Edges SQL¶	184
Result columns¶	185
See Also¶	185
Description¶	185
Inner Queries¶	185
See Also¶	186
Metrics - Family of functions¶	186
pgr_degree¶	186
Description¶	187
Signatures¶	187
Edges¶	187
Edges and Vertices¶	187
Parameters¶	188
Optional parameters¶	188
Inner Queries¶	188
Edges SQL¶	188
Vertex SQL¶	188
Result columns¶	188
Additional Examples¶	188
Degree of a loop¶	189
Degree of a sub graph¶	189
Using a vertex table¶	189
Dry run execution¶	189
Finding dead ends¶	190
Finding linear vertices¶	190
See Also¶	190
pgr_betweennessCentrality - Experimental¶	191
Description¶	191
Signatures¶	191
Parameters¶	191
Optional parameters¶	191
Inner Queries¶	191
Edges SQL¶	191
Result columns¶	192
See Also¶	192
See Also¶	192
Prim - Family of functions¶	192
pgr_prim¶	192
Description¶	192
Signatures¶	192
Parameters¶	193
Inner Queries¶	193
Edges SQL¶	193
Result columns¶	193
See Also¶	193
pgr_primBFS¶	193
Description¶	194
Signatures¶	194
Single vertex¶	194
Multiple vertices¶	194
Parameters¶	194
BFS optional parameters¶	195
Inner Queries¶	195
Edges SQL¶	195
Result columns¶	195
See Also¶	195
pgr_primDD¶	196
Description¶	196
Signatures¶	196

Single vertex¶	196
Multiple vertices¶	196
Parameters¶	197
Inner Queries¶	197
Edges SQL¶	197
Result columns¶	197
See Also¶	197
pgr_primDFS¶	198
Description¶	198
Signatures¶	198
Single vertex¶	198
Multiple vertices¶	198
Parameters¶	199
DFS optional parameters¶	199
Inner Queries¶	199
Edges SQL¶	199
Result columns¶	199
See Also¶	200
Description¶	200
Inner Queries¶	200
See Also¶	200
Reference¶	201
pgr_version¶	201
Description¶	201
Signature¶	201
Result columns¶	201
See Also¶	201
pgr_full_version¶	201
Description¶	201
Signatures¶	201
Result columns¶	201
See Also¶	202
See Also¶	202
Topology - Family of Functions¶	202
Utility functions¶	202
pgr_createTopology - Deprecated since v3.8.0¶	202
Migration of pgr_createTopology¶	203
Build a routing topology¶	203
Description¶	203
Signatures¶	203
Parameters¶	203
Usage when the edge table's columns MATCH the default values.¶	204
Usage when the edge table's columns DO NOT MATCH the default values.¶	205
Additional Examples¶	206
Create a routing topology¶	206
Make sure the database does not have the vertices_table¶	206
Clean up the columns of the routing topology to be created¶	206
Create the vertices table¶	206
Inspect the vertices table¶	206
Create the routing topology on the edge table¶	206
Inspect the routing topology¶	207
With full output¶	207
See Also¶	207
pgr_createVerticesTable - Deprecated since 3.8.0¶	208
Migration of pgr_createVerticesTable¶	208
Description¶	208
Signatures¶	208
Parameters¶	208
Additional Examples¶	209
Usage when the edge table's columns DO NOT MATCH the default values.¶	210
See Also¶	212
pgr_analyzeGraph -- Deprecated since 3.8.0¶	212
Migration of pgr_analyzeGraph¶	212
Description¶	213
Parameters¶	214
Usage when the edge table's columns MATCH the default values.¶	214
Usage when the edge table's columns DO NOT MATCH the default values.¶	216
Additional Examples¶	218
See Also¶	219
pgr_analyzeOneWay - Deprecated since 3.8.0¶	219
Migration of pgr_analyzeOneWay¶	219
Description¶	219
Signatures¶	220
Parameters¶	220
Additional Examples¶	220
See Also¶	221
pgr_nodeNetwork - Deprecated since 3.8.0¶	221
Migration of pgr_nodeNetwork¶	221
Description¶	221
Parameters¶	221
Examples¶	222
Fixing an intersection¶	222
Fixing a gap¶	223
See Also¶	224
pgr_extractVertices¶	224
Description¶	224
Signatures¶	224
Parameters¶	225
Optional parameters¶	225
Inner Queries¶	225
Edges SQL¶	225
When line geometry is known¶	225
When vertex geometry is known¶	225
When identifiers of vertices are known¶	225
Result columns¶	226
Additional Examples¶	226
Dry run execution¶	226
Create a routing topology¶	226
Make sure the database does not have the vertices_table¶	226
Clean up the columns of the routing topology to be created¶	227
Create the vertices table¶	227
Inspect the vertices table¶	227
Create the routing topology on the edge table¶	227
Inspect the routing topology¶	227
See Also¶	227
pgr_findCloseEdges¶	228
Description¶	228
Signatures¶	228
One point¶	228
Many points¶	228
Parameters¶	229
Optional parameters¶	229
Inner Queries¶	229
Edges SQL¶	229
Result columns¶	229
Additional Examples¶	229
One point in an edge¶	229
One point dry run execution¶	230
Many points in an edge¶	230
Many points dry run execution¶	230
Find at most two routes to a given point¶	231
A point of interest table¶	231
Points of interest¶	231
Points of interest fill up¶	231
See Also¶	231
pgr_separateCrossing¶	232
Description¶	232
Signature¶	232
Parameters¶	232
Optional parameters¶	232
Inner Queries¶	232
Edges SQL¶	232
Examples¶	232
Get the code for further refinement.¶	233
Fixing an intersection¶	233
See Also¶	234
pgr_separateTouching¶	234
Description¶	234
Signature¶	234
Parameters¶	234
Optional parameters¶	234
Inner Queries¶	234
Edges SQL¶	234
Examples¶	234

Get the code for further refinement.¶	235
Fixing a gap¶	235
See Also¶	236
See Also¶	236
Traveling Sales Person - Family of functions¶	236
pgr_TSP¶	236
Description¶	236
Problem Definition¶	236
Characteristics¶	236
Signatures¶	237
Parameters¶	237
TSP optional parameters¶	237
Inner Queries¶	237
Matrix SQL¶	238
Result columns¶	238
Additional Examples¶	238
Start from vertex \{1\}¶	238
Using points of interest to generate an asymmetric matrix.¶	238
Connected incomplete data¶	238
See Also¶	239
pgr_TSPeuclidean¶	239
Description¶	239
Problem Definition¶	239
Characteristics¶	239
Signatures¶	240
Parameters¶	240
TSP optional parameters¶	240
Inner Queries¶	240
Coordinates SQL¶	240
Result columns¶	240
Additional Examples¶	241
Test 29 cities of Western Sahara¶	241
Creating a table for the data and storing the data¶	241
Adding a geometry (for visual purposes)¶	241
Total tour cost¶	241
Getting a geometry of the tour¶	241
Visual results¶	241
See Also¶	241
General Information¶	242
Problem Definition¶	242
Origin¶	242
Characteristics¶	242
TSP optional parameters¶	242
See Also¶	242
BFS - Category¶	243
Parameters¶	243
BFS optional parameters¶	243
Inner Queries¶	243
Edges SQL¶	243
Result columns¶	243
See Also¶	244
Cost - Category¶	244
General Information¶	244
Characteristics¶	244
See Also¶	245
Cost Matrix - Category¶	245
General Information¶	245
Synopsis¶	245
Characteristics¶	245
Parameters¶	245
Optional parameters¶	246
Inner Queries¶	246
Edges SQL¶	246
Points SQL¶	246
Result columns¶	247
See Also¶	247
DFS - Category¶	247
See Also¶	247
Driving Distance - Category¶	248
pgr_alphaShape¶	248
Migration of pgr_alphaShape¶	248
Description¶	248
Signatures¶	249
Parameters¶	249
Return Value¶	249
See Also¶	249
Parameters¶	249
Inner Queries¶	249
Edges SQL¶	249
Result columns¶	250
See Also¶	250
K shortest paths - Category¶	250
Spanning Tree - Category¶	250
See Also¶	251
Via - Category¶	251
General Information¶	251
Parameters¶	251
Via optional parameters¶	252
Inner Queries¶	252
Edges SQL¶	252
Restrictions SQL¶	252
Points SQL¶	253
Result columns¶	253
See Also¶	253
Vehicle Routing Functions - Category¶	253
pgr_pickDeliver - Experimental¶	254
Synopsis¶	255
Characteristics¶	255
Signature¶	255
Parameters¶	255
Pick-Deliver optional parameters¶	256
Orders SQL¶	256
Vehicles SQL¶	257
Matrix SQL¶	257
Result columns¶	257
See Also¶	258
pgr_pickDeliverEuclidean - Experimental¶	258
Synopsis¶	259
Characteristics¶	259
Signature¶	259
Parameters¶	260
Pick-Deliver optional parameters¶	260
Orders SQL¶	260
Vehicles SQL¶	261
Result columns¶	261
Example¶	262
The vehicles¶	262
The original orders¶	263
The orders¶	263
The query¶	264
See Also¶	264
pgr_vrpOneDepot - Experimental¶	264
Description¶	264
Signatures¶	264
Parameters¶	264
Inner Queries¶	264
Result columns¶	265
Additional Example¶	265
See Also¶	265
Introduction¶	265
Characteristics¶	265
Pick & Delivery¶	266
Parameters¶	266
Pick & deliver¶	266
Pick-Deliver optional parameters¶	266
Inner Queries¶	266
Orders SQL¶	266
Vehicles SQL¶	267
Matrix SQL¶	267
Result columns¶	268
Summary Flow¶	269
Handling Parameters¶	270
Capacity and Demand Units Handling¶	270



Locations¶	270
Time Handling¶	270
Factor handling¶	270
See Also¶	271
withPoints - Category¶	271
Introduction¶	271
Parameters¶	272
Optional parameters¶	272
Inner Queries¶	272
Edges SQL¶	272
Points SQL¶	273
Combinations SQL¶	273
Advanced documentation¶	273
About points¶	273
Driving side¶	274
Right driving side¶	274
Left driving side¶	274
Driving side does not matter¶	275
Creating temporary vertices¶	275
On a right hand side driving network¶	275
On a left hand side driving network¶	276
When driving side does not matter¶	276
See Also¶	277
See Also¶	277

## Functions by categories¶278

## Available Functions but not official pgRouting functions¶279

### Proposed Functions¶279

#### TRSP - Family of functions¶279

##### pgr\_trsp - Proposed¶280

Description¶	281
Signatures¶	281
One to One¶	281
One to Many¶	281
Many to One¶	281
Many to Many¶	281
Combinations¶	282
Parameters¶	282
Optional parameters¶	282
Inner Queries¶	283
Edges SQL¶	283
Restrictions SQL¶	283
Combinations SQL¶	283
Result columns¶	283
See Also¶	284

##### pgr\_trspVia - Proposed¶284

Description¶	284
Signatures¶	284
One Via¶	284
Parameters¶	285
Optional parameters¶	285
Via optional parameters¶	285
Inner Queries¶	285
Edges SQL¶	285
Restrictions SQL¶	285
Result columns¶	286
Additional Examples¶	286
The main query¶	286
Aggregate cost of the third path.¶	286
Route's aggregate cost of the route at the end of the third path.¶	286
Nodes visited in the route.¶	286
The aggregate costs of the route when the visited vertices are reached.¶	287
Status of "passes in front" or "visits" of the nodes.¶	287
Simulation of how algorithm works.¶	287
See Also¶	288

##### pgr\_trsp\_withPoints - Proposed¶288

Description¶	289
Signatures¶	289
One to One¶	289
One to Many¶	289
Many to One¶	289
Many to Many¶	290
Combinations¶	290
Parameters¶	290
Optional parameters¶	291
With points optional parameters¶	291
Inner Queries¶	291
Edges SQL¶	291
Restrictions SQL¶	291
Points SQL¶	292
Combinations SQL¶	292
Result columns¶	292
Additional Examples¶	292
Use pgr_findCloseEdges for points on the fly¶	293
Pass in front or visits.¶	293
Show details on undirected graph.¶	293
See Also¶	293

##### pgr\_trspVia\_withPoints - Proposed¶294

Description¶	294
Signatures¶	294
One Via¶	294
Parameters¶	294
Optional parameters¶	295
Via optional parameters¶	295
With points optional parameters¶	295
Inner Queries¶	295
Edges SQL¶	295
Restrictions SQL¶	295
Points SQL¶	296
Result columns¶	296
Additional Examples¶	296
Use pgr_findCloseEdges for points on the fly¶	297
Usage variations¶	297
Aggregate cost of the third path.¶	297
Route's aggregate cost of the route at the end of the third path.¶	297
Nodes visited in the route.¶	297
The aggregate costs of the route when the visited vertices are reached.¶	297
Status of "passes in front" or "visits" of the nodes and points.¶	298
Simulation of how algorithm works.¶	298
See Also¶	299

##### pgr\_turnRestrictedPath - Experimental¶299

Description¶	300
Signatures¶	300
Parameters¶	300
Optional parameters¶	300
KSP Optional parameters¶	300
Special optional parameters¶	301
Inner Queries¶	301
Edges SQL¶	301
Restrictions SQL¶	301
Result columns¶	301
Additional Examples¶	302
See Also¶	302
Introduction¶	302
TRSP algorithm¶	302
Parameters¶	302
Restrictions¶	303
Edges SQL¶	303
Restrictions SQL¶	303
See Also¶	304

#### Traversal - Family of functions¶304

##### pgr\_depthFirstSearch - Proposed¶305

Description¶	305
Signatures¶	305
Single vertex¶	305
Multiple vertices¶	305
Parameters¶	306
Optional parameters¶	306
DFS optional parameters¶	306
Inner Queries¶	306

Edges SQL ¶	306
Result columns ¶	307
Additional Examples ¶	307
See Also ¶	307
pgr_breadthFirstSearch - Experimental ¶	308
Description ¶	308
Signatures ¶	308
Single vertex ¶	308
Multiple vertices ¶	308
Parameters ¶	309
Optional parameters ¶	309
DFS optional parameters ¶	309
Inner Queries ¶	309
Edges SQL ¶	309
Result columns ¶	310
Additional Examples ¶	310
See Also ¶	310
pgr_binaryBreadthFirstSearch - Experimental ¶	311
Description ¶	311
Signatures ¶	311
One to One ¶	311
One to Many ¶	312
Many to One ¶	312
Many to Many ¶	312
Combinations ¶	312
Parameters ¶	313
Optional parameters ¶	313
Inner Queries ¶	313
Edges SQL ¶	313
Combinations SQL ¶	313
Result columns ¶	314
Additional Examples ¶	314
See Also ¶	314
See Also ¶	314
Coloring - Family of functions ¶	315
pgr_sequentialVertexColoring - Proposed ¶	315
Description ¶	315
Signatures ¶	316
Parameters ¶	316
Inner Queries ¶	316
Edges SQL ¶	316
Result columns ¶	316
See Also ¶	317
pgr_bipartite - Experimental ¶	317
Description ¶	317
Signatures ¶	317
Parameters ¶	318
Inner Queries ¶	318
Edges SQL ¶	318
Result columns ¶	318
Additional Example ¶	318
See Also ¶	319
pgr_edgeColoring - Experimental ¶	319
Description ¶	319
Signatures ¶	319
Parameters ¶	320
Inner Queries ¶	320
Edges SQL ¶	320
Result columns ¶	320
See Also ¶	320
Result columns ¶	320
See Also ¶	321
withPoints - Family of functions ¶	321
pgr_withPoints - Proposed ¶	321
Description ¶	322
Signatures ¶	322
One to One ¶	322
One to Many ¶	322
Many to One ¶	323
Many to Many ¶	323
Combinations ¶	323
Parameters ¶	323
Optional parameters ¶	324
With points optional parameters ¶	324
Inner Queries ¶	324
Edges SQL ¶	324
Points SQL ¶	324
Combinations SQL ¶	325
Result columns ¶	325
Additional Examples ¶	325
Use pgr_findCloseEdges in the Points SQL ¶	326
Usage variations ¶	326
Passes in front or visits with right side driving ¶	326
Passes in front or visits with left side driving ¶	327
See Also ¶	327
pgr_withPointsCost - Proposed ¶	327
Description ¶	327
Signatures ¶	328
One to One ¶	328
One to Many ¶	328
Many to One ¶	328
Many to Many ¶	329
Combinations ¶	329
Parameters ¶	329
Optional parameters ¶	329
With points optional parameters ¶	329
Inner Queries ¶	330
Edges SQL ¶	330
Points SQL ¶	330
Combinations SQL ¶	330
Result columns ¶	330
Additional Examples ¶	331
Use pgr_findCloseEdges in the Points SQL ¶	331
Right side driving topology ¶	331
Left side driving topology ¶	331
Does not matter driving side driving topology ¶	331
See Also ¶	332
pgr_withPointsCostMatrix - proposed ¶	332
Description ¶	332
Signatures ¶	332
Parameters ¶	333
Optional parameters ¶	333
With points optional parameters ¶	333
Inner Queries ¶	333
Edges SQL ¶	333
Points SQL ¶	333
Result columns ¶	334
Additional Examples ¶	334
Use pgr_findCloseEdges in the Points SQL ¶	334
Use with pgr_TSP ¶	334
See Also ¶	335
pgr_withPointsKSP - Proposed ¶	335
Description ¶	335
Signatures ¶	335
One to One ¶	335
One to Many ¶	336
Many to One ¶	336
Many to Many ¶	336
Combinations ¶	337
Parameters ¶	337
Optional parameters ¶	337
KSP Optional parameters ¶	337
withPointsKSP optional parameters ¶	337
Inner Queries ¶	338
Edges SQL ¶	338
Points SQL ¶	338
Combinations SQL ¶	338
Result columns ¶	338
Additional Examples ¶	339
Use pgr_findCloseEdges in the Points SQL ¶	339
Left driving side ¶	339

Right driving side¶	339
See Also¶	340
pgr_withPointsDD - Proposed¶	340
Description¶	340
Signatures¶	340
Single vertex¶	340
Multiple vertices¶	341
Parameters¶	341
Optional parameters¶	341
With points optional parameters¶	341
Driving distance optional parameters¶	342
Inner Queries¶	342
Edges SQL¶	342
Points SQL¶	342
Result columns¶	342
Additional Examples¶	343
Use pgr_findCloseEdges in the Points SQL¶	343
Driving side does not matter¶	343
See Also¶	343
pgr_withPointsVia - Proposed¶	344
Description¶	344
Signatures¶	344
One Via¶	344
Parameters¶	344
Optional parameters¶	345
Via optional parameters¶	345
With points optional parameters¶	345
Inner Queries¶	345
Edges SQL¶	345
Points SQL¶	345
Result columns¶	346
Additional Examples¶	346
Use pgr_findCloseEdges in the Points SQL¶	346
Usage variations¶	347
Aggregate cost of the third path.¶	347
Route's aggregate cost of the route at the end of the third path.¶	347
Nodes visited in the route.¶	347
The aggregate costs of the route when the visited vertices are reached.¶	347
Status of "passes in front" or "visits" of the nodes and points.¶	347
See Also¶	347
Introduction¶	348
Parameters¶	348
Optional parameters¶	348
With points optional parameters¶	348
Inner Queries¶	348
Edges SQL¶	348
Points SQL¶	349
Combinations SQL¶	349
Advanced Documentation¶	349
About points¶	349
Driving side¶	350
Right driving side¶	350
Left driving side¶	350
Driving side does not matter¶	351
Creating temporary vertices¶	351
On a right hand side driving network¶	351
On a left hand side driving network¶	352
When driving side does not matter¶	352
See Also¶	353
See Also¶	353
Experimental Functions¶	353
Chinese Postman Problem - Family of functions (Experimental)¶	354
pgr_chinesePostman - Experimental¶	354
Description¶	354
Signatures¶	355
Parameters¶	355
Inner Queries¶	355
Edges SQL¶	355
Result columns¶	355
See Also¶	356
pgr_chinesePostmanCost - Experimental¶	356
Description¶	356
Signatures¶	356
Parameters¶	356
Inner Queries¶	357
Edges SQL¶	357
Result columns¶	357
See Also¶	357
Description¶	357
Parameters¶	358
Inner Queries¶	358
Edges SQL¶	358
See Also¶	358
Transformation - Family of functions¶	358
pgr_lineGraph - Proposed¶	359
Description¶	359
Signatures¶	359
Parameters¶	359
Optional parameters¶	359
Inner Queries¶	360
Edges SQL¶	360
Result columns¶	360
Additional Examples¶	360
Representation as directed with shared edge identifiers¶	360
Line Graph of a directed graph represented with shared edges¶	361
Representation as directed with unique edge identifiers¶	361
Line Graph of a directed graph represented with unique edges¶	361
See Also¶	362
pgr_lineGraphFull - Experimental¶	362
Description¶	362
Signatures¶	363
Parameters¶	363
Inner Queries¶	363
Edges SQL¶	363
Result columns¶	363
Additional Examples¶	364
The data¶	364
The transformation¶	364
Creating table that identifies transformed vertices¶	365
Store edge results¶	365
Create the mapping table¶	365
Filling the mapping table¶	365
Adding a soft restriction¶	366
Identifying the restriction¶	366
Adding a value to the restriction¶	366
Simplifying leaf vertices¶	366
Using the vertex map give the leaf vertices their original value.¶	367
Removing self loops on leaf nodes¶	367
Complete routing graph¶	367
Add edges from the original graph¶	367
Add the newly calculated edges¶	367
Using the routing graph¶	367
See Also¶	368
Introduction¶	368
See Also¶	368
Ordering - Family of functions¶	368
pgr_cuthillMckeeOrdering - Experimental¶	368
Description¶	369
Signatures¶	369
Parameters¶	369
Inner Queries¶	369
Edges SQL¶	369
Result columns¶	370
See Also¶	370
pgr_topologicalSort - Experimental¶	370
Description¶	370
Signatures¶	371
Parameters¶	371

Inner Queries¶	371
Edges SQL¶	371
Result columns¶	371
Additional examples¶	372
See Also¶	372
See Also¶	372
pgr_bellmanFord - Experimental¶	372
Description¶	373
Signatures¶	373
One to One¶	373
One to Many¶	373
Many to One¶	374
Many to Many¶	374
Combinations¶	374
Parameters¶	374
Optional parameters¶	375
Inner Queries¶	375
Edges SQL¶	375
Combinations SQL¶	375
Result columns¶	375
Additional Examples¶	376
See Also¶	376
pgr_dagShortestPath - Experimental¶	377
Description¶	377
Signatures¶	377
One to One¶	377
One to Many¶	378
Many to One¶	378
Many to Many¶	378
Combinations¶	378
Parameters¶	379
Inner Queries¶	379
Edges SQL¶	379
Combinations SQL¶	379
Return columns¶	379
Additional Examples¶	380
See Also¶	380
pgr_edwardMoore - Experimental¶	380
Description¶	381
Signatures¶	381
One to One¶	381
One to Many¶	382
Many to One¶	382
Many to Many¶	382
Combinations¶	382
Parameters¶	383
Optional parameters¶	383
Inner Queries¶	383
Edges SQL¶	383
Combinations SQL¶	383
Result columns¶	383
Additional Examples¶	384
See Also¶	385
pgr_isPlanar - Experimental¶	385
Description¶	385
Signatures¶	385
Parameters¶	385
Inner Queries¶	386
Edges SQL¶	386
Result columns¶	386
Additional Examples¶	386
See Also¶	386
pgr_jengauerTarjanDominatorTree - Experimental¶	387
Description¶	387
Signatures¶	387
Parameters¶	388
Inner Queries¶	388
Edges SQL¶	388
Result columns¶	388
Additional Examples¶	388
See Also¶	388
pgr_stoerWagner - Experimental¶	388
Description¶	389
Signatures¶	389
Parameters¶	389
Inner Queries¶	389
Edges SQL¶	389
Result columns¶	390
Additional Example¶	390
See Also¶	390
pgr_transitiveClosure - Experimental¶	390
Description¶	391
Signatures¶	391
Parameters¶	391
Inner Queries¶	391
Edges SQL¶	391
Result columns¶	392
See Also¶	392
pgr_hawickCircuits - Experimental¶	392
Description¶	392
Signatures¶	393
Parameters¶	394
Optional parameters¶	394
Inner Queries¶	394
Edges SQL¶	394
Result columns¶	394
See Also¶	394
See Also¶	394
Release Notes¶	395
pgRouting 3.8.0 Release Notes¶	395
All releases¶	395
Release Notes¶	395
pgRouting 3¶	396
pgRouting 3.8¶	396
pgRouting 3.8.0 Release Notes¶	396
pgRouting 3.7¶	397
pgRouting 3.7.3 Release Notes¶	397
pgRouting 3.7.2 Release Notes¶	397
pgRouting 3.7.1 Release Notes¶	397
pgRouting 3.7.0 Release Notes¶	397
pgRouting 3.6¶	398
pgRouting 3.6.3 Release Notes¶	398
pgRouting 3.6.2 Release Notes¶	398
pgRouting 3.6.1 Release Notes¶	398
pgRouting 3.6.0 Release Notes¶	398
pgRouting 3.5¶	399
pgRouting 3.5.1 Release Notes¶	399
pgRouting 3.5.0 Release Notes¶	400
pgRouting 3.4¶	400
pgRouting 3.4.2 Release Notes¶	400
pgRouting 3.4.1 Release Notes¶	400
pgRouting 3.4.0 Release Notes¶	400
pgRouting 3.3¶	401
pgRouting 3.3.5 Release Notes¶	401
pgRouting 3.3.4 Release Notes¶	401

pgRouting 3.3.3 Release Notes¶	401
pgRouting 3.3.2 Release Notes¶	401
pgRouting 3.3.1 Release Notes¶	402
pgRouting 3.3.0 Release Notes¶	402
pgRouting 3.2¶	402
pgRouting 3.2.2 Release Notes¶	402
pgRouting 3.2.1 Release Notes¶	402
pgRouting 3.2.0 Release Notes¶	402
pgRouting 3.1¶	403
pgRouting 3.1.4 Release Notes¶	403
pgRouting 3.1.3 Release Notes¶	403
pgRouting 3.1.2 Release Notes¶	404
pgRouting 3.1.1 Release Notes¶	404
pgRouting 3.1.0 Release Notes¶	404
pgRouting 3.0¶	404
pgRouting 3.0.6 Release Notes¶	404
pgRouting 3.0.5 Release Notes¶	404
pgRouting 3.0.4 Release Notes¶	404
pgRouting 3.0.3 Release Notes¶	404
pgRouting 3.0.2 Release Notes¶	405
pgRouting 3.0.1 Release Notes¶	405
pgRouting 3.0.0 Release Notes¶	405
pgRouting 2¶	407
pgRouting 2.6¶	407
pgRouting 2.6.3 Release Notes¶	407
pgRouting 2.6.2 Release Notes¶	407
pgRouting 2.6.1 Release Notes¶	407
pgRouting 2.6.0 Release Notes¶	408
pgRouting 2.5¶	408
pgRouting 2.5.5 Release Notes¶	408
pgRouting 2.5.4 Release Notes¶	408
pgRouting 2.5.3 Release Notes¶	409
pgRouting 2.5.2 Release Notes¶	409
pgRouting 2.5.1 Release Notes¶	409
pgRouting 2.5.0 Release Notes¶	409
pgRouting 2.4¶	410
pgRouting 2.4.2 Release Notes¶	410
pgRouting 2.4.1 Release Notes¶	410
pgRouting 2.4.0 Release Notes¶	410
pgRouting 2.3¶	410
pgRouting 2.3.2 Release Notes¶	410
pgRouting 2.3.1 Release Notes¶	410
pgRouting 2.3.0 Release Notes¶	410
pgRouting 2.2¶	411
pgRouting 2.2.4 Release Notes¶	411
pgRouting 2.2.3 Release Notes¶	411
pgRouting 2.2.2 Release Notes¶	411
pgRouting 2.2.1 Release Notes¶	411
pgRouting 2.2.0 Release Notes¶	412
pgRouting 2.1¶	412
pgRouting 2.1.0 Release Notes¶	412
pgRouting 2.0¶	413
pgRouting 2.0.1 Release Notes¶	413
pgRouting 2.0.0 Release Notes¶	413
pgRouting 1¶	414
pgRouting 1.0¶	414
Changes for release 1.05¶	414
Changes for release 1.03¶	414
Changes for release 1.02¶	414
Changes for release 1.01¶	414
Changes for release 1.0¶	414
Changes for release 1.0.0b¶	414
Changes for release 1.0.0a¶	414
Changes for release 0.9.9¶	414
Changes for release 0.9.6¶	414
Migration guide¶	414
Migration of pgr_alphaShape¶	415
Migration of pgr_nodeNetwork¶	415
Migration of pgr_createTopology¶	415
Build a routing topology¶	415
Migration of pgr_createVerticesTable¶	416
Migration of pgr_analyzeOneWay¶	416
Migration of pgr_analyzeGraph¶	416
Migration of pgr_aStar¶	417
Migration of pgr_bdAstar¶	418
Migration of pgr_dijkstra¶	419
Migration of pgr_drivingDistance¶	420
pgr_drivingDistance(Single vertex)¶	420
pgr_drivingDistance(Multiple vertices)¶	421
Migration of pgr_kruskalBD / pgr_kruskalBFS / pgr_kruskalDFS¶	421
Kruskal single vertex¶	422
Kruskal multiple vertices¶	422
Migration of pgr_KSP¶	422
pgr_KSP (One to One)¶	423
Migration of pgr_maxCardinalityMatch¶	423
Migration of pgr_primDD / pgr_primBFS / pgr_primDFS¶	424
Prim single vertex¶	424
Prim multiple vertices¶	424
Migration of pgr_withPointsDD¶	425
pgr_withPointsDD (Single vertex)¶	426
pgr_withPointsDD (Multiple vertices)¶	426
Migration of pgr_withPointsKSP¶	427
pgr_withPointsKSP (One to One)¶	427
Migration of pgr_trsp (Vertices)¶	428
Use pgr_dijkstra when there are no restrictions.¶	428
Use pgr_trsp when there are restrictions.¶	428
Migration of pgr_trsp (Edges)¶	429
Use pgr_withPoints when there are no restrictions.¶	429
Use pgr_trsp_withPoints when there are restrictions.¶	429
Migration of pgr_trspViaVertices¶	429
Use pgr_dijkstraVia when there are no restrictions¶	430
Use pgr_trspVia when there are restrictions¶	430
Migration of pgr_trspViaEdges¶	430
Use pgr_withPointsVia when there are no restrictions¶	431
Use pgr_trspVia_withPoints when there are restrictions¶	431
Migration of restrictions¶	432
Old restrictions structure¶	432
Old restrictions contents¶	432
New restrictions structure¶	432
Restrictions data¶	433
Migration¶	433
See Also¶	433

# pgRouting Manual (3.8)

## pgRouting Manual (3.8)

[Contents](#)

### Table of Contents¶

pgRouting extends the [PostGIS/PostgreSQL](#) geospatial database to provide geospatial routing and other network analysis functionality.

This is the manual for pgRouting v3.8.0.

The pgRouting Manual is licensed under a [Creative Commons Attribution-Share Alike 3.0 License](#). Feel free to use this material any way you like, but we ask that you attribute credit to the pgRouting Project and wherever possible, a link back to <https://pgrouting.org>. For other licenses used in pgRouting see the [Licensing](#) page.

### General¶

#### Introduction¶

pgRouting is an extension of [PostGIS](#) and [PostgreSQL](#) geospatial database and adds routing and other network analysis functionality. A predecessor of pgRouting – pgDijkstra, written by Sylvain Pasche from [Camptocamp](#), was later extended by Orkney and renamed to pgRouting. The project is now supported and maintained by [Georepublic](#), [Paragon Corporation](#) and a broad user community.

pgRouting is part of [OSGeo Community Projects](#) from the [OSGeo Foundation](#) and included on [OSGeoLive](#).

#### Licensing¶

The following licenses can be found in pgRouting:

##### License

GNU General Public License v2.0 or later	Most features of pgRouting are available under <a href="#">GNU General Public License v2.0 or later</a> .
Boost Software License - Version 1.0	Some Boost extensions are available under <a href="#">Boost Software License - Version 1.0</a>
MIT-X License	Some code contributed by iMaptools.com is available under MIT-X license.
Creative Commons Attribution-Share Alike 3.0 License	The pgRouting Manual is licensed under a <a href="#">Creative Commons Attribution-Share Alike 3.0 License</a> .

In general license information should be included in the header of each source file.

#### Contributors¶

##### This Release Contributors¶

Individuals in this release v3.8.x (in alphabetical order)¶

Aur lie Bousquet, Regina Obe, Vicky Vergara

And all the people that give us a little of their time making comments, finding issues, making pull requests etc. in any of our products: osm2pgrouting, pgRouting, pgRoutingLayer, workshop.

Translators (in alphabetical order)¶

Dapeng Wang

Corporate Sponsors in this release (in alphabetical order)¶

These are corporate entities that have contributed developer time, hosting, or direct monetary funding to the pgRouting project:

- [OSGeo](#)
- [OSGeo UK](#)
- [Google Summer of Code](#)
- [HighGo Software](#)
- [Paragon Corporation](#)

##### Contributors Past & Present:¶

Individuals (in alphabetical order)¶

Aasheesh Tiwari, Abhinav Jain, Aditya Pratap Singh, Adrien Berchet, Akio Takubo, Andrea Nardelli, Anthony Tasca, Anton Patrushev, Aryan Gupta, Ashraf Hossain, Ashish Kumar, Aur lie Bousquet, Cayetano Benavent, Christian Gonzalez, Daniel Kastl, Dapeng Wang, Dave Potts, David Techer, Denis Rykov, Ema Miyawaki, Esteban Zimanyi, Florian Therkow, Frederic Junod, Gerald Fenoy, Gudes Venkata Sai Akhil, Hang Wu, Himanshu Raj, Imre Samu, Jay Mahadeokar, Jinfu Leng, Kai Behncke, Kishore Kumar, Ko Nagase, Mahmoud Sakr, Manikata Kondeti, Mario Basa, Martin Wiesenhaan, Maxim Dubinin, Maoguang Wang, Mohamed Bakli, Mohamed Zia, Mukul Priya, Nitish Chauhan, Rajat Shinde, Razequl Islam, Regina Obe, Rohith Reddy, Sarthak Agarwal, Shobhit Chaurasia, Sourabh Garg, Stephen Woodbridge, Swapnil Joshi, Sylvain Housseman, Sylvain Pasche, Veenit Kumar, Vidhan Jain, Virginia Vergara, Yige Huang

Corporate Sponsors (in alphabetical order)¶

These are corporate entities that have contributed developer time, hosting, or direct monetary funding to the pgRouting project:

- Camptocamp
- CSIS (University of Tokyo)
- Georepublic
- Google Summer of Code
- HighGo Software
- iMaptools
- Leopark

- Orkney
- OSGeo
- OSGeo UK
- Paragon Corporation
- Versatarm Inc.

## More Information¶

- The latest software, documentation and news items are available at the pgRouting web site <https://pgrouting.org>.
- PostgreSQL database server at the PostgreSQL main site <https://www.postgresql.org>.
- PostGIS extension at the PostGIS project web site <https://postgis.net>.
- Boost C++ source libraries at <https://www.boost.org>.
- [Migration guide](#)

## Installation¶

### Table of Contents

- [Short Version](#)
- [Get the sources](#)
- [Enabling and upgrading in the database](#)
- [Dependencies](#)
- [Configuring](#)
- [Building](#)
- [Testing](#)

Instructions for downloading and installing binaries for different operating systems, additional notes and corrections not included in this documentation can be found in [installation wiki](#)

To use pgRouting PostGIS needs to be installed, please read the information about installation in this [install Guide](#)

## Short Version¶

### Extracting the tar ball

```
tar xvfz pgrouting-3.8.0.tar.gz
cd pgrouting-3.8.0
```

To compile assuming you have all the dependencies in your search path:

```
mkdir build
cd build
cmake ..
make
sudo make install
```

Once pgRouting is installed, it needs to be enabled in each individual database you want to use it in.

```
createdb routing
psql routing -c 'CREATE EXTENSION PostGIS'
psql routing -c 'CREATE EXTENSION pgRouting'
```

## Get the sources¶

The pgRouting latest release can be found in <https://github.com/pgRouting/pgrouting/releases/latest>

To download this release:

```
wget -O pgrouting-3.8.0.tar.gz https://github.com/pgRouting/pgrouting/archive/v3.8.0.tar.gz
```

Go to [Short Version](#) for more instructions on extracting tar ball and compiling pgRouting.

git

To download the repository

```
git clone git://github.com/pgRouting/pgrouting.git
cd pgrouting
git checkout v3.8.0
```

Go to [Short Version](#) for more instructions on compiling pgRouting (there is no tar ball involved while downloading pgRouting repository from GitHub).

## Enabling and upgrading in the database¶

### Enabling the database

pgRouting is a PostgreSQL extension and depends on PostGIS to provide functionalities to end user. Below given code demonstrates enabling PostGIS and pgRouting in the database.

```
CREATE EXTENSION postgis;
CREATE EXTENSION pgrouting;
```

Checking PostGIS and pgRouting version after enabling them in the database.

```
SELECT PostGIS_full_version();
SELECT * FROM pgr_version();
```

### Upgrading the database

To upgrade pgRouting in the database to version 3.8.0 use the following command:

```
ALTER EXTENSION pgrouting UPDATE TO "3.8.0";
```

More information can be found in <https://www.postgresql.org/docs/current/sql-createextension.html>

## Dependencies¶

### Compilation Dependencies

To be able to compile pgRouting, make sure that the following dependencies are met:

- C and C++0x compilers
  - Compiling with Boost 1.56 up to Boost 1.74 requires C++ Compiler with C++03 or C++11 standard support

- Compiling with Boost 1.75 requires C++ Compiler with C++14 standard support

- Postgresql version = Supported versions by PostgreSQL
- The Boost Graph Library (BGL). Version >= 1.56
- CMake >= 3.2

optional dependencies

For user's documentation

- Sphinx >= 1.1
- Latex

For developer's documentation

- Doxygen >= 1.7

For testing

- pgtap
- pg\_prove

For using:

- PostGIS version >= 2.2

Example: Installing dependencies on linux

Installing the compilation dependencies

Database dependencies

```
sudo apt install postgresql-15
sudo apt install postgresql-server-dev-15
sudo apt install postgresql-15-postgis
```

Configuring PostgreSQL

Entering psql console

```
sudo systemctl start postgresql.service
sudo -i -u postgres
psql
```

To exit psql console

```
q
```

Entering psql console directly without switching roles can be done by the following commands

```
sudo -u postgres psql
```

Then use the above given method to exit out of the psql console

Checking PostgreSQL version

```
psql --version
```

or

Enter the psql console using above given method and then enter

```
SELECT VERSION();
```

Creating PostgreSQL role

```
sudo -i -u postgres
createuser --interactive
```

or

```
sudo -u postgres createuser --interactive
```

Default role provided by PostgreSQL is postgres. To create new roles you can use the above provided commands. The prompt will ask the user to type name of the role and then provide affirmation. Proceed with the steps and you will succeed in creating PostgreSQL role successfully.

To add password to the role or change previously created password of the role use the following commands

```
ALTER USER <role name> PASSWORD <password>
```

To get additional details on the flags associated with createuser below given command can be used

```
man createuser
```

Creating Database in PostgreSQL

```
sudo -i -u postgres
createdb <database name>
```

or

```
sudo -u postgres createdb <database name>
```

Connecting to a PostgreSQL Database

Enter the psql console and type the following commands

```
connect <database name>
```

Build dependencies

```
sudo apt install cmake
sudo apt install g++
sudo apt install libboost-graph-dev
```

Optional dependencies

For documentation and testing

```
pip install sphinx
pip install sphinx-bootstrap-theme
sudo apt install texlive
sudo apt install doxygen
sudo apt install libtap-parser-sourcehandler-pgtap-perl
sudo apt install postgresql-15-pgtap
```

**Configuring**



pgRouting uses the *cmake* system to do the configuration.

The build directory is different from the source directory

Create the build directory

```
$ mkdir build
```

**Configurable variables**

To see the variables that can be configured

```
$ cd build
$ cmake -L ..
```

Configuring The Documentation

Most of the effort of the documentation has been on the HTML files. Some variables for building documentation:

Variable	Default	Comment
WITH_DOC	BOOL=OFF	Turn on/off building the documentation
BUILD_HTML	BOOL=ON	If ON, turn on/off building HTML for user's documentation
BUILD_DOXY	BOOL=ON	If ON, turn on/off building HTML for developer's documentation
BUILD_LATEX	BOOL=OFF	If ON, turn on/off building PDF
BUILD_MAN	BOOL=OFF	If ON, turn on/off building MAN pages
DOC_USE_BOOTSTRAP	BOOL=OFF	If ON, use sphinx-bootstrap for HTML pages of the users documentation

Configuring cmake to create documentation before building pgRouting

```
$ cmake -DWITH_DOC=ON -DDOC_USE_BOOTSTRAP=ON ..
```

**Note**

Most of the effort of the documentation has been on the html files.

**Building**

Using make to build the code and the documentation

The following instructions start from *path/to/pgrouting/build*

```
$ make      # build the code but not the documentation
$ make doc  # build only the user's documentation
$ make all doc # build both the code and the user's documentation
$ make doxy # build only the developer's documentation
```

We have tested on several platforms, For installing or reinstalling all the steps are needed.

**Warning**

The sql signatures are configured and build in the *cmake* command.

**MinGW on Windows**

```
$ mkdir build
$ cd build
$ cmake -G"MSYS Makefiles" ..
$ make
$ make install
```

**Linux**

The following instructions start from *path/to/pgrouting*

```
mkdir build
cd build
cmake ..
make
sudo make install
```

To remove the build when the configuration changes, use the following code:

```
rm -rf build
```

and start the build process as mentioned previously.

**Testing**

Currently there is no make test and testing is done as follows

The following instructions start from *path/to/pgrouting/*

```
tools/testers/doc_queries_generator.pl
createdb -U <user> __pgr__test__
sh ./tools/testers/pg_prove_tests.sh <user>
dropdb -U <user> __pgr__test__
```

**See Also**

Indices and tables

- [Index](#)
- [Search Page](#)

**Support**

pgRouting community support is available through the [pgRouting website](#), [documentation](#), tutorials, mailing lists and others. If you're looking for [commercial support](#), find below a list of companies providing pgRouting development and consulting services.

## Reporting Problems¶

Bugs are reported and managed in an [issue tracker](#). Please follow these steps:

1. Search the tickets to see if your problem has already been reported. If so, add any extra context you might have found, or at least indicate that you too are having the problem. This will help us prioritize common issues.
2. If your problem is unreported, create a [new issue](#) for it.
3. In your report include explicit instructions to replicate your issue. The best tickets include the exact SQL necessary to replicate a problem.
4. If you can test older versions of PostGIS for your problem, please do. On your ticket, note the earliest version the problem appears.
5. For the versions where you can replicate the problem, note the operating system and version of pgRouting, PostGIS and PostgreSQL.
6. It is recommended to use the following wrapper on the problem to pin point the step that is causing the problem.

```
SET client_min_messages TO debug;  
<your code>  
SET client_min_messages TO notice;
```

## Mailing List and GIS StackExchange¶

There are two mailing lists for pgRouting hosted on OSGeo mailing list server:

- User mailing list: <https://lists.osgeo.org/mailman/listinfo/pgrouting-users>
- Developer mailing list: <https://discourse.osgeo.org/c/pgrouting/pgrouting-dev/>
  - Subscribe: <https://discourse.osgeo.org/g/pgrouting-dev>

For general questions and topics about how to use pgRouting, please write to the user mailing list.

You can also ask at [GIS StackExchange](#) and tag the question with `pgrouting`. Find all questions tagged with `pgrouting` under <https://gis.stackexchange.com/questions/tagged/pgrouting> or subscribe to the [pgRouting questions feed](#).

## Commercial Support¶

For users who require professional support, development and consulting services, consider contacting any of the following organizations, which have significantly contributed to the development of pgRouting:

Company	Offices in	Website
Georepublic	Germany, Japan	<a href="https://georepublic.info">https://georepublic.info</a>
Paragon Corporation	United States	<a href="https://www.paragoncorporation.com">https://www.paragoncorporation.com</a>
Netlab	Capranica, Italy	<a href="https://www.osgeo.org/service-providers/netlab/">https://www.osgeo.org/service-providers/netlab/</a>

- [Sample Data](#) that is used in the examples of this manual.

## Sample Data¶

The documentation provides very simple example queries based on a small sample network that resembles a city. To be able to execute the majority of the examples queries, follow the instructions below.

- [Main graph](#)
  - [Edges](#)
    - [Edges data](#)
  - [Vertices](#)
    - [Vertices data](#)
  - [The topology](#)
    - [Topology data](#)
  - [Points outside the graph](#)
    - [Points of interest](#)
    - [Points of interest fill up](#)
- [Support tables](#)
  - [Combinations](#)
    - [Combinations data](#)
  - [Restrictions](#)
    - [Restrictions data](#)
- [Images](#)
  - [Directed graph with cost and reverse\\_cost](#)
  - [Undirected graph with cost and reverse\\_cost](#)
  - [Directed graph with cost](#)
  - [Undirected graph with cost](#)
- [Pick & Deliver Data](#)
  - [The vehicles](#)
  - [The original orders](#)
  - [The orders](#)

## Main graph¶

A graph consists of a set of edges and a set of vertices.

The following city is to be inserted into the database:

Information known at this point is the geometry of the edges, cost values, capacity values, category values and some locations that are not in the graph.

The process to have working topology starts by inserting the edges. After that everything else is calculated.

Edges¶

The database design for the documentation of pgRouting, keeps in the same row 2 segments, one in the direction of the geometry and the second in the opposite direction. Therefore some information has the `reverse_` prefix which corresponds to the segment on the opposite direction of the geometry.

Column	Description
id	A unique identifier.
source	Identifier of the starting vertex of the geometrygeom.
target	Identifier of the ending vertex of the geometrygeom
cost	Cost to traverse from <i>source</i> to target.
reverse_cost	Cost to traverse from <i>target</i> to source.
capacity	Flow capacity from <i>source</i> to target.
reverse_capacity	Flow capacity from <i>target</i> to source.
category	Flow capacity from <i>target</i> to source.
reverse_category	Flow capacity from <i>target</i> to source.
x1	<div><div>\(x\) coordinate of the starting vertex of the geometry.</div><div><ul style="list-style-type: none"><li>For convenience it is saved on the table but can be calculated as <code>ST_X(ST_StartPoint(geom))</code>.</li></ul></div></div>
y2	<div><div>\(y\) coordinate of the ending vertex of the geometry.</div><div><ul style="list-style-type: none"><li>For convenience it is saved on the table but can be calculated as <code>ST_Y(ST_EndPoint(geom))</code>.</li></ul></div></div>
geom	The geometry of the segments.

```
CREATE TABLE edges (  
  id BIGSERIAL PRIMARY KEY,  
  source BIGINT,  
  target BIGINT,  
  cost FLOAT,  
  reverse_cost FLOAT,  
  capacity BIGINT,  
  reverse_capacity BIGINT,  
  x1 FLOAT,  
  y1 FLOAT,  
  x2 FLOAT,  
  y2 FLOAT,  
  geom geometry  
);  
CREATE TABLE
```

Starting on PostgreSQL 12:

```
...  
x1 FLOAT GENERATED ALWAYS AS (ST_X(ST_StartPoint(geom))) STORED,  
y1 FLOAT GENERATED ALWAYS AS (ST_Y(ST_StartPoint(geom))) STORED,  
x1 FLOAT GENERATED ALWAYS AS (ST_X(ST_EndPoint(geom))) STORED,  
y1 FLOAT GENERATED ALWAYS AS (ST_Y(ST_EndPoint(geom))) STORED,  
...
```

Optionally indexes on different columns can be created. The recommendation is to have

- id indexed.
- source and target columns indexed to speed up pgRouting queries.
- geom indexed to speed up geometry processes that might be needed in the front end.

For this small example the indexes are skipped, except for id

Edges data¶

Inserting into the database the information of the edges:

```
INSERT INTO edges (  
  cost, reverse_cost,  
  capacity, reverse_capacity, geom) VALUES  
(1, 1, 80, 130, ST_MakeLine(ST_POINT(2, 0), ST_POINT(2, 1))),  
(-1, 1, -1, 100, ST_MakeLine(ST_POINT(2, 1), ST_POINT(3, 1))),  
(-1, 1, -1, 130, ST_MakeLine(ST_POINT(3, 1), ST_POINT(4, 1))),  
(1, 1, 100, 50, ST_MakeLine(ST_POINT(2, 1), ST_POINT(2, 2))),  
(1, -1, 130, -1, ST_MakeLine(ST_POINT(3, 1), ST_POINT(3, 2))),  
(1, 1, 50, 100, ST_MakeLine(ST_POINT(0, 2), ST_POINT(1, 2))),  
(1, 1, 50, 130, ST_MakeLine(ST_POINT(1, 2), ST_POINT(2, 2))),  
(1, 1, 100, 130, ST_MakeLine(ST_POINT(2, 2), ST_POINT(3, 2))),  
(1, 1, 130, 80, ST_MakeLine(ST_POINT(3, 2), ST_POINT(4, 2))),  
(1, 1, 130, 50, ST_MakeLine(ST_POINT(2, 2), ST_POINT(2, 3))),  
(1, -1, 130, -1, ST_MakeLine(ST_POINT(3, 2), ST_POINT(3, 3))),  
(1, -1, 100, -1, ST_MakeLine(ST_POINT(2, 3), ST_POINT(3, 3))),  
(1, -1, 100, -1, ST_MakeLine(ST_POINT(3, 3), ST_POINT(4, 3))),  
(1, 1, 80, 130, ST_MakeLine(ST_POINT(2, 3), ST_POINT(2, 4))),  
(1, 1, 80, 50, ST_MakeLine(ST_POINT(4, 2), ST_POINT(4, 3))),  
(1, 1, 80, 80, ST_MakeLine(ST_POINT(4, 1), ST_POINT(4, 2))),  
(1, 1, 130, 100, ST_MakeLine(ST_POINT(0.5, 3.5), ST_POINT(1.9999999999999999, 3.5))),  
(1, 1, 50, 130, ST_MakeLine(ST_POINT(3.5, 2.3), ST_POINT(3.5, 4)));
```

INSERT 0 18

Negative values on the cost, capacity and category means that the edge do not exist.

## Vertices¶

The vertex information is calculated based on the identifier of the edge and the geometry and saved on a table. Saving all the information provided by `pggr::extractVertices()`

```
SELECT * INTO vertices
FROM pgr_extractVertices('SELECT id, geom FROM edges ORDER BY id');
SELECT 17
```

In this case the because the CREATE statement was not used, the definition of an index on the table is needed

```
CREATE SEQUENCE vertices_id_seq;
CREATE SEQUENCE
ALTER TABLE vertices ALTER COLUMN id SET DEFAULT nextval('vertices_id_seq');
ALTER TABLE
ALTER SEQUENCE vertices_id_seq OWNED BY vertices.id;
ALTER SEQUENCE
SELECT setval('vertices_id_seq', (SELECT coalesce(max(id)) FROM vertices));
setval
-----
17
(1 row)
```

The structure of the table is:

Column	Type	Collation	Nullable	Default
id	bigint			nextval('vertices_id_seq')::regclass
in_edges	bigint[]			
out_edges	bigint[]			
x	double precision			
y	double precision			
geom	geometry			

### Vertices data

The saved information of the vertices is:

[illegible]

Here is where adding more columns to the vertices table can be done. Additional columns names and types will depend on the application.

## The topology

This queries based on the vertices data create a topology by filling the `source` and `target` columns in the edges table.

```
/* -- set the source information */
UPDATE edges AS e
SET source = v.id, x1 = x, y1 = y
FROM vertices AS v
WHERE ST_StartPoint(e.geom) = v.geom
UPDATE 18

/* -- set the target information */
UPDATE edges AS e
SET target = v.id, x2 = x, y2 = y
FROM vertices AS v
WHERE ST_EndPoint(e.geom) = v.geom;
UPDATE 18
```

## Topology data

```
SELECT id, source, target
FROM edges ORDER BY id;
id | source | target
```

1	5	6
2	6	10
3	10	15
4	6	7
5	10	11
6	1	3
7	3	7
8	7	11
9	11	16
10	7	8
11	11	12
12	8	12
13	12	17
14	8	9
15	16	17
16	15	16
17	2	4
18	13	14

### Points outside the graph

### Points of interest

Some times the applications work “on the fly” starting from a location that is not a vertex in the graph. Those locations, in pgRouting are called points of interest.

The information needed in the points of interest is `pid`, `edge`, `id`, `side`, `fraction`.

On this documentation there will be some 6 fixed points of interest and they will be stored on a table.

Column	Description
pid	A unique identifier.
edge_id	Identifier of the nearest segment.
side	Is it on the left, right or both sides of the segmentedge_id.
fraction	Where in the segment is the point located.
geom	The geometry of the points.
distance	The distance between geom and the segment edge_id.
edge	A segment that connects the geom of the point to the closest point on the segment edge_id.
newPoint	A point on segment edge_id that is the closest to geom.

```
CREATE TABLE pointsOfInterest(  
  pid BIGSERIAL PRIMARY KEY,  
  edge_id BIGINT,  
  side CHAR,  
  fraction FLOAT,  
  distance FLOAT,  
  edge geometry,  
  newPoint geometry,  
  geom geometry);  
IF v > 3.4 THEN
```

Points of interest fill up¶

Inserting the points of interest.

```
INSERT INTO pointsOfInterest (geom) VALUES  
(ST_Point(1.8, 0.4)),  
(ST_Point(4.2, 2.4)),  
(ST_Point(2.6, 3.2)),  
(ST_Point(0.3, 1.8)),  
(ST_Point(2.9, 1.8)),  
(ST_Point(2.2, 1.7));
```

Filling the rest of the table.

```
UPDATE pointsofinterest SET  
  edge_id = poi.edge_id,  
  side = poi.side,  
  fraction = round(poi.fraction::numeric, 2),  
  distance = round(poi.distance::numeric, 2),  
  edge = poi.edge,  
  newPoint = ST_EndPoint(poi.edge)  
FROM (  
  SELECT *  
  FROM pgr_findCloseEdges(  
    $$SELECT id, geom FROM edges$$,(SELECT array_agg(geom) FROM pointsOfInterest), 0.5) ) AS poi  
WHERE pointsOfInterest.geom = poi.geom;
```

Any other additional modification: In this manual, point\(\6\) can be reached from both sides.

```
UPDATE pointsOfInterest SET side = 'b' WHERE pid = 6;
```

The points of interest:

```
SELECT  
  pid, ST_AsText(geom) geom,  
  edge_id, fraction AS frac, side, distance AS dist,  
  ST_AsText(edge) edge, ST_AsText(newPoint) newPoint  
FROM pointsOfInterest;  
pid | geom | edge_id | frac | side | dist | edge | newpoint  
-----+-----+-----+-----+-----+-----+-----+-----  
1 | POINT(1.8 0.4) | 1 | 0.4 | l | 0.2 | LINESTRING(1.8 0.4,2 0.4) | POINT(2 0.4)  
4 | POINT(0.3 1.8) | 6 | 0.3 | r | 0.2 | LINESTRING(0.3 1.8,0.3 2) | POINT(0.3 2)  
3 | POINT(2.6 3.2) | 12 | 0.6 | l | 0.2 | LINESTRING(2.6 3.2,2.6 3) | POINT(2.6 3)  
2 | POINT(4.2 2.4) | 15 | 0.4 | r | 0.2 | LINESTRING(4.2 2.4,4 2.4) | POINT(4 2.4)  
5 | POINT(2.9 1.8) | 5 | 0.8 | l | 0.1 | LINESTRING(2.9 1.8,3 1.8) | POINT(3 1.8)  
6 | POINT(2.2 1.7) | 4 | 0.7 | b | 0.2 | LINESTRING(2.2 1.7,2 1.7) | POINT(2 1.7)  
(6 rows)
```

Support tables¶

Combinations¶

Many functions can be used with a combinations of(source, target) pairs when wanting a route fromsource to target.

For convenience of this documentation, some combinations will be stored on a table:

```
CREATE TABLE combinations (  
  source BIGINT,  
  target BIGINT  
);  
CREATE TABLE
```

Inserting the data:

```
INSERT INTO combinations (  
  source, target) VALUES  
(5, 6),  
(5, 10),  
(6, 5),  
(6, 15),  
(6, 14);  
INSERT 0 5
```

Combinations data¶

```
SELECT * FROM combinations;  
source | target  
-----+-----  
5 | 6  
5 | 10  
6 | 5
```

6		15
6		14

(5 rows)

Restrictions¶

Some functions accept soft restrictions about the segments.

The creation of the restrictions table

```
CREATE TABLE restrictions (  
  id SERIAL PRIMARY KEY,  
  path BIGINT[],  
  cost FLOAT  
);  
CREATE TABLE
```

Adding the restrictions

```
INSERT INTO restrictions (path, cost) VALUES  
(ARRAY[4, 7], 100),  
(ARRAY[8, 11], 100),  
(ARRAY[7, 10], 100),  
(ARRAY[3, 5, 9], 4),  
(ARRAY[9, 16], 100);  
INSERT 0 5
```

Restrictions data¶


```
SELECT * FROM restrictions;  
id | path | cost  
-----+-----+-----  
1 | {4,7} | 100  
2 | {8,11} | 100  
3 | {7,10} | 100  
4 | {3,5,9} | 4  
5 | {9,16} | 100  
(5 rows)
```

Images¶

- Red arrows correspond when cost > 0 in the edge table.
- Blue arrows correspond when reverse\_cost > 0 in the edge table.
- Points are outside the graph.
- Click on the graph to enlarge.

Directed graph with cost and reverse\_cost¶


When working with city networks, this is recommended for point of view of vehicles.



Directed, with cost and reverse\_cost¶

Undirected graph with cost and reverse\_cost¶

When working with city networks, this is recommended for point of view of pedestrians.



Undirected, with cost and reverse cost¶

Directed graph with cost¶

[\\_images/Fig2-cost.png](#)

Directed, with cost

[Undirected graph with cost](#)

Undirected, with cost

[Pick & Deliver Data](#)

This data example **lc101** is from data published at <https://www.sintef.no/projectweb/top/pdptw/li-lim-benchmark/>

[The vehicles](#)

There are 25 vehicles in the problem all with the same characteristics.

```
CREATE TABLE v_lc101(  
  id BIGINT NOT NULL primary key,  
  capacity BIGINT DEFAULT 200,  
  start_x FLOAT DEFAULT 30,  
  start_y FLOAT DEFAULT 50,  
  start_open INTEGER DEFAULT 0,  
  start_close INTEGER DEFAULT 1236);  
CREATE TABLE  
/* create 25 vehicles */  
INSERT INTO v_lc101 (id)  
(SELECT * FROM generate_series(1, 25));  
INSERT 0 25
```

[The original orders](#)

The data comes in different rows for the pickup and the delivery of the same order.

```
CREATE table lc101_c(  
  id BIGINT not null primary key,  
  x DOUBLE PRECISION,  
  y DOUBLE PRECISION,  
  demand INTEGER,  
  open INTEGER,  
  close INTEGER,  
  service INTEGER,  
  pindex BIGINT,  
  dindex BIGINT  
);  
CREATE TABLE  
/* the original data */  
INSERT INTO lc101_c(  
  id, x, y, demand, open, close, service, pindex, dindex) VALUES  
( 1, 45, 68, -10, 912, 967, 90, 11, 0),  
( 2, 45, 70, -20, 825, 870, 90, 6, 0),  
( 3, 42, 66, 10, 65, 146, 90, 0, 75),  
( 4, 42, 68, -10, 727, 782, 90, 9, 0),  
( 5, 42, 65, 10, 15, 67, 90, 0, 7),  
( 6, 40, 69, 20, 621, 702, 90, 0, 2),  
( 7, 40, 66, -10, 170, 225, 90, 5, 0),  
( 8, 38, 68, 20, 255, 324, 90, 0, 10),  
( 9, 38, 70, 10, 534, 605, 90, 0, 4),  
(10, 35, 66, -20, 357, 410, 90, 8, 0),  
(11, 35, 69, 10, 448, 505, 90, 0, 1),  
(12, 25, 85, -20, 652, 721, 90, 18, 0),  
(13, 22, 75, 30, 30, 92, 90, 0, 17),  
(14, 22, 85, -40, 567, 620, 90, 16, 0),  
(15, 20, 80, -10, 384, 429, 90, 19, 0),  
(16, 20, 85, 40, 475, 528, 90, 0, 14),  
(17, 18, 75, -30, 99, 148, 90, 13, 0),  
(18, 15, 75, 20, 179, 254, 90, 0, 12),  
(19, 15, 80, 10, 278, 345, 90, 0, 15),  
(20, 30, 50, 10, 10, 73, 90, 0, 24),  
(21, 30, 52, -10, 914, 965, 90, 30, 0),  
(22, 28, 52, -20, 812, 883, 90, 28, 0),  
(23, 28, 55, 10, 732, 777, 0, 0, 103),  
(24, 25, 50, -10, 65, 144, 90, 20, 0),  
(25, 25, 52, 40, 169, 224, 90, 0, 27),  
(26, 25, 55, -10, 622, 701, 90, 29, 0),  
(27, 23, 52, -40, 261, 316, 90, 25, 0),  
(28, 23, 55, 20, 546, 593, 90, 0, 22),  
(29, 20, 50, 10, 358, 405, 90, 0, 26),  
(30, 20, 55, 10, 449, 504, 90, 0, 21),  
(31, 10, 35, -30, 200, 237, 90, 32, 0),  
(32, 10, 40, 30, 31, 100, 90, 0, 31),  
(33, 8, 40, 40, 87, 158, 90, 0, 37),  
(34, 8, 45, -30, 751, 816, 90, 38, 0),  
(35, 5, 35, 10, 283, 344, 90, 0, 39),  
(36, 5, 45, 10, 665, 716, 0, 0, 105),  
(37, 2, 40, -40, 383, 434, 90, 33, 0),  
(38, 0, 40, 30, 479, 522, 90, 0, 34),  
(39, 0, 45, -10, 567, 624, 90, 35, 0),  
(40, 35, 30, -20, 264, 321, 90, 42, 0),  
(41, 35, 32, -10, 166, 235, 90, 43, 0),  
(42, 33, 32, 20, 68, 149, 90, 0, 40),  
(43, 33, 35, 10, 16, 80, 90, 0, 41),  
(44, 32, 30, 10, 359, 412, 90, 0, 46),  
(45, 30, 30, 10, 541, 600, 90, 0, 48),  
(46, 30, 32, -10, 448, 509, 90, 44, 0),  
(47, 30, 35, -10, 1054, 1127, 90, 49, 0),  
(48, 28, 30, -10, 632, 693, 90, 45, 0),  
(49, 28, 35, 10, 1001, 1066, 90, 0, 47),  
(50, 26, 32, 10, 815, 880, 90, 0, 52),  
(51, 25, 30, 10, 725, 786, 0, 0, 101),  
(52, 25, 35, -10, 912, 969, 90, 50, 0),  
(53, 44, 5, 20, 286, 347, 90, 0, 58),  
(54, 42, 10, 40, 186, 257, 90, 0, 60),  
(55, 42, 15, -40, 95, 158, 90, 57, 0),  
(56, 40, 5, 30, 385, 436, 90, 0, 59),
```

```
(57, 40, 15, 40, 35, 87, 90, 0, 55),
(58, 38, 5, -20, 471, 534, 90, 53, 0),
(59, 38, 15, -30, 651, 740, 90, 56, 0),
(60, 35, 5, -40, 562, 629, 90, 54, 0),
(61, 50, 30, -10, 531, 610, 90, 67, 0),
(62, 50, 35, 20, 262, 317, 90, 0, 68),
(63, 50, 40, 50, 171, 218, 90, 0, 74),
(64, 48, 30, 10, 632, 693, 0, 0, 102),
(65, 48, 40, 10, 76, 129, 90, 0, 72),
(66, 47, 35, 10, 826, 875, 90, 0, 69),
(67, 47, 40, 10, 12, 77, 90, 0, 61),
(68, 45, 30, -20, 734, 777, 90, 62, 0),
(69, 45, 35, -10, 916, 969, 90, 66, 0),
(70, 95, 30, -30, 387, 456, 90, 61, 0),
(71, 95, 35, 20, 293, 360, 90, 0, 77),
(72, 53, 30, -10, 450, 505, 90, 65, 0),
(73, 92, 30, -10, 478, 551, 90, 76, 0),
(74, 53, 35, -50, 353, 412, 90, 63, 0),
(75, 45, 65, -10, 997, 1068, 90, 3, 0),
(76, 90, 35, 10, 203, 260, 90, 0, 73),
(77, 88, 30, -20, 574, 643, 90, 71, 0),
(78, 88, 35, 20, 109, 170, 0, 0, 104),
(79, 87, 30, 10, 668, 731, 90, 0, 80),
(80, 85, 25, -10, 769, 820, 90, 79, 0),
(81, 85, 35, 30, 47, 124, 90, 0, 70),
(82, 75, 55, 20, 369, 420, 90, 0, 85),
(83, 72, 55, -20, 265, 338, 90, 87, 0),
(84, 70, 58, 20, 458, 523, 90, 0, 89),
(85, 68, 60, -20, 555, 612, 90, 82, 0),
(86, 66, 55, 10, 173, 238, 90, 0, 91),
(87, 65, 55, 20, 85, 144, 90, 0, 83),
(88, 65, 60, -10, 645, 708, 90, 90, 0),
(89, 63, 58, -20, 737, 802, 90, 84, 0),
(90, 60, 55, 10, 20, 84, 90, 0, 88),
(91, 60, 60, -10, 836, 889, 90, 86, 0),
(92, 67, 85, 20, 368, 441, 90, 0, 93),
(93, 65, 85, -20, 475, 518, 90, 92, 0),
(94, 65, 82, -10, 285, 336, 90, 96, 0),
(95, 62, 80, -20, 196, 239, 90, 98, 0),
(96, 60, 80, 10, 95, 156, 90, 0, 94),
(97, 60, 85, 30, 561, 622, 0, 0, 106),
(98, 58, 75, 20, 30, 84, 90, 0, 95),
(99, 55, 80, -20, 743, 820, 90, 100, 0),
(100, 55, 85, 20, 647, 726, 90, 0, 99),
(101, 25, 30, -10, 725, 786, 90, 51, 0),
(102, 48, 30, -10, 632, 693, 90, 64, 0),
(103, 28, 55, -10, 732, 777, 90, 23, 0),
(104, 88, 35, -20, 109, 170, 90, 78, 0),
(105, 5, 45, -10, 665, 716, 90, 36, 0),
(106, 60, 85, -30, 561, 622, 90, 97, 0);
INSERT 0 106
```

[The orders¶](#)

The original data needs to be converted to an appropriate table:

```
WITH deliveries AS (SELECT * FROM lc101_c WHERE dindex = 0)
SELECT
  row_number() over() AS id, p.demand,
  p.id as p_node_id, p.x AS p_x, p.y AS p_y, p.open AS p_open, p.close as p_close, p.service as p_service,
  d.id as d_node_id, d.x AS d_x, d.y AS d_y, d.open AS d_open, d.close as d_close, d.service as d_service
INTO c_lc101
FROM deliveries as d JOIN lc101_c as p ON (d.pindex = p.id);
SELECT 53
SELECT * FROM c_lc101 LIMIT 1;
id | demand | p_node_id | p_x | p_y | p_open | p_close | p_service | d_node_id | d_x | d_y | d_open | d_close | d_service
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1  | 10      | 3         | 42  | 66   | 65      | 146    | 90       | 75       | 45  | 65   | 997    | 1068    | 90
(1 row)
```

# Pgrouting Concepts¶

## pgRouting Concepts¶

This is a simple guide that go through some of the steps for getting started with pgRouting. This guide covers:

- [Graphs](#)
  - [Graph definition](#)
  - [Graph with cost](#)
  - [Graph with cost and reverse cost](#)
- [Graphs without geometries](#)
  - [Wiki example](#)
- [Graphs with geometries](#)
  - [Create a routing Database](#)
  - [Load Data](#)
  - [Build a routing topology](#)
  - [Adjust costs](#)
- [Check the Routing Topology](#)
  - [Crossing edges](#)
  - [Touching edges](#)
  - [Connecting components](#)
  - [Contraction of a graph](#)
- [Function's structure](#)
- [Function's overloads](#)
  - [One to One](#)
  - [One to Many](#)
  - [Many to One](#)
  - [Many to Many](#)



- [Combinations](#)
- [Inner Queries](#)
  - [Edges SQL](#)
  - [Combinations SQL](#)
  - [Restrictions SQL](#)
  - [Points SQL](#)
- [Parameters](#)
  - [Parameters for the Via functions](#)
  - [For the TRSP functions](#)
- [Result columns](#)
  - [Result columns for a path](#)
  - [Multiple paths](#)
  - [Result columns for cost functions](#)
  - [Result columns for flow functions](#)
  - [Result columns for spanning tree functions](#)
- [Performance Tips](#)
  - [For the Routing functions](#)
- [How to contribute](#)

**Graphs¶**

- [Graph definition](#)
- [Graph with cost](#)
- [Graph with cost and reverse\\_cost](#)

**Graph definition¶**

A graph is an ordered pair  $G = (V, E)$  where:

- $V$  is a set of vertices, also called nodes.
- $E \subseteq \{(u, v) \mid u, v \in V\}$

There are different kinds of graphs:

- Undirected graph
  - $E \subseteq \{(u, v) \mid u, v \in V\}$
- Undirected simple graph
  - $E \subseteq \{(u, v) \mid u, v \in V, u \neq v\}$
- Directed graph
  - $E \subseteq \{(u, v) \mid (u, v) \in (V \times V)\}$
- Directed simple graph
  - $E \subseteq \{(u, v) \mid (u, v) \in (V \times V), u \neq v\}$

Graphs:

- Do not have geometries.
- Some graph theory problems require graphs to have weights, called **cost** in pgRouting.

In pgRouting there are several ways to represent a graph on the database:

- With cost
  - (id, source, target, cost)
- With cost and reverse\_cost
  - (id, source, target, cost, reverse\_cost)

Where:

Column	Description
id	Identifier of the edge. Requirement to use the database in a consistent manner.
source	Identifier of a vertex.
target	Identifier of a vertex.
Weight of the edge (source, target):	
cost	<ul style="list-style-type: none"> <li>• When negative the edge (source, target) do not exist on the graph.</li> <li>• cost must exist in the query.</li> </ul>
Weight of the edge (target, source)	
reverse_cost	<ul style="list-style-type: none"> <li>• When negative the edge (target, source) do not exist on the graph.</li> </ul>

The decision of the graph to be **directed** or **undirected** is done when executing a pgRouting algorithm.

**[Graph with cost¶](#)**

The weighted directed graph,  $(G_d(V,E))$ :

- Graph data is obtained with a query  

```
SELECT id, source, target, cost FROM edges
```
- the set of edges  $(E)$ 
  - $E = \{(source\_id, target\_id, cost\_id) \mid \text{when } cost\_id \geq 0\}$
  - Edges where `cost` is non negative are part of the graph.
- the set of vertices  $(V)$ 
  - $V = \{source\_id \cup target\_id\}$
  - All vertices in `source` and `target` are part of the graph.

Directed graph

In a directed graph the edge  $((source\_id, target\_id, cost\_id))$  has directionality:  $(source\_id \rightarrow target\_id)$

For the following data:

```
SELECT *
FROM (VALUES (1, 1, 2, 5), (2, 1, 3, -3))
AS t(id, source, target, cost);
id | source | target | cost
-----+-----+-----+-----
1 | 1 | 2 | 5
2 | 1 | 3 | -3
(2 rows)
```

Edge  $(2) \nrightarrow (1 \rightarrow 3)$  is not part of the graph.

The data is representing the following graph:



Undirected graph

In an undirected graph the edge  $((source\_id, target\_id, cost\_id))$  does not have directionality:  $(source\_id \leftrightarrow target\_id)$

- In terms of a directed graph is like having two edges:  $(source\_id \rightarrow target\_id)$

For the following data:

```
SELECT *
FROM (VALUES (1, 1, 2, 5), (2, 1, 3, -3))
AS t(id, source, target, cost);
id | source | target | cost
-----+-----+-----+-----
1 | 1 | 2 | 5
2 | 1 | 3 | -3
(2 rows)
```

Edge  $(2) \nleftrightarrow (1 \leftrightarrow 3)$  is not part of the graph.

The data is representing the following graph:



[Graph with cost and reverse\\_cost](#)

The weighted directed graph,  $(G_d(V,E))$ , is defined by:

- Graph data is obtained with a query  

```
SELECT id, source, target, cost, reverse_cost FROM edges
```
- The set of edges  $(E)$ :
  - $E = \begin{matrix} \text{split} \\ \text{align} \end{matrix} \{ (source\_id, target\_id, cost\_id) \mid \text{when } cost\_id \geq 0 \} \cup \{ (target\_id, source\_id, reverse\_cost\_id) \mid \text{when } reverse\_cost\_id \geq 0 \} \end{matrix}$
  - Edges  $((source \rightarrow target))$  where `cost` is non negative are part of the graph.
  - Edges  $((target \rightarrow source))$  where `reverse_cost` is non negative are part of the graph.
- The set of vertices  $(V)$ :
  - $V = \{source\_id \cup target\_id\}$
  - All vertices in `source` and `target` are part of the graph.

Directed graph

In a directed graph both edges have directionality

- edge  $((source\_id, target\_id, cost\_id))$  has directionality:  $(source\_id \rightarrow target\_id)$
- edge  $((target\_id, source\_id, reverse\_cost\_id))$  has directionality:  $(target\_id \rightarrow source\_id)$

For the following data:

```
SELECT *
FROM (VALUES (1, 1, 2, 5, 2), (2, 1, 3, -3, 4), (3, 2, 3, 7, -1))
AS t(id, source, target, cost, reverse_cost);
id | source | target | cost | reverse_cost
-----+-----+-----+-----+-----
1 | 1 | 2 | 5 | 2
2 | 1 | 3 | -3 | 4
3 | 2 | 3 | 7 | -1
(3 rows)
```

Edges not part of the graph:

- $(2) \nrightarrow (1 \rightarrow 3)$
- $(3) \nrightarrow (3 \rightarrow 2)$

The data is representing the following graph:



Undirected graph

In a directed graph both edges do not have directionality

- Edge  $((source\_id, target\_id, cost\_id))$  is  $(source\_id \frac{1}{3} target\_id)$
- Edge  $((target\_id, source\_id, reverse\_cost\_id))$  is  $(target\_id \frac{1}{3} source\_id)$
- In terms of a directed graph is like having four edges:
  - $(source\_i \rightarrow target\_i)$
  - $(target\_i \rightarrow source\_i)$

For the following data:

```
SELECT *
FROM (VALUES (1, 1, 2, 5, 2), (2, 1, 3, -3, 4), (3, 2, 3, 7, -1))
AS t(id, source, target, cost, reverse_cost);
id | source | target | cost | reverse_cost
---+-----+-----+-----+-----
1 | 1 | 2 | 5 | 2
2 | 1 | 3 | -3 | 4
3 | 2 | 3 | 7 | -1
(3 rows)
```

Edges not part of the graph:

- $(2) (1 \frac{1}{3} 3)$
- $(3) (3 \frac{1}{3} 2)$

The data is representing the following graph:

Graphs without geometries¶

Personal relationships, genealogy, file dependency problems can be solved using pgRouting. Those problems, normally, do not come with geometries associated with the graph.

- [Wiki example](#)
  - [Prepare the database](#)
  - [Create a table](#)
  - [Insert the data](#)
  - [Find the shortest path](#)
  - [Vertex information](#)

Wiki example¶

Solve the example problem taken from [wikipedia](#):

Where:

- Problem is to find the shortest path from  $(1)$  to  $(5)$ .
- Is an undirected graph.
- Although visually looks like to have geometries, the drawing is not to scale.
  - No geometries associated to the vertices or edges
- Has 6 vertices  $(1,2,3,4,5,6)$
- Has 9 edges:  
$$E = \{(1,2,7), (1,3,9), (1,6,14), \& (2,3,10), (2,4,13), \& (3,4,11), (3,6,2), \& (4,5,6), \& (5,6,9)\}$$
- The graph can be represented in many ways for example:

Prepare the database¶

Create a database for the example, access the database and install pgRouting:

```
$ createdb wiki
$ psql wiki
wiki=# CREATE EXTENSION pgRouting CASCADE;
```

Create a table¶

The basic elements needed to perform basic routing on an undirected graph are:

Column	Type	Description
id	ANY-INTEGER	Identifier of the edge.
source	ANY-INTEGER	Identifier of the first end point vertex of the edge.
target	ANY-INTEGER	Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL	Weight of the edge (source, target)

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Using this table design for this example:

```
CREATE TABLE wiki (  
  id SERIAL,  
  source INTEGER,  
  target INTEGER,  
  cost INTEGER);  
CREATE TABLE
```

[Insert the data¶](#)

```
INSERT INTO wiki (source, target, cost) VALUES  
(1, 2, 7), (1, 3, 9), (1, 6, 14),  
(2, 3, 10), (2, 4, 15),  
(3, 6, 2), (3, 4, 11),  
(4, 5, 6),  
(5, 6, 9);  
INSERT 0 9
```

[Find the shortest path¶](#)

To solve this example [pgr\\_dijkstra](#) is used:

```
SELECT * FROM pgr_dijkstra(  
  'SELECT id, source, target, cost FROM wiki',  
  1, 5, false);  
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost  
-----  
1 | 1 | 1 | 5 | 1 | 2 | 7 | 0  
2 | 2 | 1 | 5 | 3 | 6 | 2 | 9  
3 | 3 | 1 | 5 | 6 | 9 | 11  
4 | 4 | 1 | 5 | 5 | -1 | 0 | 20  
(4 rows)
```

To go from \1) to \5) the path goes thru the following vertices:\1 \rightarrow 3 \rightarrow 6 \rightarrow 5)

[Vertex information¶](#)

To obtain the vertices information, use [pgr\\_extractVertices](#)

```
SELECT id, in_edges, out_edges  
FROM pgr_extractVertices('SELECT id, source, target FROM wiki');  
id | in_edges | out_edges  
-----  
3 | {2,4} | {6,7}  
5 | {8} | {9}  
4 | {5,7} | {8}  
2 | {1} | {4,5}  
1 | | {1,2,3}  
6 | {3,6,9} |  
(6 rows)
```

[Graphs with geometries¶](#)

- [Create a routing Database](#)
- [Load Data](#)
- [Build a routing topology](#)
- [Adjust costs](#)
  - [Update costs to length of geometry](#)
  - [Update costs based on codes](#)

[Create a routing Database¶](#)

The first step is to create a database and load pgRouting in the database.  
Typically create a database for each project.  
Once having the database to work in, load your data and build the routing application in that database.

```
createdb sampledata  
psql sampledata -c "CREATE EXTENSION pgrouting CASCADE"
```

[Load Data¶](#)

There are several ways to load your data into pgRouting.

- Manually creating a database.
  - [Graphs without geometries](#)
  - [Sample Data](#): a small graph used in the documentation examples
- Using [osm2pgrouting](#)

There are various open source tools that can help, like:

shp2pgsql:

- postgresql shapefile loader

ogr2ogr:

- vector data conversion utility

osm2pgsql:

- load OSM data into postgresql

Please note that these tools will **not** import the data in a structure compatible with pgRouting and when this happens the topology needs to be adjusted.

- Breakup a segments on each segment-segment intersection
- When missing, add columns and assign values to source, target, cost, reverse\_cost.
- Connect a disconnected graph.
- Create the complete graph topology
- Create one or more graphs based on the application to be developed.
  - Create a contracted graph for the high speed roads

- Create graphs per state/country

In few words:

Prepare the graph

What and how to prepare the graph, will depend on the application and/or on the quality of the data and/or on how close the information is to have a topology usable by pgRouting and/or some other factors not mentioned.

The steps to prepare the graph involve geometry operations using [PostGIS](#) and some others involve graph operations like [pgr\\_contraction](#) to contract a graph.

The [workshop](#) has a step by step on how to prepare a graph using Open Street Map data, for a small application.

The use of indexes on the database design in general:

- Have the geometries indexed.
- Have the identifiers columns indexed.

Please consult the [PostgreSQL](#) documentation and the [PostGIS](#) documentation.

#### [Build a routing topology¶](#)

The basic information to use the majority of the pgRouting functions `id, source, target, cost, [reverse_cost]` is what in pgRouting is called the routing topology.

`reverse_cost` is optional but strongly recommended to have in order to reduce the size of the database due to the size of the geometry columns. Having said that, in this documentation `reverse_cost` is used in this documentation.

When the data comes with geometries and there is no routing topology, then this step is needed.

All the start and end vertices of the geometries need an identifier that is to be stored in `source` and `target` columns of the table of the data. Likewise, `cost` and `reverse_cost` need to have the value of traversing the edge in both directions.

If the columns do not exist they need to be added to the table in question. (see [ALTER TABLE](#))

The function [pgr\\_extractVertices](#) is used to create a vertices table based on the edge identifier and the geometry of the edge of the graph.

```
SELECT * INTO vertices
FROM pgr_extractVertices('SELECT id, geom FROM edges ORDER BY id');
SELECT 18
```

Finally using the data stored on the vertices tables the `source` and `target` are filled up.

```
/* -- set the source information */
UPDATE edges AS e
SET source = v.id, x1 = x, y1 = y
FROM vertices AS v
WHERE ST_StartPoint(e.geom) = v.geom;
UPDATE 24
/* -- set the target information */
UPDATE edges AS e
SET target = v.id, x2 = x, y2 = y
FROM vertices AS v
WHERE ST_EndPoint(e.geom) = v.geom;
UPDATE 24
```

Data coming from OSM and using [osm2pgrouting](#) as an import tool, comes with the routing topology. See an example of using [osm2pgrouting](#) on the [workshop](#).

#### [Adjust costs¶](#)

For this example the `cost` and `reverse_cost` values are going to be the double of the length of the geometry.

##### [Update costs to length of geometry¶](#)

Suppose that `cost` and `reverse_cost` columns in the sample data represent:

- $\sqrt{1}$  when the edge exists in the graph
- $\sqrt{-1}$  when the edge does not exist in the graph

Using that information updating to the length of the geometries:

```
UPDATE edges SET
cost = sign(cost) * ST_length(geom) * 2,
reverse_cost = sign(reverse_cost) * ST_length(geom) * 2;
UPDATE 18
```

Which gives the following results:

```
SELECT id, cost, reverse_cost FROM edges;
```

id	cost	reverse_cost
6	2	2
7	2	2
4	2	2
5	2	-2
8	2	2
12	2	-2
11	2	-2
10	2	2
17	2.99999999999998	2.99999999999998
14	2	2
18	3.4000000000000004	3.4000000000000004
13	2	-2
15	2	2
16	2	2
9	2	2
3	-2	2
1	2	2
2	-2	2

(18 rows)

Note that to be able to follow the documentation examples, everything is based on the original graph.

Returning to the original data:

```
UPDATE edges SET
cost = sign(cost),
reverse_cost = sign(reverse_cost);
UPDATE 18
```

##### [Update costs based on codes¶](#)

Other datasets, can have a column with values like

- FT vehicle flow on the direction of the geometry
- TF vehicle flow opposite of the direction of the geometry
- B vehicle flow on both directions

Preparing a code column for the example:

```
ALTER TABLE edges ADD COLUMN direction TEXT;
ALTER TABLE
UPDATE edges SET
direction = CASE WHEN (cost>0 AND reverse_cost>0) THEN 'B' /* both ways */
              WHEN (cost>0 AND reverse_cost<0) THEN 'FT' /* direction of the LINESTRING */
              WHEN (cost<0 AND reverse_cost>0) THEN 'TF' /* reverse direction of the LINESTRING */
              ELSE " END;
UPDATE 18
/* unknown */
```

Adjusting the costs based on the codes:

```
UPDATE edges SET
cost = CASE WHEN (direction = 'B' OR direction = 'FT')
            THEN ST_length(geom) * 2
            ELSE -1 END,
reverse_cost = CASE WHEN (direction = 'B' OR direction = 'TF')
                    THEN ST_length(geom) * 2
                    ELSE -1 END;
UPDATE 18
```

Which gives the following results:

```
SELECT id, cost, reverse_cost FROM edges;
```

id	cost	reverse_cost
6	2	2
7	2	2
4	2	2
5	2	-1
8	2	2
12	2	-1
11	2	-1
10	2	2
17	2.99999999999998	2.99999999999998
14	2	2
18	3.4000000000000004	3.4000000000000004
13	2	-1
15	2	2
16	2	2
9	2	2
3	-1	2
1	2	2
2	-1	2

(18 rows)

Returning to the original data:

```
UPDATE edges SET
cost = sign(cost),
reverse_cost = sign(reverse_cost);
UPDATE 18
ALTER TABLE edges DROP COLUMN direction;
ALTER TABLE
```

Check the Routing Topology¶

- [Crossing edges](#)
  - [Fixing an intersection](#)
- [Touching edges](#)
  - [Fixing a gap](#)
- [Connecting components](#)
- [Contraction of a graph](#)
  - [Dead ends](#)
  - [Linear edges](#)


There are lots of possible problems in a graph.

- The data used may not have been designed with routing in mind.
- A graph has some very specific requirements.
- The graph is disconnected.
- There are unwanted intersections.
- The graph is too large and needs to be contracted.
- A sub graph is needed for the application.
- and many other problems that the pgRouting user, that is the application developer might encounter.

Crossing edges¶

To get the crossing edges:

```
SELECT a.id, b.id
FROM edges AS a, edges AS b
WHERE a.id < b.id AND st_crosses(a.geom, b.geom);
id | id
----+----
13 | 18
(1 row)
```



That information is correct, for example, when in terms of vehicles, is it a tunnel or bridge crossing over another road.

It might be incorrect, for example:

1. When it is actually an intersection of roads, where vehicles can make turns.
2. When in terms of electrical lines, the electrical line is able to switch roads even on a tunnel or bridge.

When it is incorrect, it needs fixing:

1. For vehicles and pedestrians
  - If the data comes from OSM and was imported to the database using `psm2pgrouting`, the fix needs to be done in the [OSM portal](#) and the data imported again.
  - In general when the data comes from a supplier that has the data prepared for routing vehicles, and there is a problem, the data is to be fixed from the supplier

2. For very specific applications
  - The data is correct when from the point of view of routing vehicles or pedestrians.
  - The data needs a local fix for the specific application.

Once analyzed one by one the crossings, for the ones that need a local fix, the edges need to be **split**.

The new edges need to be added to the edges table, the rest of the attributes need to be updated in the new edges, the old edges need to be removed and the routing topology needs to be updated.

#### Fixing an intersection

In this example the original edge table will be used to store the additional geometries.

An example use without results

Routing from \1) to \18\ gives no solution.

```
SELECT *
FROM pgr_dijkstra(SELECT id, source, target, cost, reverse_cost FROM edges', 1, 18);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
(0 rows)
```

Analyze the network for intersections.

```
SELECT
  e1.id id1, e2.id id2,
  ST_AsText(ST_Intersection(e1.geom, e2.geom)) AS point
FROM edges e1, edges e2
WHERE e1.id < e2.id AND ST_Crosses(e1.geom, e2.geom);
id1 | id2 | point
-----+-----+-----
13 | 18 | POINT(3.5 3)
(1 row)
```

The analysis tell us that the network has an intersection.

Prepare tables

Additional columns to control the origin of the segments.

```
ALTER TABLE edges ADD old_id BIGINT;
ALTER TABLE
```

Adding new segments.

Calling **pgr\_separateCrossing** and adding the new segments to the edges table.

```
INSERT INTO edges (old_id, geom)
SELECT id, geom
FROM pgr_separateCrossing(SELECT id, geom FROM edges');
INSERT 0 4
```

Update other values

In this example only **cost** and **reverse\_cost** are updated, where they are based on the length of the geometry and the directionality is kept using the **sign** function.

```
WITH
costs AS (
  SELECT e2.id, sign(e1.cost) * ST_Length(e2.geom) AS cost,
    sign(e1.reverse_cost) * ST_Length(e2.geom) AS reverse_cost
  FROM edges e1 JOIN edges e2 ON (e1.id = e2.old_id)
)
UPDATE edges e
SET (cost, reverse_cost) = (c.cost, c.reverse_cost)
FROM costs AS c WHERE e.id = c.id;
UPDATE 4
```

Update the topology

Insert the new vertices if any.

```
WITH
new_vertex AS (
  SELECT ev.*
  FROM pgr_extractVertices(SELECT id, geom FROM edges WHERE old_id IS NOT NULL) ev
  LEFT JOIN vertices v using(geom)
  WHERE v IS NULL)
INSERT INTO vertices (in_edges, out_edges,x,y,geom)
SELECT in_edges, out_edges,x,y,geom FROM new_vertex;
INSERT 0 1
```

Update source and target information on the edges table.

```
/* -- set the source information */
UPDATE edges AS e
SET source = v.id, x1 = x, y1 = y
FROM vertices AS v
WHERE source IS NULL AND ST_StartPoint(e.geom) = v.geom;
UPDATE 4
/* -- set the target information */
UPDATE edges AS e
SET target = v.id, x2 = x, y2 = y
FROM vertices AS v
WHERE target IS NULL AND ST_EndPoint(e.geom) = v.geom;
UPDATE 4
```

The example has results

Routing from \1) to \18\ gives a solution.

```
SELECT *
FROM pgr_dijkstra(SELECT id, source, target, cost, reverse_cost FROM edges', 1, 18);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 18 | 1 | 6 | 1 | 0
2 | 2 | 1 | 18 | 3 | 7 | 1 | 1
3 | 3 | 1 | 18 | 7 | 10 | 1 | 2
4 | 4 | 1 | 18 | 8 | 12 | 1 | 3
5 | 5 | 1 | 18 | 12 | 19 | 0.5 | 4
6 | 6 | 1 | 18 | 18 | -1 | 0 | 4.5
(6 rows)
```

#### Touching edges

Visually the edges seem to be connected, but internally they are not.

```
SELECT id, ST_AsText(geom)
FROM edges where id IN (14,17);
id | st_astext
-----+-----
```

```
17 | LINESTRING(0.5 3.5,1.999999999999 3.5)
14 | LINESTRING(2 3,2 4)
(2 rows)
```



The validity of the information is application dependent.

- Maybe there is a small barrier for vehicles but not for pedestrians.

Once analyzed one by one the touchings, for the ones that need a local fix, the edges need to be [split](#).

The new edges need to be added to the edges table, the rest of the attributes need to be updated in the new edges, the old edges need to be removed and the routing topology needs to be updated.

[Fixing a gap¶](#)

In this example the original edge table will be used to store the additional geometries.

An example use without results

Routing from \{1\} to \{2\} gives no solution.

```
SELECT *
FROM pgr_dijkstra('SELECT id, source, target, cost, reverse_cost FROM edges', 1, 2);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
(0 rows)
```

Analyze the network for gaps.

```
WITH
deadends AS (
  SELECT id AS vid, (in_edges || out_edges)[1] AS edge, geom AS vgeom
  FROM vertices
  WHERE array_length(in_edges || out_edges, 1) = 1
)
SELECT id, ST_AsText(geom), vid, ST_AsText(vgeom), ST_Distance(geom, vgeom)
FROM edges, deadends
WHERE id != edge AND ST_Distance(geom, vgeom) < 0.1;
id | st_astext | vid | st_astext | st_distance
-----+-----+-----+-----+-----
14 | LINESTRING(2 3,2 4) | 4 | POINT(1.999999999999 3.5) | 1.000088900582341e-12
(1 row)
```

The analysis tell us that the network has a gap.

Prepare tables

Additional columns to control the origin of the segments.

```
ALTER TABLE edges ADD old_id BIGINT;
ALTER TABLE
```

Adding new segments.

Calling [pgr\\_separateTouching](#) and adding the new segments to the edges table.

```
INSERT INTO edges (old_id, geom)
SELECT id, geom
FROM pgr_separateTouching('SELECT id, geom FROM edges');
INSERT 0 2
```

Update other values

In this example only cost and reverse\_cost are updated, where they are based on the length of the geometry and the directionality is kept using the [sign](#) function.

```
WITH
costs AS (
  SELECT e2.id,
  sign(e1.cost) * ST_Length(e2.geom) AS cost,
  sign(e1.reverse_cost) * ST_Length(e2.geom) AS reverse_cost
  FROM edges e1
  JOIN edges e2 ON (e1.id = e2.old_id)
)
UPDATE edges e SET (cost, reverse_cost) = (c.cost, c.reverse_cost)
FROM costs AS c
WHERE e.id = c.id;
UPDATE 2
```

Update the topology

Insert the new vertices if any.

```
WITH new_vertex AS (
  SELECT ev.*
  FROM pgr_extractVertices('SELECT id, geom FROM edges WHERE old_id IS NOT NULL') ev
  LEFT JOIN vertices v using(geom)
  WHERE v IS NULL
)
INSERT INTO vertices (in_edges, out_edges,x,y,geom)
SELECT in_edges, out_edges,x,y,geom
FROM new_vertex;
INSERT 0 0
```

Update source and target information on the edges table.

```
/* -- set the source information */
UPDATE edges AS e
SET source = v.id, x1 = x, y1 = y
FROM vertices AS v
WHERE source IS NULL AND ST_StartPoint(e.geom) = v.geom;
UPDATE 2
/* -- set the target information */
UPDATE edges AS e
SET target = v.id, x2 = x, y2 = y
FROM vertices AS v
WHERE target IS NULL AND ST_EndPoint(e.geom) = v.geom;
UPDATE 2
```

The example has results

Routing from \{1\} to \{2\} gives a solution.

```
SELECT *
FROM pgr_dijkstra('SELECT id, source, target, cost, reverse_cost FROM edges', 1, 2);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | 2 | 1 | 6 | 0
2 | 2 | 1 | 2 | 3 | 7 | 1 | 1
3 | 3 | 1 | 2 | 7 | 10 | 1 | 2
4 | 4 | 1 | 2 | 8 | 19 | 0.5 | 3
5 | 5 | 1 | 2 | 4 | 17 | 1 | 3.5
6 | 6 | 1 | 2 | 2 | -1 | 0 | 4.5
```



(6 rows)

Connecting components¶

To get the graph connectivity:

```
SELECT * FROM pgr_connectedComponents(
  'SELECT id, source, target, cost, reverse_cost FROM edges'
);
```

seq	component	node
1	1	1
2	1	3
3	1	5
4	1	6
5	1	7
6	1	8
7	1	9
8	1	10
9	1	11
10	1	12
11	1	15
12	1	16
13	1	17
14	2	2
15	2	4
16	13	13
17	13	14

(17 rows)

There are three basic ways to connect components:

- From the vertex to the starting point of the edge
- From the vertex to the ending point of the edge
- From the vertex to the closest vertex on the edge
  - This solution requires the edge to be split.

In this example [pgr\\_separateCrossing](#) and [pgr\\_separateTouching](#) will be used.

Get the connectivity

```
SELECT * FROM pgr_connectedComponents(
  'SELECT id, source, target, cost, reverse_cost FROM edges'
);
```

seq	component	node
1	1	1
2	1	3
3	1	5
4	1	6
5	1	7
6	1	8
7	1	9
8	1	10
9	1	11
10	1	12
11	1	15
12	1	16
13	1	17
14	2	2
15	2	4
16	13	13
17	13	14

(17 rows)

Prepare tables

In this example: the edges table will need an additional column and the vertex table will be rebuilt completely.

```
ALTER TABLE edges ADD old_id BIGINT;
ALTER TABLE
DROP TABLE vertices;
DROP TABLE
```

Insert new edges

Using [pgr\\_separateCrossing](#) and [pgr\\_separateTouching](#) insert the results into the edges table.

```
INSERT INTO edges (old_id, geom)
SELECT id, geom FROM pgr_separateCrossing('SELECT * FROM edges')
UNION
SELECT id, geom FROM pgr_separateTouching('SELECT * FROM edges');
INSERT 0 6
```

Create the vertices table

Using [pgr\\_extractVertices](#) create the table.

```
CREATE TABLE vertices AS
SELECT * FROM pgr_extractVertices('SELECT id, geom FROM edges');
SELECT 18
```

Update the topology

```
/* -- set the source information */
UPDATE edges AS e
SET source = v.id, x1 = x, y1 = y
FROM vertices AS v
WHERE ST_StartPoint(e.geom) = v.geom;
UPDATE 24
/* -- set the target information */
UPDATE edges AS e
SET target = v.id, x2 = x, y2 = y
FROM vertices AS v
WHERE ST_EndPoint(e.geom) = v.geom;
UPDATE 24
```

Update other values

In this example only cost and reverse\_cost are updated, where they are based on the length of the geometry and the directionality is kept using the sign function.

```
UPDATE edges e
SET cost = ST_length(e.geom)*sign(e1.cost),
    reverse_cost = ST_length(e.geom)*sign(e1.reverse_cost)
FROM edges e1
WHERE e.cost IS NULL AND e1.id = e.old_id;
UPDATE 6

SELECT * FROM pgr_connectedComponents(
  'SELECT id, source, target, cost, reverse_cost FROM edges'
);
```

seq	component	node
1	1	1
2	1	2
3	1	3
4	1	4
5	1	5
6	1	6
7	1	7
8	1	8
9	1	9
10	1	10
11	1	11
12	1	12
13	1	13
14	1	14
15	1	15
16	1	16
17	1	17
18	1	18

(18 rows)

#### Contraction of a graph¶

The graph can be reduced in size using [Contraction - Family of functions](#)

When to contract will depend on the size of the graph, processing times, correctness of the data, on the final application, or any other factor not mentioned.

A fairly good method of finding out if contraction can be useful is because of the number of dead ends and/or the number of linear edges.

A complete method on how to contract and how to use the contracted graph is described on [Contraction - Family of functions](#)

#### Dead ends¶


To get the dead ends:

```
SELECT id FROM vertices
WHERE array_length(in_edges || out_edges, 1) = 1;
id
----
1
2
5
(3 rows)
```

A dead end happens when

- The vertex is the limit of a cul-de-sac, a no-through road or a no-exit road.
- The vertex is on the limit of the imported graph.
  - If a larger graph is imported then the vertex might not be a dead end

Node \4\), is a dead end on the query, even that it visually looks like an end point of 3 edges.



Is node \4\) a dead end or not?

The answer to that question will depend on the application.

- Is there such a small curb:
  - That does not allow a vehicle to use that visual intersection?
  - Is the application for pedestrians and therefore the pedestrian can easily walk on the small curb?
  - Is the application for the electricity and the electrical lines than can easily be extended on top of the small curb?
- Is there a big cliff and from eagles view look like the dead end is close to the segment?

Depending on the answer, modification of the data might be needed.

When there are many dead ends, to speed up processing, the [Contraction - Family of functions](#) functions can be used to contract the graph.

#### Linear edges¶

To get the linear edges:

```
SELECT id FROM vertices
WHERE array_length(in_edges || out_edges, 1) = 2;
id
----
3
9
13
15
16
(5 rows)
```

These linear vertices are correct, for example, when those the vertices are speed bumps, stop signals and the application is taking them into account.

When there are many linear vertices, that need not to be taken into account, to speed up the processing, the [Contraction - Family of functions](#) functions can be used to contract the problem.

#### Function's structure¶

Once the graph preparation work has been done above, it is time to use a

The general form of a pgRouting function call is:

pgr\_<name>([Inner queries](#), [parameters](#), [ Optional parameters])

Where:

- **Inner queries:** Are compulsory parameters that are TEXT strings containing SQL queries.
- **parameters:** Additional compulsory parameters needed by the function.
- Optional parameters: Are non compulsory **named** parameters that have a default value when omitted.

The compulsory parameters are positional parameters, the optional parameters are named parameters.

For example, for this [pgr\\_dijkstra](#) signature:

`pgr_dijkstra(Edges SQL, start vids, end vids, [directed])`

- **Edges SQL:**
  - Is the first parameter.
  - It is compulsory.
  - It is an inner query.
  - It has no name, so **Edges SQL** gives an idea of what kind of inner query needs to be used
- **start vid:**
  - Is the second parameter.
  - It is compulsory.
  - It has no name, so **start vid** gives an idea of what the second parameter's value should contain.
- **end vid**
  - Is the third parameter.
  - It is compulsory.
  - It has no name, so **end vid** gives an idea of what the third parameter's value should contain
- **directed**
  - Is the fourth parameter.
  - It is optional.
  - It has a name.

The full description of the parameters are found on the [Parameters](#) section of each function.

### [Function's overloads¶](#)

A function might have different overloads. The most common are called:

- [One to One](#)
- [One to Many](#)
- [Many to One](#)
- [Many to Many](#)
- [Combinations](#)

Depending on the overload the parameters types change.

- **One: ANY-INTEGER**
- **Many: ARRAY [ANY-INTEGER]**

Depending of the function the overloads may vary. But the concept of parameter type change remains the same.

#### [One to One¶](#)

When routing from:

- From **one** starting vertex
- to **one** ending vertex

#### [One to Many¶](#)

When routing from:

- From **one** starting vertex
- to **many** ending vertices

#### [Many to One¶](#)

When routing from:

- From **many** starting vertices
- to **one** ending vertex

#### [Many to Many¶](#)

When routing from:

- From **many** starting vertices
- to **many** ending vertices

#### [Combinations¶](#)

When routing from:

- From **many** different starting vertices
- to **many** different ending vertices
- Every tuple specifies a pair of a start vertex and an end vertex
- Users can define the combinations as desired.
- Needs a [Combinations SQL](#)

Inner Queries¶

- Edges SQL
  - General
  - General without id
  - General with (X,Y)
  - Flow
- Combinations SQL
- Restrictions SQL
- Points SQL

There are several kinds of valid inner queries and also the columns returned are depending of the function. Which kind of inner query will depend on the function's requirements. To simplify the variety of types, **ANY-INTEGER** and **ANY-NUMERICAL** is used.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Edges SQL¶

General¶

Edges SQL for

- Dijkstra - Family of functions
- withPoints - Family of functions
- Bidirectional Dijkstra - Family of functions
- Components - Family of functions
- Kruskal - Family of functions
- Prim - Family of functions
- Some uncategorised functions

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
			Weight of the edge (target, source)
reverse_cost	ANY-NUMERICAL	-1	<ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

General without id¶

Edges SQL for

- All Pairs - Family of Functions

Column	Type	Default	Description
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
			Weight of the edge (target, source)
reverse_cost	ANY-NUMERICAL	-1	<ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

[General with \(X,Y\)¶](#)

Edges SQL for

- [A\\* - Family of functions](#)
- [Bidirectional A\\* - Family of functions](#)

Parameter	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target) <ul style="list-style-type: none"><li>• When negative: edge (source, target) does not exist, therefore it's not part of the graph.</li></ul>
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source), <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>
x1	ANY-NUMERICAL		X coordinate of source vertex.
y1	ANY-NUMERICAL		Y coordinate of source vertex.
x2	ANY-NUMERICAL		X coordinate of target vertex.
y2	ANY-NUMERICAL		Y coordinate of target vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

[Flow¶](#)

Edges SQL for [Flow - Family of functions](#)

Edges SQL for

- [pgr\\_pushRelabel](#)
- [pgr\\_edmondsKarp](#)
- [pgr\\_boykovKolmogorov](#)

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
capacity	ANY-INTEGER		Weight of the edge (source, target)
reverse_capacity	ANY-INTEGER	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Edges SQL for the following functions of [Flow - Family of functions](#)

- [pgr\\_maxFlowMinCost - Experimental](#)
- [pgr\\_maxFlowMinCost\\_Cost - Experimental](#)

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.

Column	Type	Default	Description
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
target	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
capacity	<b>ANY-INTEGER</b>		Capacity of the edge (source, target) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>
reverse_capacity	<b>ANY-INTEGER</b>	-1	Capacity of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>
cost	<b>ANY-NUMERICAL</b>		Weight of the edge (source, target) if it exist
reverse_cost	<b>ANY-NUMERICAL</b>	\(-1\)	Weight of the edge (target, source) if it exist

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

[Combinations SQL¶](#)

Used in combination signatures

Parameter	Type	Description
source	<b>ANY-INTEGER</b>	Identifier of the departure vertex.
target	<b>ANY-INTEGER</b>	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

[Restrictions SQL¶](#)

Column	Type	Description
path	ARRAY <b>[ANY-INTEGER]</b>	Sequence of edge identifiers that form a path that is not allowed to be taken. - Empty arrays or NULL arrays are ignored. - Arrays that have a NULL element will raise an exception.
Cost	<b>ANY-NUMERICAL</b>	Cost of taking the forbidden path.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

[Points SQL¶](#)

Points SQL for

- [withPoints - Family of functions](#)

Parameter	Type	Default	Description
			Identifier of the point. <ul style="list-style-type: none"> <li>Use with positive value, as internally will be converted to negative value</li> </ul>
pid	<b>ANY-INTEGER</b>	<b>value</b>	<ul style="list-style-type: none"> <li>If column is present, it can not be NULL.</li> <li>If column is not present, a sequential negative <b>value</b> will be given automatically.</li> </ul>
edge_id	<b>ANY-INTEGER</b>		Identifier of the "closest" edge to the point.
fraction	<b>ANY-NUMERICAL</b>		Value in <0,1> that indicates the relative position from the first end point of the edge.

Parameter	Type	Default	Description
side	CHAR	b	Value in [b, r, l, NULL] indicating if the point is:
			• In the right <i>r</i> ,
			• In the left <i>l</i> ,
			• In both sides <i>b</i> , NULL

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Parameters¶

The main parameter of the majority of the pgRouting functions is a query that selects the edges of the graph.

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below.

Depending on the family or category of a function it will have additional parameters, some of them are compulsory and some are optional.

The compulsory parameters are nameless and must be given in the required order. The optional parameters are named parameters and will have a default value.

Parameters for the Via functions¶

- [pgr\\_dijkstraVia - Proposed](#)

Parameter	Type	Default	Description
<a href="#">Edges SQL</a>	TEXT		SQL query as described.
<b>via vertices</b>	ARRAY [ANY-INTEGER]		Array of ordered vertices identifiers that are going to be visited.
directed	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true Graph is considered <i>Directed</i></li> <li>When false the graph is considered as Undirected.</li> </ul>
strict	BOOLEAN	false	<ul style="list-style-type: none"> <li>When true if a path is missing stops and returns <b>EMPTY SET</b></li> <li>When false ignores missing paths returning all paths found</li> </ul>
U_turn_on_edge	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true departing from a visited vertex will not try to avoid using the edge used to reach it. In other words, U turn using the edge with same identifier is allowed.</li> <li>When false when a departing from a visited vertex tries to avoid using the edge used to reach it. In other words, U turn using the edge with same identifier is used when no other path is found.</li> </ul>

For the TRSP functions¶

- [pgr\\_trsp - Proposed](#)

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	SQL query as described.
<a href="#">Restrictions SQL</a>	TEXT	SQL query as described.
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
<b>start vid</b>	ANY-INTEGER	Identifier of the departure vertex.
<b>start vids</b>	ARRAY [ANY-INTEGER]	Array of identifiers of destination vertices.
<b>end vid</b>	ANY-INTEGER	Identifier of the departure vertex.
<b>end vids</b>	ARRAY [ANY-INTEGER]	Array of identifiers of destination vertices.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns¶

- [Result columns for a path](#)
- [Multiple paths](#)
  - [Selective for multiple paths.](#)

- [Non selective for multiple paths](#)
- [Result columns for cost functions](#)
- [Result columns for flow functions](#)
- [Result columns for spanning tree functions](#)

There are several kinds of columns returned are depending of the function.

#### [Result columns for a path¶](#)

Used in functions that return one path solution

Returns set of (seq, path\_seq [, start\_vid] [, end\_vid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex. Returned when multiple starting vetrices are in the query. <ul style="list-style-type: none"> <li>• <a href="#">Many to One</a></li> <li>• <a href="#">Many to Many</a></li> </ul>
end_vid	BIGINT	Identifier of the ending vertex. Returned when multiple ending vertices are in the query. <ul style="list-style-type: none"> <li>• <a href="#">One to Many</a></li> <li>• <a href="#">Many to Many</a></li> </ul>
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Used in functions the following:

- [pgr\\_withPoints - Proposed](#)

Returns set of (seq, path\_seq [, start\_pid] [, end\_pid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_seq	INTEGER	Relative position in the path. <ul style="list-style-type: none"> <li>• 1 For the first row of the path.</li> </ul>
start_pid	BIGINT	Identifier of a starting vertex/point of the path. <ul style="list-style-type: none"> <li>• When positive is the identifier of the starting vertex.</li> <li>• When negative is the identifier of the starting point.</li> <li>• Returned on <a href="#">Many to One</a> and <a href="#">Many to Many</a></li> </ul>
end_pid	BIGINT	Identifier of an ending vertex/point of the path. <ul style="list-style-type: none"> <li>• When positive is the identifier of the ending vertex.</li> <li>• When negative is the identifier of the ending point.</li> <li>• Returned on <a href="#">One to Many</a> and <a href="#">Many to Many</a></li> </ul>
node	BIGINT	Identifier of the node in the path fromstart_pid to end_pid. <ul style="list-style-type: none"> <li>• When positive is the identifier of the a vertex.</li> <li>• When negative is the identifier of the a point.</li> </ul>
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence. <ul style="list-style-type: none"> <li>• -1 for the last row of the path.</li> </ul>
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence. <ul style="list-style-type: none"> <li>• 0 For the first row of the path.</li> </ul>
agg_cost	FLOAT	Aggregate cost from start_vid to node. <ul style="list-style-type: none"> <li>• 0 For the first row of the path.</li> </ul>

Used in functions the following:

- [pgr\\_dijkstraNear - Proposed](#)



Returns (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex of the current path.
end_vid	BIGINT	Identifier of the ending vertex of the current path.
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

[Multiple paths¶](#)

[Selective for multiple paths¶](#)

The columns depend on the function call.

Set of (seq, path\_id, path\_seq [, start\_vid] [, end\_vid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Path identifier. <ul style="list-style-type: none"><li>Has value 1 for the first of a path fromstart_vid to end_vid.</li></ul>
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex. Returned when multiple starting vetrices are in the query. <ul style="list-style-type: none"><li><a href="#">Many to One</a></li><li><a href="#">Many to Many</a></li><li><a href="#">Combinations</a></li></ul>
end_vid	BIGINT	Identifier of the ending vertex. Returned when multiple ending vertices are in the query. <ul style="list-style-type: none"><li><a href="#">One to Many</a></li><li><a href="#">Many to Many</a></li><li><a href="#">Combinations</a></li></ul>
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

[Non selective for multiple paths¶](#)

Regardless of the call, al the columns are returned.

- [pgr\\_trsp - Proposed](#)

Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Path identifier. <ul style="list-style-type: none"><li>Has value 1 for the first of a path fromstart_vid to end_vid.</li></ul>
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.

Column	Type	Description
node	BIGINT	Identifier of the node in the path from start_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go from node to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

#### [Result columns for cost functions¶](#)

Used in the following

- [Cost - Category](#)
- [Cost Matrix - Category](#)
- [All Pairs - Family of Functions](#)

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

Note

When start\_vid or end\_vid columns have negative values, the identifier is for a Point.

#### [Result columns for flow functions¶](#)

Edges SQL for the following

- [Flow - Family of functions](#)

Column	Type	Description
seq	INT	Sequential value starting from 1.
edge	BIGINT	Identifier of the edge in the original query (edges_sql).
start_vid	BIGINT	Identifier of the first end point vertex of the edge.
end_vid	BIGINT	Identifier of the second end point vertex of the edge.
flow	BIGINT	Flow through the edge in the direction (start_vid, end_vid).
residual_capacity	BIGINT	Residual capacity of the edge in the direction (start_vid, end_vid).

Edges SQL for the following functions of [Flow - Family of functions](#)

- [pgr\\_maxFlowMinCost - Experimental](#)

Column	Type	Description
seq	INT	Sequential value starting from 1.
edge	BIGINT	Identifier of the edge in the original query (edges_sql).
source	BIGINT	Identifier of the first end point vertex of the edge.
target	BIGINT	Identifier of the second end point vertex of the edge.
flow	BIGINT	Flow through the edge in the direction (source, target).
residual_capacity	BIGINT	Residual capacity of the edge in the direction (source, target).
cost	FLOAT	The cost of sending this flow through the edge in the direction (source, target).
agg_cost	FLOAT	The aggregate cost.

#### [Result columns for spanning tree functions¶](#)

Edges SQL for the following

- [pgr\\_prim](#)

- [pgr\\_kruskal](#)

Returns set of (edge, cost)

Column	Type	Description
edge	BIGINT	Identifier of the edge.
cost	FLOAT	Cost to traverse the edge.

[Performance Tips¶](#)

- [For the Routing functions](#)

[For the Routing functions¶](#)

To get faster results bound the queries to an area of interest of routing.  
In this example Use an inner query SQL that does not include some edges in the routing function and is within the area of the results.  
Given this area:

```
SELECT id, source, target, cost, reverse_cost
FROM edges
WHERE geom && (
  SELECT st_buffer(geom, 1) as myarea
  FROM edges where id = 6) ORDER BY id;
id | source | target | cost | reverse_cost
-----+-----+-----+-----+-----
1 | 5 | 6 | 1 | 1
2 | 6 | 10 | -1 | 1
4 | 6 | 7 | 1 | 1
6 | 1 | 3 | 1 | 1
7 | 3 | 7 | 1 | 1
8 | 7 | 11 | 1 | 1
10 | 7 | 8 | 1 | 1
12 | 8 | 12 | 1 | -1
14 | 8 | 9 | 1 | 1
20 | 8 | 4 | 0.5 | 0.5
(10 rows)
```

Calculate a route:

```
SELECT * FROM pgr_dijkstra($$
SELECT id, source, target, cost, reverse_cost
FROM edges
WHERE geom && (
  SELECT st_buffer(geom, 1) AS myarea
  FROM edges WHERE id = 6)$$,
7, 8);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 7 | 8 | 7 | 10 | 1 | 0
2 | 2 | 7 | 8 | 8 | -1 | 0 | 1
(2 rows)
```

[How to contribute¶](#)

Wiki

- Edit an existing [pgRouting Wiki](#) page.
- Or create a new Wiki page
  - Create a page on the [pgRouting Wiki](#)
  - Give the title an appropriate name
- [Example](#)

Adding Functionality to pgRouting

Consult the [developer's documentation](#)

Indices and tables

- [Index](#)
- [Search Page](#)

[Function Families¶](#)

[Function Families¶](#)

[All Pairs - Family of Functions](#)

- [pgr\\_floydWarshall](#) - Floyd-Warshall's algorithm.
- [pgr\\_johnson](#) - Johnson's algorithm

[A\\* - Family of functions](#)

- [pgr\\_aStar](#) - A\* algorithm for the shortest path.
- [pgr\\_aStarCost](#) - Get the aggregate cost of the shortest paths.
- [pgr\\_aStarCostMatrix](#) - Get the cost matrix of the shortest paths.

[Bidirectional A\\* - Family of functions](#)

- [pgr\\_bdAstar](#) - Bidirectional A\* algorithm for obtaining paths.
- [pgr\\_bdAstarCost](#) - Bidirectional A\* algorithm to calculate the cost of the paths.
- [pgr\\_bdAstarCostMatrix](#) - Bidirectional A\* algorithm to calculate a cost matrix of paths.

[Bidirectional Dijkstra - Family of functions](#)

- [pgr\\_bdDijkstra](#) - Bidirectional Dijkstra algorithm for the shortest paths.
- [pgr\\_bdDijkstraCost](#) - Bidirectional Dijkstra to calculate the cost of the shortest paths

- [pgr\\_bdDijkstraCostMatrix](#) - Bidirectional Dijkstra algorithm to create a matrix of costs of the shortest paths.

#### [Components - Family of functions](#)

- [pgr\\_connectedComponents](#) - Connected components of an undirected graph.
- [pgr\\_strongComponents](#) - Strongly connected components of a directed graph.
- [pgr\\_biconnectedComponents](#) - Biconnected components of an undirected graph.
- [pgr\\_articulationPoints](#) - Articulation points of an undirected graph.
- [pgr\\_bridges](#) - Bridges of an undirected graph.

#### [Contraction - Family of functions](#)

- [pgr\\_contraction](#)

#### [Dijkstra - Family of functions](#)

- [pgr\\_dijkstra](#) - Dijkstra's algorithm for the shortest paths.
- [pgr\\_dijkstraCost](#) - Get the aggregate cost of the shortest paths.
- [pgr\\_dijkstraCostMatrix](#) - Use [pgr\\_dijkstra](#) to create a costs matrix.
- [pgr\\_drivingDistance](#) - Use [pgr\\_dijkstra](#) to calculate catchment information.
- [pgr\\_KSP](#) - Use Yen algorithm with [pgr\\_dijkstra](#) to get the K shortest paths.

#### [Flow - Family of functions](#)

- [pgr\\_maxFlow](#) - Only the Max flow calculation using Push and Relabel algorithm.
- [pgr\\_boykovKolmogorov](#) - Boykov and Kolmogorov with details of flow on edges.
- [pgr\\_edmondsKarp](#) - Edmonds and Karp algorithm with details of flow on edges.
- [pgr\\_pushRelabel](#) - Push and relabel algorithm with details of flow on edges.
- Applications
  - [pgr\\_edgeDisjointPaths](#) - Calculates edge disjoint paths between two groups of vertices.
  - [pgr\\_maxCardinalityMatch](#) - Calculates a maximum cardinality matching in a graph.

#### [Kruskal - Family of functions](#)

- [pgr\\_kruskal](#)
- [pgr\\_kruskalBFS](#)
- [pgr\\_kruskalDD](#)
- [pgr\\_kruskalDFS](#)

#### [Metrics - Family of functions](#)

- [pgr\\_degree](#) - Returns a set of vertices and corresponding count of incident edges to the vertex.

#### [Prim - Family of functions](#)

- [pgr\\_prim](#)
- [pgr\\_primBFS](#)
- [pgr\\_primDD](#)
- [pgr\\_primDFS](#)

#### [Reference](#)

- [pgr\\_version](#)
- [pgr\\_full\\_version](#)

#### [Topology - Family of Functions](#)

The following functions modify the database directly therefore the user must have special permissions given by the administrators to use them.

- [pgr\\_createTopology](#) - [Deprecated since v3.8.0](#) - create a topology based on the geometry.
- [pgr\\_createVerticesTable](#) - [Deprecated since 3.8.0](#) - reconstruct the vertices table based on the source and target information.
- [pgr\\_analyzeGraph](#) - [Deprecated since 3.8.0](#) - to analyze the edges and vertices of the edge table.
- [pgr\\_analyzeOneWay](#) - [Deprecated since 3.8.0](#) - to analyze directionality of the edges.
- [pgr\\_nodeNetwork](#) - [Deprecated since 3.8.0](#) - to create nodes to a not noded edge table.

#### [Traveling Sales Person - Family of functions](#)

- [pgr\\_TSP](#) - When input is given as matrix cell information.
- [pgr\\_TSPeuclidean](#) - When input are coordinates.

[pgr\\_trsp](#) - [Proposed](#) - Turn Restriction Shortest Path (TRSP)

#### Utilities

- [pgr\\_extractVertices](#) - Extracts vertex information based on the edge table information.
- [pgr\\_findCloseEdges](#) - Finds close edges of points on the fly
- [pgr\\_separateCrossing](#) - Breaks geometries that cross each other.
- [pgr\\_separateTouching](#) - Breaks geometries that (almost) touch each other.

#### Functions by categories¶

##### [Cost - Category](#)

- [pgr\\_aStarCost](#)
- [pgr\\_bdAStarCost](#)
- [pgr\\_dijkstraCost](#)
- [pgr\\_bdDijkstraCost](#)

- [pgr\\_dijkstraNearCost - Proposed](#)

#### [Cost Matrix - Category](#)

- [pgr\\_aStarCostMatrix](#)
- [pgr\\_dijkstraCostMatrix](#)
- [pgr\\_bdAstarCostMatrix](#)
- [pgr\\_bdDijkstraCostMatrix](#)

#### [Driving Distance - Category](#)

- [pgr\\_drivingDistance](#) - Driving Distance based on Dijkstra's algorithm
- [pgr\\_primDD](#) - Driving Distance based on Prim's algorithm
- [pgr\\_kruskalDD](#) - Driving Distance based on Kruskal's algorithm
- Post processing
  - [pgr\\_alphaShape](#) - Alpha shape computation

#### [K shortest paths - Category](#)

- [pgr\\_KSP](#) - Yen's algorithm based on pgr\_dijkstra

#### [Spanning Tree - Category](#)

- [Kruskal - Family of functions](#)
- [Prim - Family of functions](#)

#### [BFS - Category](#)

- [pgr\\_kruskalBFS](#)
- [pgr\\_primBFS](#)

#### [DFS - Category](#)

- [pgr\\_kruskalDFS](#)
- [pgr\\_primDFS](#)

#### **All Pairs - Family of Functions ¶**

The following functions work on all vertices pair combinations

- [pgr\\_floydWarshall](#) - Floyd-Warshall's algorithm.
- [pgr\\_johnson](#) - Johnson's algorithm

#### **`pgr_floydWarshall` ¶**

`pgr_floydWarshall` - Returns the sum of the costs of the shortest path for each pair of nodes in the graph using Floyd-Warshall algorithm.

Availability

- Version 2.2.0
  - Signature change
  - Old signature no longer supported
- Version 2.0.0
  - New official function.

#### **Description ¶**

The Floyd-Warshall algorithm, also known as Floyd's algorithm, is a good choice to calculate the sum of the costs of the shortest path for each pair of nodes in the graph, *fatense graphs*. We use Boost's implementation which runs in  $\mathcal{O}(\Theta(V^3))$  time,

The main characteristics are:

- It does not return a path.
- Returns the sum of the costs of the shortest path for each pair of nodes in the graph.
- Process is done only on edges with positive costs.
- Boost returns a  $(V \times V)$  matrix, where the infinity values. Represent the distance between vertices for which there is no path.
  - We return only the non infinity values in form of a set `o(start_vid, end_vid, agg_cost)`.
- Let be the case the values returned are stored in a table, so the unique index would be the pair `(start_vid, end_vid)`.
- For the undirected graph, the results are symmetric.
  - The `agg_cost` of  $(u, v)$  is the same as for  $(v, u)$ .
- When `start_vid = end_vid`, the `agg_cost` = 0.
- **Recommended, use a bounding box of no more than 3500 edges.**

Boost Graph Inside

#### **Signatures ¶**

Summary

`pgr_floydWarshall`([Edges SQL](#), [directed])

Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

For a directed subgraph with edges  $\{(1, 2, 3, 4)\}$ .

```
SELECT * FROM pgr_floydWarshall(
  'SELECT id, source, target, cost, reverse_cost
  FROM edges where id < 5'
) ORDER BY start_vid, end_vid,
start_vid | end_vid | agg_cost
-----+-----+-----
5 | 6 | 1
5 | 7 | 2
6 | 5 | 1
6 | 7 | 1
7 | 5 | 2
7 | 6 | 1
10 | 5 | 2
10 | 6 | 1
10 | 7 | 2
15 | 5 | 3
15 | 6 | 2
15 | 7 | 3
15 | 10 | 1
(13 rows)
```

Parameters1

Parameter	Type	Default	Description
<a href="#">Edges_SQL</a>	TEXT		<a href="#">Edges_SQL</a> as described below.

Optional parameters1

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries1

Edges SQL1

Column	Type	Default	Description
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns1

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

See Also1

- [pgr\\_johnson](#)
- [Sample Data](#)
- Boost [floyd-Warshall](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_johnson1

pgr\_johnson - Returns the sum of the costs of the shortest path for each pair of nodes in the graph using Floyd-Warshall algorithm.

Availability

- Version 2.2.0
  - Signature change
  - Old signature no longer supported
- Version 2.0.0
  - New official function.

Description

The Johnson algorithm, is a good choice to calculate the sum of the costs of the shortest path for each pair of nodes in the graph, for *sparse graphs*. It uses the Boost's implementation which runs in  $O(V E \log V)$  time,

The main characteristics are:

- It does not return a path.
- Returns the sum of the costs of the shortest path for each pair of nodes in the graph.
- Process is done only on edges with positive costs.
- Boost returns a  $(V \times V)$  matrix, where the infinity values. Represent the distance between vertices for which there is no path.
  - We return only the non infinity values in form of a set of  $(start\_vid, end\_vid, agg\_cost)$ .
- Let be the case the values returned are stored in a table, so the unique index would be the pair  $(start\_vid, end\_vid)$ .
- For the undirected graph, the results are symmetric.
  - The  $agg\_cost$  of  $(u, v)$  is the same as for  $(v, u)$ .
- When  $start\_vid = end\_vid$ , the  $agg\_cost = 0$ .
- **Recommended, use a bounding box of no more than 3500 edges.**

 Boost Graph Inside

Signatures

Summary

pgr\_johnson([Edges SQL](#), [directed])

Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

For a directed subgraph with edges  $\{(1, 2, 3, 4)\}$ .

```
SELECT * FROM pgr_johnson(
'SELECT source, target, cost FROM edges
WHERE id < 5'
) ORDER BY start_vid, end_vid;
start_vid | end_vid | agg_cost
-----+-----+-----
5 | 6 | 1
5 | 7 | 2
6 | 7 | 1
(3 rows)
```

Parameters

Parameter	Type	Default	Description
<a href="#">Edges SQL</a>	TEXT		<a href="#">Edges SQL</a> as described below.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>• When true the graph is considered <i>Directed</i></li><li>• When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

See Also

- [pgr\\_floydWarshall](#)
- [Sample Data](#)
- Boost [Johnson](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Introduction

The main characteristics are:

- It does not return a path.
- Returns the sum of the costs of the shortest path for each pair of nodes in the graph.
- Process is done only on edges with positive costs.
- Boost returns a  $(V \times V)$  matrix, where the infinity values. Represent the distance between vertices for which there is no path.
  - We return only the non infinity values in form of a set of (start\_vid, end\_vid, agg\_cost).
- Let be the case the values returned are stored in a table, so the unique index would be the pair (start\_vid, end\_vid).
- For the undirected graph, the results are symmetric.
  - The agg\_cost of (u, v) is the same as for (v, u).
- When start\_vid = end\_vid, the agg\_cost = 0.
- **Recommended, use a bounding box of no more than 3500 edges.**

Parameters

Parameter	Type	Default	Description
<a href="#">Edges SQL</a>	TEXT		<a href="#">Edges SQL</a> as described below.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>• When true the graph is considered <i>Directed</i></li><li>• When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT



ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns1

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

Performance1

The following tests:

- non server computer
- with AMD 64 CPU
- 4G memory
- trusty
- PostgreSQL version 9.3

Data1

The following data was used

BBOX="-122.8,45.4,-122.5,45.6"  
wget --progress=dot:mega -O "sampledata.osm" "<https://www.overpass-api.de/api/xapi?lbbox=-11@meta>"

Data processing was done with osm2pgrouting-alpha

createdb portland  
psql -c "create extension postgis" portland  
psql -c "create extension pgrouting" portland  
osm2pgrouting -f sampledata.osm -d portland -s 0

Results1

Test:

One

This test is not with a bounding box The density of the passed graph is extremely low. For each <SIZE> 30 tests were executed to get the average The tested query is:

SELECT count(\*) FROM pgr\_floydWarshall(  
 'SELECT gid as id, source, target, cost, reverse\_cost  
 FROM ways where id <= <SIZE>');  
  
SELECT count(\*) FROM pgr\_johnson(  
 'SELECT gid as id, source, target, cost, reverse\_cost  
 FROM ways where id <= <SIZE>');

The results of this tests are presented as:

SIZE:

is the number of edges given as input.

EDGES:

is the total number of records in the query.

DENSITY:

is the density of the data  $\frac{E}{V \times (V-1)}$ .

OUT ROWS:

is the number of records returned by the queries.

Floyd-Warshall:

is the average execution time in seconds of pgr\_floydWarshall.

Johnson:

is the average execution time in seconds of pgr\_johnson.

SIZE EDGES DENSITY OUT ROWS Floyd-Warshall Johnson

500	500	0.18E-7	1346	0.14	0.13
1000	1000	0.36E-7	2655	0.23	0.18
1500	1500	0.55E-7	4110	0.37	0.34
2000	2000	0.73E-7	5676	0.56	0.37
2500	2500	0.89E-7	7177	0.84	0.51
3000	3000	1.07E-7	8778	1.28	0.68
3500	3500	1.24E-7	10526	2.08	0.95

SIZE EDGES DENSITY OUT ROWS Floyd-Warshall Johnson

4000	4000	1.41E-7	12484	3.16	1.24
4500	4500	1.58E-7	14354	4.49	1.47
5000	5000	1.76E-7	16503	6.05	1.78
5500	5500	1.93E-7	18623	7.53	2.03
6000	6000	2.11E-7	20710	8.47	2.37
6500	6500	2.28E-7	22752	9.99	2.68
7000	7000	2.46E-7	24687	11.82	3.12
7500	7500	2.64E-7	26861	13.94	3.60
8000	8000	2.83E-7	29050	15.61	4.09
8500	8500	3.01E-7	31693	17.43	4.63
9000	9000	3.17E-7	33879	19.19	5.34
9500	9500	3.35E-7	36287	20.77	6.24
10000	10000	3.52E-7	38491	23.26	6.51

Test:

Two

This test is with a bounding box The density of the passed graph higher than of the Test One. For each <SIZE> 30 tests were executed to get the average The tested edge query is:

```
WITH
buffer AS (
  SELECT ST_Buffer(ST_Centroid(ST_Extent(the_geom)), SIZE) AS geom
  FROM ways),
bbox AS (
  SELECT ST_Envelope(ST_Extent(geom)) as box FROM buffer)
SELECT gid as id, source, target, cost, reverse_cost
FROM ways where the_geom && (SELECT box from bbox);
```

The tested queries

```
SELECT count(*) FROM pgr_floydWarshall(<edge query>)
SELECT count(*) FROM pgr_johnson(<edge query>)
```

The results of this tests are presented as:

SIZE:

is the size of the bounding box.

EDGES:

is the total number of records in the query.

DENSITY:

is the density of the data  $\frac{E}{V \times (V-1)}$ .

OUT ROWS:

is the number of records returned by the queries.

Floyd-Warshall:

is the average execution time in seconds of pgr\_floydWarshall.

Johnson:

is the average execution time in seconds of pgr\_johnson.

SIZE EDGES DENSITY OUT ROWS Floyd-Warshall Johnson

0.001	44	0.0608	1197	0.10	0.10
0.002	99	0.0251	4330	0.10	0.10
0.003	223	0.0122	18849	0.12	0.12
0.004	358	0.0085	71834	0.16	0.16
0.005	470	0.0070	116290	0.22	0.19
0.006	639	0.0055	207030	0.37	0.27
0.007	843	0.0043	346930	0.64	0.38
0.008	996	0.0037	469936	0.90	0.49

SIZE EDGES DENSITY OUT ROWS Floyd-Warshall Johnson

0.009	1146	0.0032	613135	1.26	0.62
0.010	1360	0.0027	849304	1.87	0.82
0.011	1573	0.0024	1147101	2.65	1.04
0.012	1789	0.0021	1483629	3.72	1.35
0.013	1975	0.0019	1846897	4.86	1.68
0.014	2281	0.0017	2438298	7.08	2.28
0.015	2588	0.0015	3156007	10.28	2.80
0.016	2958	0.0013	4090618	14.67	3.76
0.017	3247	0.0012	4868919	18.12	4.48

See Also¶

- [pgr\\_johnson](#)
- [pgr\\_floydWarshall](#)
- Boost [floyd-Warshall](#)

Indices and tables

- [Index](#)
- [Search Page](#)

A\* - Family of functions¶

The A\* (pronounced "A Star") algorithm is based on Dijkstra's algorithm with a heuristic that allow it to solve most shortest path problems by evaluation only a sub-set of the overall graph.

- [pgr\\_aStar](#) - A\* algorithm for the shortest path.
- [pgr\\_aStarCost](#) - Get the aggregate cost of the shortest paths.
- [pgr\\_aStarCostMatrix](#) - Get the cost matrix of the shortest paths.

pgr\_aStar¶

pgr\_aStar — Shortest path using the A\* algorithm.

Availability

- Version 3.6.0
  - Standardizing output columns to (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)
    - [pgr\\_aStar\(One to One\)](#) added start\_vid and end\_vid columns.
    - [pgr\\_aStar\(One to Many\)](#) added end\_vid column.
    - [pgr\\_aStar\(Many to One\)](#) added start\_vid column.
- Version 3.2.0
  - New proposed signature:
    - [pgr\\_aStar\(Combinations\)](#)
- Version 3.0.0
  - Function promoted to official.
- Version 2.4.0
  - New proposed signatures:
    - [pgr\\_aStar\(One to Many\)](#)
    - [pgr\\_aStar\(Many to One\)](#)
    - [pgr\\_aStar\(Many to Many\)](#)
- Version 2.3.0
  - Signature change on [pgr\\_aStar\(One to One\)](#)
    - Old signature no longer supported
- Version 2.0.0
  - New official function.

Description¶

The main characteristics are:

- Process works for directed and undirected graphs.
- Ordering is:
  - first by start\_vid (if exists)

- then by `end_vid`
- Values are returned when there is a path.
- Let  $\backslash(v)$  and  $\backslash(u)$  be nodes on the graph:
  - If there is no path from  $\backslash(v)$  to  $\backslash(u)$ :
    - no corresponding row is returned
    - `agg_cost` from  $\backslash(v)$  to  $\backslash(u)$  is  $\backslash(infty)$
  - There is no path when  $\backslash(v = u)$  therefore
    - no corresponding row is returned
    - `agg_cost` from  $v$  to  $u$  is  $\backslash(0)$
- When  $\backslash((x,y))$  coordinates for the same vertex identifier differ:
  - A random selection of the vertex's  $\backslash((x,y))$  coordinates is used.
- Running time:  $\backslash(O((E + V) * \log V))$
- The results are equivalent to the union of the results of the `pgr_aStar(One to One)` on the:
  - `pgr_aStar(One to Many)`
  - `pgr_aStar(Many to One)`
  - `pgr_aStar(Many to Many)`
  - `pgr_aStar(Combinations)`

Boost Graph Inside

Signatures

Summary

`pgr_aStar(Edges SQL, start vid, end vid, [options])`  
`pgr_aStar(Edges SQL, start vid, end vids, [options])`  
`pgr_aStar(Edges SQL, start vids, end vid, [options])`  
`pgr_aStar(Edges SQL, start vids, end vids, [options])`  
`pgr_aStar(Edges SQL, Combinations SQL, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Optional parameters are *named parameters* and have a default value.

One to One

`pgr_aStar(Edges SQL, start vid, end vid, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex  $\backslash(6)$  to vertex  $\backslash(12)$  on a **directed** graph with heuristic  $\backslash(2)$

```
SELECT * FROM pgr_aStar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
6, 12,
directed => true, heuristic => 2);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 12 | 6 | 4 | 1 | 0
2 | 2 | 6 | 12 | 7 | 10 | 1 | 1
3 | 3 | 6 | 12 | 8 | 12 | 1 | 2
4 | 4 | 6 | 12 | 12 | -1 | 0 | 3
(4 rows)
```

One to Many

`pgr_aStar(Edges SQL, start vid, end vids, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex  $\backslash(6)$  to vertices  $\backslash(\{10, 12\})$  on a **directed** graph with heuristic  $\backslash(3)$  and factor  $\backslash(3.5)$

```
SELECT * FROM pgr_aStar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
6, ARRAY[10, 12],
heuristic => 3, factor := 3.5);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
7 | 1 | 6 | 12 | 6 | 4 | 1 | 0
8 | 2 | 6 | 12 | 7 | 8 | 1 | 1
9 | 3 | 6 | 12 | 11 | 11 | 1 | 2
10 | 4 | 6 | 12 | 12 | -1 | 0 | 3
(10 rows)
```

Many to One

`pgr_aStar(Edges SQL, start vids, end vid, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices  $\backslash(\{6, 8\})$  to vertex  $\backslash(10)$  on an **undirected** graph with heuristic  $\backslash(4)$

```
SELECT * FROM pgr_aStar(
  'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
  FROM edges',
  ARRAY[6, 8], 10,
  false, heuristic => 4);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 2 | 1 | 0
2 | 2 | 6 | 10 | 10 | -1 | 0 | 1
3 | 1 | 8 | 10 | 8 | 12 | 1 | 0
4 | 2 | 8 | 10 | 12 | 11 | 1 | 1
5 | 3 | 8 | 10 | 11 | 5 | 1 | 2
6 | 4 | 8 | 10 | 10 | -1 | 0 | 3
(6 rows)
```

Many to Many

pgr\_aStar([Edges SQL](#), start\_vids, end\_vids, [options])  
options: [directed, heuristic, factor, epsilon]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \{(6, 8)\} to vertices \{(10, 12)\} on a **directed** graph with factor \{0.5\}

```
SELECT * FROM pgr_aStar(
  'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
  FROM edges',
  ARRAY[6, 8], ARRAY[10, 12],
  factor => 0.5);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
7 | 1 | 6 | 12 | 6 | 4 | 1 | 0
8 | 2 | 6 | 12 | 7 | 10 | 1 | 1
9 | 3 | 6 | 12 | 8 | 12 | 1 | 2
10 | 4 | 6 | 12 | 12 | -1 | 0 | 3
11 | 1 | 8 | 10 | 8 | 10 | 1 | 0
12 | 2 | 8 | 10 | 7 | 8 | 1 | 1
13 | 3 | 8 | 10 | 11 | 9 | 1 | 2
14 | 4 | 8 | 10 | 16 | 16 | 1 | 3
15 | 5 | 8 | 10 | 15 | 3 | 1 | 4
16 | 6 | 8 | 10 | 10 | -1 | 0 | 5
17 | 1 | 8 | 12 | 8 | 12 | 1 | 0
18 | 2 | 8 | 12 | 12 | -1 | 0 | 1
(18 rows)
```

Combinations

pgr\_aStar([Edges SQL](#), [Combinations SQL](#), [options])  
options: [directed, heuristic, factor, epsilon]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table on a **directed** graph with factor \{0.5\}.

The combinations table:

```
SELECT * FROM combinations;
source | target
-----+-----
5 | 6
5 | 10
6 | 5
6 | 15
6 | 14
(5 rows)
```

The query:

```
SELECT * FROM pgr_aStar(
  'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
  FROM edges',
  'SELECT * FROM combinations',
  factor => 0.5);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 5 | 6 | 5 | 1 | 1 | 0
2 | 2 | 5 | 6 | 6 | -1 | 0 | 1
3 | 1 | 5 | 10 | 5 | 1 | 1 | 0
4 | 2 | 5 | 10 | 6 | 4 | 1 | 1
5 | 3 | 5 | 10 | 7 | 8 | 1 | 2
6 | 4 | 5 | 10 | 11 | 9 | 1 | 3
7 | 5 | 5 | 10 | 16 | 16 | 1 | 4
8 | 6 | 5 | 10 | 15 | 3 | 1 | 5
9 | 7 | 5 | 10 | 10 | -1 | 0 | 6
10 | 1 | 6 | 5 | 6 | 1 | 1 | 0
11 | 2 | 6 | 5 | 5 | -1 | 0 | 1
12 | 1 | 6 | 15 | 6 | 4 | 1 | 0
13 | 2 | 6 | 15 | 7 | 8 | 1 | 1
14 | 3 | 6 | 15 | 11 | 9 | 1 | 2
15 | 4 | 6 | 15 | 16 | 16 | 1 | 3
16 | 5 | 6 | 15 | 15 | -1 | 0 | 4
(16 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start_vid	BIGINT	Identifier of the starting vertex of the path.

Column	Type	Description
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li></ul>
			<ul style="list-style-type: none"><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

sStar optional parameters

Parameter	Type	Default	Description
heuristic	INTEGER	5	Heuristic number. Current valid values 0~5. <ul style="list-style-type: none"><li>0: <math>\backslash(h(v) = 0\backslash)</math> (Use this value to compare with pgr_dijkstra)</li><li>1: <math>\backslash(h(v) = \text{abs}(\text{max}(\Delta x, \Delta y))\backslash)</math></li><li>2: <math>\backslash(h(v) = \text{abs}(\text{min}(\Delta x, \Delta y))\backslash)</math></li><li>3: <math>\backslash(h(v) = \Delta x * \Delta x + \Delta y * \Delta y\backslash)</math></li><li>4: <math>\backslash(h(v) = \text{sqrt}(\Delta x * \Delta x + \Delta y * \Delta y)\backslash)</math></li><li>5: <math>\backslash(h(v) = \text{abs}(\Delta x) + \text{abs}(\Delta y)\backslash)</math></li></ul>
factor	FLOAT	1	For units manipulation. $\backslash(\text{factor} > 0\backslash)$ .
epsilon	FLOAT	1	For less restricted results. $\backslash(\text{epsilon} \geq 1\backslash)$ .

See [heuristics](#) available and [factor](#) handling.

Inner Queries

Edges SQL

Parameter	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target) <ul style="list-style-type: none"><li>When negative: edge (source, target) does not exist, therefore it's not part of the graph.</li></ul>
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source), <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>
x1	ANY-NUMERICAL		X coordinate of source vertex.
y1	ANY-NUMERICAL		Y coordinate of source vertex.
x2	ANY-NUMERICAL		X coordinate of target vertex.
y2	ANY-NUMERICAL		Y coordinate of target vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns<sup>1</sup>

Returns set of (seq, path\_seq [, start\_vid] [, end\_vid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_seq	INTEGER	Relative position in the path. Has value 1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex. Returned when multiple starting vertices are in the query. <ul style="list-style-type: none"><li><a href="#">Many to One</a></li><li><a href="#">Many to Many</a></li></ul>
end_vid	BIGINT	Identifier of the ending vertex. Returned when multiple ending vertices are in the query. <ul style="list-style-type: none"><li><a href="#">One to Many</a></li><li><a href="#">Many to Many</a></li></ul>
node	BIGINT	Identifier of the node in the path from start_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go from node to the next node in the path sequence. -1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Additional Examples<sup>1</sup>

Example 1:

Demonstration of repeated values are ignored, and result is sorted.

```
SELECT * FROM pgr_astar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
ARRAY[7, 10, 15, 10, 10, 15], ARRAY[10, 7, 10, 15]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1	1	7	10	7	8	1	0
2	2	7	10	11	9	1	1
3	3	7	10	16	16	1	2
4	4	7	10	15	3	1	3
5	5	7	10	10	-1	0	4
6	1	7	15	7	8	1	0
7	2	7	15	11	9	1	1
8	3	7	15	16	16	1	2
9	4	7	15	15	-1	0	3
10	1	10	7	10	5	1	0
11	2	10	7	11	8	1	1
12	3	10	7	7	-1	0	2
13	1	10	15	10	5	1	0
14	2	10	15	11	9	1	1
15	3	10	15	16	16	1	2
16	4	10	15	15	-1	0	3
17	1	15	7	15	3	1	0
18	2	15	7	10	2	1	1
19	3	15	7	6	4	1	2
20	4	15	7	7	-1	0	3
21	1	15	10	15	3	1	0
22	2	15	10	10	-1	0	1
(22 rows)							

Example 2:

Making start vids the same as end vids.

```
SELECT * FROM pgr_astar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
ARRAY[7, 10, 15], ARRAY[7, 10, 15]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1	1	7	10	7	8	1	0
2	2	7	10	11	9	1	1
3	3	7	10	16	16	1	2
4	4	7	10	15	3	1	3
5	5	7	10	10	-1	0	4
6	1	7	15	7	8	1	0
7	2	7	15	11	9	1	1
8	3	7	15	16	16	1	2
9	4	7	15	15	-1	0	3
10	1	10	7	10	5	1	0
11	2	10	7	11	8	1	1
12	3	10	7	7	-1	0	2
13	1	10	15	10	5	1	0
14	2	10	15	11	9	1	1

15		3		10		15		16		16		1		2
16		4		10		15		15		-1		0		3
17		1		15		7		15		3		1		0
18		2		15		7		10		2		1		1
19		3		15		7		6		4		1		2
20		4		15		7		7		-1		0		3
21		1		15		10		15		3		1		0
22		2		15		10		10		-1		0		1
(22 rows)														

Example 3:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_aStar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)');
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1		1		6		7		6		4		1		0
2		2		6		7		7		-1		0		1
3		1		6		10		6		4		1		0
4		2		6		10		7		8		1		1
5		3		6		10		11		9		1		2
6		4		6		10		16		16		1		3
7		5		6		10		15		3		1		4
8		6		6		10		10		-1		0		5
9		1		12		10		12		13		1		0
10		2		12		10		17		15		1		1
11		3		12		10		16		16		1		2
12		4		12		10		15		3		1		3
13		5		12		10		10		-1		0		4
(13 rows)														

See Also

- [A\\* - Family of functions](#)
- [Bidirectional A\\* - Family of functions](#)
- [Sample Data](#)
- [Boost: A\\* search](#)
- [https://en.wikipedia.org/wiki/A\\*\\_search\\_algorithm](https://en.wikipedia.org/wiki/A*_search_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_aStarCost

pgr\_aStarCost - Total cost of the shortest path using the A\* algorithm.

Availability

- Version 3.2.0
  - New proposed signature:
    - pgr\_aStarCost(Combinations)
- Version 3.0.0
  - Function promoted to official.
- Version 2.4.0
  - New proposed function.

Description

The pgr\_aStarCost function summarizes of the cost of the shortest path using the A\* algorithm.

The main characteristics are:

- Process works for directed and undirected graphs.
- Ordering is:
  - first by start\_vid (if exists)
  - then by end\_vid
- Values are returned when there is a path.
- Let  $\backslash(v)$  and  $\backslash(u)$  be nodes on the graph:
  - If there is no path from  $\backslash(v)$  to  $\backslash(u)$ :
    - no corresponding row is returned
    - agg\_cost from  $\backslash(v)$  to  $\backslash(u)$  is  $\backslash(\infty)$
  - There is no path when  $\backslash(v = u)$  therefore
    - no corresponding row is returned
    - agg\_cost from  $v$  to  $u$  is  $\backslash(0)$
- When  $\backslash((x,y))$  coordinates for the same vertex identifier differ:
  - A random selection of the vertex's  $\backslash((x,y))$  coordinates is used.
- Running time:  $\backslash(O((E + V) * \log V))$
- It does not return a path.
- Returns the sum of the costs of the shortest path of each pair combination of nodes requested.
- Let be the case the values returned are stored in a table, so the unique index would be the pair(*start\_vid*, *end\_vid*)
- For undirected graphs, the results are symmetric.



- The *agg\_cost* of  $(u, v)$  is the same as for  $(v, u)$ .
- The returned values are ordered in ascending order:
  - *start\_vid* ascending
  - *end\_vid* ascending

Boost Graph Inside

Signatures1

Summary

`pgr_aStarCost(Edges SQL, start vid, end vid, [options])`  
`pgr_aStarCost(Edges SQL, start vid, end vids, [options])`  
`pgr_aStarCost(Edges SQL, start vids, end vid, [options])`  
`pgr_aStarCost(Edges SQL, start vids, end vids, [options])`  
`pgr_aStarCost(Edges SQL, Combinations SQL, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

One to One1

`pgr_aStarCost(Edges SQL, start vid, end vid, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertex \12\ on a **directed** graph with heuristic \2\)

```
SELECT * FROM pgr_aStarCost(
  'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
   FROM edges',
  6, 12,
  directed => true, heuristic => 2);
 start_vid | end_vid | agg_cost
-----+-----+-----
        6 |      12 |         3
(1 row)
```

One to Many1

`pgr_aStarCost(Edges SQL, start vid, end vids, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertices \(\10, 12\\) on a **directed** graph with heuristic \3\ and factor \3.5\)

```
SELECT * FROM pgr_aStarCost(
  'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
   FROM edges',
  6, ARRAY[10, 12],
  heuristic => 3, factor => 3.5);
 start_vid | end_vid | agg_cost
-----+-----+-----
        6 |      10 |         5
        6 |      12 |         3
(2 rows)
```

Many to One1

`pgr_aStarCost(Edges SQL, start vids, end vid, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertices \(\6, 8\\) to vertex \10\ on an **undirected** graph with heuristic \4\)

```
SELECT * FROM pgr_aStarCost(
  'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
   FROM edges',
  ARRAY[6, 8], 10,
  false, heuristic => 4);
 start_vid | end_vid | agg_cost
-----+-----+-----
        6 |      10 |         1
        8 |      10 |         3
(2 rows)
```

Many to Many1

`pgr_aStarCost(Edges SQL, start vids, end vids, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertices \(\6, 8\\) to vertices \(\10, 12\\) on a **directed** graph with factor \0.5\)

```
SELECT * FROM pgr_aStarCost(
  'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
   FROM edges',
  ARRAY[6, 8], ARRAY[10, 12],
  factor => 0.5);
 start_vid | end_vid | agg_cost
-----+-----+-----
        6 |      10 |         5
        6 |      12 |         3
        8 |      10 |         5
        8 |      12 |         1
(4 rows)
```

Combinations1

pgr\_aStarCost([Edges SQL](#), [Combinations SQL](#), [options])  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table on a **directed** graph with factor \0.5).

The combinations table:

```
SELECT * FROM combinations;
source | target
-----+-----
5 | 6
5 | 10
6 | 5
6 | 15
6 | 14
(5 rows)
```

The query:

```
SELECT * FROM pgr_aStarCost(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
'SELECT * FROM combinations',
factor => 0.5);
start_vid | end_vid | agg_cost
-----+-----+-----
5 | 6 | 1
5 | 10 | 6
6 | 5 | 1
6 | 15 | 4
(4 rows)
```

Parameters1

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters1

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

aStar optional parameters1

Parameter	Type	Default	Description
heuristic	INTEGER	5	<p>Heuristic number. Current valid values 0~5.</p> <ul style="list-style-type: none"><li>0: <math>\backslash(h(v) = 0\backslash)</math> (Use this value to compare with <code>pgr_dijkstra</code>)</li><li>1: <math>\backslash(h(v) = \text{abs}(\text{max}(\backslash\Delta x, \backslash\Delta y)\backslash))</math></li><li>2: <math>\backslash(h(v) = \text{abs}(\text{min}(\backslash\Delta x, \backslash\Delta y)\backslash))</math></li><li>3: <math>\backslash(h(v) = \backslash\Delta x * \backslash\Delta x + \backslash\Delta y * \backslash\Delta y\backslash)</math></li><li>4: <math>\backslash(h(v) = \text{sqrt}(\backslash\Delta x * \backslash\Delta x + \backslash\Delta y * \backslash\Delta y)\backslash)</math></li><li>5: <math>\backslash(h(v) = \text{abs}(\backslash\Delta x) + \text{abs}(\backslash\Delta y)\backslash)</math></li></ul>
factor	FLOAT	1	For units manipulation. $\backslash(\text{factor} > 0\backslash)$ .
epsilon	FLOAT	1	For less restricted results. $\backslash(\text{epsilon} >= 1\backslash)$ .

See [heuristics](#) available and [factor](#) handling.

Inner Queries1

Edges SQL1

Parameter	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.

Parameter	Type	Default	Description
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target) <ul style="list-style-type: none"><li>When negative: edge (source, target) does not exist, therefore it's not part of the graph.</li></ul>
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source), <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>
x1	ANY-NUMERICAL		X coordinate of source vertex.
y1	ANY-NUMERICAL		Y coordinate of source vertex.
x2	ANY-NUMERICAL		X coordinate of target vertex.
y2	ANY-NUMERICAL		Y coordinate of target vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations [SQL](#)

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

Additional Examples

Example 1:

Demonstration of repeated values are ignored, and result is sorted.

```
SELECT * FROM pgr_aStarCost(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
ARRAY[7, 10, 15, 10, 15], ARRAY[10, 7, 10, 15]);
start_vid | end_vid | agg_cost
-----+-----+-----
7 | 10 | 4
7 | 15 | 3
10 | 7 | 2
10 | 15 | 3
15 | 7 | 3
15 | 10 | 1
(6 rows)
```

Example 2:

Making **start vids** the same as **end vids**.

```
SELECT * FROM pgr_aStarCost(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
ARRAY[7, 10, 15], ARRAY[7, 10, 15]);
start_vid | end_vid | agg_cost
-----+-----+-----
7 | 10 | 4
7 | 15 | 3
10 | 7 | 2
10 | 15 | 3
15 | 7 | 3
```

15		10		1
(6 rows)				

Example 3:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_aStarCost(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)');
start_vid | end_vid | agg_cost
-----+-----+-----
6 | 7 | 1
6 | 10 | 5
12 | 10 | 4
(3 rows)
```

See Also

- [A\\* - Family of functions](#)
- [Cost - Category](#)
- [Sample Data](#)
- [Boost: A\\* search](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_aStarCostMatrix

pgr\_aStarCostMatrix - Calculates the a cost matrix using [pgr\\_aStar](#).

Availability

- Version 3.0.0
  - Function promoted to official.
- Version 2.4.0
  - New proposed function.

Description

The main characteristics are:

- Using internally the [pgr\\_aStar](#) algorithm
- Returns a cost matrix.
- No ordering is performed
- let  $v$  and  $u$  are nodes on the graph:
  - when there is no path from  $v$  to  $u$ :
    - no corresponding row is returned
    - cost from  $v$  to  $u$  is  $\infty$
  - when  $(v = u)$  then
    - no corresponding row is returned
    - cost from  $v$  to  $u$  is  $0$
- When the graph is **undirected** the cost matrix is symmetric

 Boost Graph Inside

Signatures

Summary

```
pgr_aStarCostMatrix(Edges SQL, start_vids, [options])
options: [directed, heuristic, factor, epsilon]
Returns set of (start_vid, end_vid, agg_cost)
OR EMPTY SET
```

Example:

Symmetric cost matrix for vertices  $\{5, 6, 10, 15\}$  on an **undirected** graph using heuristic  $2$

```
SELECT * FROM pgr_aStarCostMatrix(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2 FROM edges',
(SELECT array_agg(id) FROM vertices WHERE id IN (5, 6, 10, 15)),
directed => false, heuristic => 2);
start_vid | end_vid | agg_cost
-----+-----+-----
5 | 6 | 1
5 | 10 | 2
5 | 15 | 3
6 | 5 | 1
6 | 10 | 1
6 | 15 | 2
10 | 5 | 2
10 | 6 | 1
10 | 15 | 1
15 | 5 | 3
15 | 6 | 2
15 | 10 | 1
(12 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below

**start vids** ARRAY[BIGINT] Array of identifiers of starting vertices.

Optional parameters<sup>1</sup>

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

aStar optional parameters<sup>1</sup>

Parameter	Type	Default	Description
heuristic	INTEGER	5	<p>Heuristic number. Current valid values 0~5.</p> <ul style="list-style-type: none"><li>0: <math>\backslash(h(v) = 0\backslash)</math> (Use this value to compare with <code>pgr_dijkstra</code>)</li><li>1: <math>\backslash(h(v) = \text{abs}(\text{max}(\backslash\Delta x, \backslash\Delta y\backslash))\backslash)</math></li><li>2: <math>\backslash(h(v) = \text{abs}(\text{min}(\backslash\Delta x, \backslash\Delta y\backslash))\backslash)</math></li><li>3: <math>\backslash(h(v) = \backslash\Delta x * \backslash\Delta x + \backslash\Delta y * \backslash\Delta y\backslash)</math></li><li>4: <math>\backslash(h(v) = \text{sqrt}(\backslash\Delta x * \backslash\Delta x + \backslash\Delta y * \backslash\Delta y\backslash))\backslash)</math></li><li>5: <math>\backslash(h(v) = \text{abs}(\Delta x) + \text{abs}(\Delta y)\backslash)</math></li></ul>
factor	FLOAT	1	For units manipulation. $\backslash(\text{factor} > 0\backslash)$ .
epsilon	FLOAT	1	For less restricted results. $\backslash(\text{epsilon} >= 1\backslash)$ .

See [heuristics](#) available and [factor](#) handling.

Inner Queries<sup>1</sup>

Edges SQL<sup>1</sup>

Parameter	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		<p>Weight of the edge (source, target)</p> <ul style="list-style-type: none"><li>When negative: edge (source, target) does not exist, therefore it's not part of the graph.</li></ul>
reverse_cost	ANY-NUMERICAL	-1	<p>Weight of the edge (target, source),</p> <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>
x1	ANY-NUMERICAL		X coordinate of source vertex.
y1	ANY-NUMERICAL		Y coordinate of source vertex.
x2	ANY-NUMERICAL		X coordinate of target vertex.
y2	ANY-NUMERICAL		Y coordinate of target vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns<sup>1</sup>

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
--------	------	-------------

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

Additional Examples¶

Example:

```
Use with pgr_TSP

SELECT * FROM pgr_TSP(
$$
SELECT * FROM pgr_aStarCostMatrix(
  'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2 FROM edges',
  (SELECT array_agg(id) FROM vertices WHERE id IN (5, 6, 10, 15)),
  directed=> false, heuristic => 2)
$$);
NOTICE: pgr_TSP no longer solving with simulated annealing
HINT: Ignoring annealing parameters
seq | node | cost | agg_cost
-----+-----+-----+-----
1 | 5 | 0 | 0
2 | 6 | 1 | 1
3 | 10 | 1 | 2
4 | 15 | 1 | 3
5 | 5 | 3 | 6
(5 rows)
```

See Also¶

- [A\\* - Family of functions](#)
- [Cost Matrix - Category](#)
- [Traveling Sales Person - Family of functions](#)
- [Sample Data](#)
- [Boost: A\\* search](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Description¶

The main Characteristics are:

- Process works for directed and undirected graphs.
- Ordering is:
  - first by start\_vid (if exists)
  - then by end\_vid
- Values are returned when there is a path.
- Let  $\backslash(v)$  and  $\backslash(u)$  be nodes on the graph:
  - If there is no path from  $\backslash(v)$  to  $\backslash(u)$ :
    - no corresponding row is returned
    - agg\_cost from  $\backslash(v)$  to  $\backslash(u)$  is  $\backslash(infty)$
  - There is no path when  $\backslash(v = u)$  therefore
    - no corresponding row is returned
    - agg\_cost from  $v$  to  $u$  is  $\backslash(0)$
- When  $\backslash((x,y))$  coordinates for the same vertex identifier differ:
  - A random selection of the vertex's  $\backslash((x,y))$  coordinates is used.
- Running time:  $\backslash(O((E + V) * \log V))$

aStar optional parameters¶

Parameter	Type	Default	Description
heuristic	INTEGER	5	Heuristic number. Current valid values 0~5. <ul style="list-style-type: none"><li>• 0: <math>\backslash(h(v) = 0)</math> (Use this value to compare with pgr_dijkstra)</li><li>• 1: <math>\backslash(h(v) = \text{abs}(\text{max}(\Delta x, \Delta y)))</math></li><li>• 2: <math>\backslash(h(v) = \text{abs}(\text{min}(\Delta x, \Delta y)))</math></li><li>• 3: <math>\backslash(h(v) = \Delta x * \Delta x + \Delta y * \Delta y)</math></li><li>• 4: <math>\backslash(h(v) = \text{sqrt}(\Delta x * \Delta x + \Delta y * \Delta y))</math></li><li>• 5: <math>\backslash(h(v) = \text{abs}(\Delta x) + \text{abs}(\Delta y))</math></li></ul>
factor	FLOAT	1	For units manipulation. $\backslash(\text{factor} > 0)$ .
epsilon	FLOAT	1	For less restricted results. $\backslash(\text{epsilon} >= 1)$ .

See [heuristics](#) available and [factor](#) handling.

[Advanced documentation¶](#)

[Heuristic¶](#)

Currently the heuristic functions available are:

- 0:  $h(v) = 0$  (Use this value to compare with `pgr_dijkstra`)
- 1:  $h(v) = \text{abs}(\max(\Delta x, \Delta y))$
- 2:  $h(v) = \text{abs}(\min(\Delta x, \Delta y))$
- 3:  $h(v) = \Delta x * \Delta x + \Delta y * \Delta y$
- 4:  $h(v) = \sqrt{\Delta x * \Delta x + \Delta y * \Delta y}$
- 5:  $h(v) = \text{abs}(\Delta x) + \text{abs}(\Delta y)$

where  $\Delta x = x_1 - x_0$  and  $\Delta y = y_1 - y_0$

[Factor¶](#)

Analysis 1

Working with `cost/reverse_cost` as length in degrees, x/y in lat/lon: Factor = 1 (no need to change units)

Analysis 2

Working with `cost/reverse_cost` as length in meters, x/y in lat/lon: Factor = would depend on the location of the points:

Latitude	Conversion	Factor
45	1 longitude degree is 78846.81 m	78846
0	1 longitude degree is 111319.46 m	111319

Analysis 3

Working with `cost/reverse_cost` as time in seconds, x/y in lat/lon: Factor: would depend on the location of the points and on the average speed say 25m/s is the speed.

Latitude	Conversion	Factor
45	1 longitude degree is (78846.81m)/(25m/s)	3153 s
0	1 longitude degree is (111319.46 m)/(25m/s)	4452 s

[See Also¶](#)

- [Bidirectional A\\* - Family of functions](#)
- [Boost: A\\* search](#)
- [https://en.wikipedia.org/wiki/A\\*\\_search\\_algorithm](https://en.wikipedia.org/wiki/A*_search_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

[Bidirectional A\\* - Family of functions¶](#)

The bidirectional A\* (pronounced “A Star”) algorithm is based on the A\* algorithm.

- [pgr\\_bdAstar](#) - Bidirectional A\* algorithm for obtaining paths.
- [pgr\\_bdAstarCost](#) - Bidirectional A\* algorithm to calculate the cost of the paths.
- [pgr\\_bdAstarCostMatrix](#) - Bidirectional A\* algorithm to calculate a cost matrix of paths.

[pgr\\_bdAstar¶](#)

`pgr_bdAstar` — Shortest path using the bidirectional A\* algorithm.

Availability

- Version 3.6.0
  - Standardizing output columns to (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)
    - `pgr_bdAstar(One to One)` added `start_vid` and `end_vid` columns.
    - `pgr_bdAstar(One to Many)` added `end_vid` column.
    - `pgr_bdAstar(Many to One)` added `start_vid` column.
- Version 3.2.0
  - New proposed signature:
    - `pgr_bdAstar(Combinations)`
- Version 3.0.0
  - Function promoted to official.
- Version 2.5.0
  - New proposed signatures:

- `pgr_bdAstar(One to Many)`
  - `pgr_bdAstar(Many to One)`
  - `pgr_bdAstar(Many to Many)`
- Signature change on `pgr_bdAstar(One to One)`
  - Old signature no longer supported
- Version 2.0.0
  - New official function.

Description

The main characteristics are:

- Process works for directed and undirected graphs.
- Ordering is:
  - first by `start_vid` (if exists)
  - then by `end_vid`
- Values are returned when there is a path.
- Let  $\backslash(v)$  and  $\backslash(u)$  be nodes on the graph:
  - If there is no path from  $\backslash(v)$  to  $\backslash(u)$ :
    - no corresponding row is returned
    - `agg_cost` from  $\backslash(v)$  to  $\backslash(u)$  is  $\backslash(infty)$
  - There is no path when  $\backslash(v = u)$  therefore
    - no corresponding row is returned
    - `agg_cost` from  $v$  to  $u$  is  $\backslash(0)$
- When  $\backslash((x,y))$  coordinates for the same vertex identifier differ:
  - A random selection of the vertex's  $\backslash((x,y))$  coordinates is used.
- Running time:  $\backslash(O((E + V) * \log V))$
- The results are equivalent to the union of the results of the `pgr_bdAstar(One to One)` on the:
  - `pgr_bdAstar(One to Many)`
  - `pgr_bdAstar(Many to One)`
  - `pgr_bdAstar(Many to Many)`
  - `pgr_bdAstar(Combinations)`

 Boost Graph Inside

Signatures

Summary

`pgr_bdAstar(Edges SQL, start vid, end vid, [options])`  
`pgr_bdAstar(Edges SQL, start vid, end vids, [options])`  
`pgr_bdAstar(Edges SQL, start vids, end vid, [options])`  
`pgr_bdAstar(Edges SQL, start vids, end vids, [options])`  
`pgr_bdAstar(Edges SQL, Combinations SQL, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns **set** of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Optional parameters are *named parameters* and have a default value.

One to One

`pgr_bdAstar(Edges SQL, start vid, end vid, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns **set** of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex  $\backslash(6)$  to vertex  $\backslash(12)$  on a **directed** graph with heuristic  $\backslash(2)$

```
SELECT * FROM pgr_bdAstar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
6, 12,
directed => true, heuristic => 2
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 12 | 6 | 4 | 1 | 0
2 | 2 | 6 | 12 | 7 | 10 | 1 | 1
3 | 3 | 6 | 12 | 8 | 12 | 1 | 2
4 | 4 | 6 | 12 | 12 | -1 | 0 | 3
(4 rows)
```

One to Many

`pgr_bdAstar(Edges SQL, start vid, end vids, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns **set** of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex  $\backslash(6)$  to vertices  $\backslash(\backslash(10, 12))$  on a **directed** graph with heuristic  $\backslash(3)$  and factor  $\backslash(3.5)$

```
SELECT * FROM pgr_bdAstar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
6, ARRAY[10, 12],
heuristic => 3, factor := 3.5
);
```



```
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
7 | 1 | 6 | 12 | 6 | 4 | 1 | 0
8 | 2 | 6 | 12 | 7 | 8 | 1 | 1
9 | 3 | 6 | 12 | 11 | 11 | 1 | 2
10 | 4 | 6 | 12 | 12 | -1 | 0 | 3
(10 rows)
```

Many to One

`pgr_bdAstar(Edges SQL, start vids, end vid, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices  $\{6, 8\}$  to vertex  $\{10\}$  on an **undirected** graph with heuristic  $\{4\}$

```
SELECT * FROM pgr_bdAstar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
ARRAY[6, 8], 10,
false, heuristic => 4
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 2 | 1 | 0
2 | 2 | 6 | 10 | 10 | -1 | 0 | 1
3 | 1 | 8 | 10 | 8 | 10 | 1 | 0
4 | 2 | 8 | 10 | 7 | 4 | 1 | 1
5 | 3 | 8 | 10 | 6 | 2 | 1 | 2
6 | 4 | 8 | 10 | 10 | -1 | 0 | 3
(6 rows)
```

Many to Many

`pgr_bdAstar(Edges SQL, start vids, end vids, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices  $\{6, 8\}$  to vertices  $\{10, 12\}$  on a **directed** graph with factor  $\{0.5\}$

```
SELECT * FROM pgr_bdAstar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
ARRAY[6, 8], ARRAY[10, 12],
factor => 0.5
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
7 | 1 | 6 | 12 | 6 | 4 | 1 | 0
8 | 2 | 6 | 12 | 7 | 8 | 1 | 1
9 | 3 | 6 | 12 | 11 | 11 | 1 | 2
10 | 4 | 6 | 12 | 12 | -1 | 0 | 3
11 | 1 | 8 | 10 | 8 | 10 | 1 | 0
12 | 2 | 8 | 10 | 7 | 8 | 1 | 1
13 | 3 | 8 | 10 | 11 | 9 | 1 | 2
14 | 4 | 8 | 10 | 16 | 16 | 1 | 3
15 | 5 | 8 | 10 | 15 | 3 | 1 | 4
16 | 6 | 8 | 10 | 10 | -1 | 0 | 5
17 | 1 | 8 | 12 | 8 | 12 | 1 | 0
18 | 2 | 8 | 12 | 12 | -1 | 0 | 1
(18 rows)
```

Combinations

`pgr_bdAstar(Edges SQL, Combinations SQL, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table on a **directed** graph with factor  $\{0.5\}$ .

The combinations table:

```
SELECT * FROM combinations;
source | target
-----+-----
5 | 6
5 | 10
6 | 5
6 | 15
6 | 14
(5 rows)
```

The query:

```
SELECT * FROM pgr_bdAstar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
'SELECT * FROM combinations',
factor => 0.5
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 5 | 6 | 5 | 1 | 1 | 0
2 | 2 | 5 | 6 | 6 | -1 | 0 | 1
3 | 1 | 5 | 10 | 5 | 1 | 1 | 0
4 | 2 | 5 | 10 | 6 | 4 | 1 | 1
5 | 3 | 5 | 10 | 7 | 8 | 1 | 2
6 | 4 | 5 | 10 | 11 | 9 | 1 | 3
```

7	5	5	10	16	16	1	4
8	6	5	10	15	3	1	5
9	7	5	10	10	-1	0	6
10	1	6	5	6	1	1	0
11	2	6	5	5	-1	0	1
12	1	6	15	6	4	1	0
13	2	6	15	7	8	1	1
14	3	6	15	11	9	1	2
15	4	6	15	16	16	1	3
16	5	6	15	15	-1	0	4

(16 rows)

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

aStar optional parameters

Parameter	Type	Default	Description
heuristic	INTEGER	5	<p>Heuristic number. Current valid values 0~5.</p> <ul style="list-style-type: none"><li>0: <math>\backslash(h(v) = 0\backslash)</math> (Use this value to compare with <code>pgr_dijkstra</code>)</li><li>1: <math>\backslash(h(v) = \text{abs}(\text{max}(\backslash\Delta x, \backslash\Delta y\backslash))\backslash)</math></li><li>2: <math>\backslash(h(v) = \text{abs}(\text{min}(\backslash\Delta x, \backslash\Delta y\backslash))\backslash)</math></li><li>3: <math>\backslash(h(v) = \backslash\Delta x * \backslash\Delta x + \backslash\Delta y * \backslash\Delta y\backslash)</math></li><li>4: <math>\backslash(h(v) = \text{sqrt}(\Delta x * \Delta x + \Delta y * \Delta y)\backslash)</math></li><li>5: <math>\backslash(h(v) = \text{abs}(\Delta x) + \text{abs}(\Delta y)\backslash)</math></li></ul>
factor	FLOAT	1	For units manipulation. $\backslash(\text{factor} > 0\backslash)$ .
epsilon	FLOAT	1	For less restricted results. $\backslash(\text{epsilon} >= 1\backslash)$ .

See [heuristics](#) available and [factor](#) handling.

Inner Queries

Edges SQL

Parameter	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		<p>Weight of the edge (source, target)</p> <ul style="list-style-type: none"><li>When negative: edge (source, target) does not exist, therefore it's not part of the graph.</li></ul>
reverse_cost	ANY-NUMERICAL	-1	<p>Weight of the edge (target, source),</p> <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>
x1	ANY-NUMERICAL		X coordinate of source vertex.

Parameter	Type	Default	Description
y1	ANY-NUMERICAL		Y coordinate of source vertex.
x2	ANY-NUMERICAL		X coordinate of target vertex.
y2	ANY-NUMERICAL		Y coordinate of target vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations [SQL](#)

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Returns set of (seq, path\_seq [, start\_vid] [, end\_vid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_seq	INTEGER	Relative position in the path. Has value 1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex. Returned when multiple starting vertices are in the query. <ul style="list-style-type: none"><li><a href="#">Many to One</a></li><li><a href="#">Many to Many</a></li></ul>
end_vid	BIGINT	Identifier of the ending vertex. Returned when multiple ending vertices are in the query. <ul style="list-style-type: none"><li><a href="#">One to Many</a></li><li><a href="#">Many to Many</a></li></ul>
node	BIGINT	Identifier of the node in the path from start_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go from node to the next node in the path sequence. -1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Additional Examples

Example 1:

Demonstration of repeated values are ignored, and result is sorted.

```
SELECT * FROM pgr_bdAstar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
ARRAY[7, 10, 15, 10, 15], ARRAY[10, 7, 10, 15]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1	1	7	10	7	8	1	0
2	2	7	10	11	9	1	1
3	3	7	10	16	16	1	2
4	4	7	10	15	3	1	3
5	5	7	10	10	-1	0	4
6	1	7	15	7	8	1	0
7	2	7	15	11	9	1	1
8	3	7	15	16	16	1	2
9	4	7	15	15	-1	0	3
10	1	10	7	10	5	1	0
11	2	10	7	11	8	1	1
12	3	10	7	7	-1	0	2
13	1	10	15	10	5	1	0
14	2	10	15	11	9	1	1
15	3	10	15	16	16	1	2
16	4	10	15	15	-1	0	3
17	1	15	7	15	3	1	0
18	2	15	7	10	5	1	1
19	3	15	7	11	8	1	2
20	4	15	7	7	-1	0	3
21	1	15	10	15	3	1	0

22		2		15		10		10		-1		0		1
----	--	---	--	----	--	----	--	----	--	----	--	---	--	---

(22 rows)

Example 2:

Making **start vids** the same as **end vids**.

```
SELECT * FROM pgr_bdAstar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
ARRAY[7, 10, 15], ARRAY[7, 10, 15]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1		1		7		10		7		8		1		0
2		2		7		10		11		9		1		1
3		3		7		10		16		16		1		2
4		4		7		10		15		3		1		3
5		5		7		10		10		-1		0		4
6		1		7		15		7		8		1		0
7		2		7		15		11		9		1		1
8		3		7		15		16		16		1		2
9		4		7		15		15		-1		0		3
10		1		10		7		10		5		1		0
11		2		10		7		11		8		1		1
12		3		10		7		7		-1		0		2
13		1		10		15		10		5		1		0
14		2		10		15		11		9		1		1
15		3		10		15		16		16		1		2
16		4		10		15		15		-1		0		3
17		1		15		7		15		3		1		0
18		2		15		7		10		5		1		1
19		3		15		7		11		8		1		2
20		4		15		7		7		-1		0		3
21		1		15		10		15		3		1		0
22		2		15		10		10		-1		0		1

(22 rows)

Example 3:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_bdAstar(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)');
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1		1		6		7		6		4		1		0
2		2		6		7		7		-1		0		1
3		1		6		10		6		4		1		0
4		2		6		10		7		8		1		1
5		3		6		10		11		9		1		2
6		4		6		10		16		16		1		3
7		5		6		10		15		3		1		4
8		6		6		10		10		-1		0		5
9		1		12		10		12		13		1		0
10		2		12		10		17		15		1		1
11		3		12		10		16		16		1		2
12		4		12		10		15		3		1		3
13		5		12		10		10		-1		0		4

(13 rows)

See Also

- [A\\* - Family of functions](#)
- [Bidirectional A\\* - Family of functions](#)
- [Sample Data](#)
- [Boost: A\\* search](#)
- [https://en.wikipedia.org/wiki/A\\*\\_search\\_algorithm](https://en.wikipedia.org/wiki/A*_search_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_bdAstarCost

pgr\_bdAstarCost - Total cost of the shortest path using the bidirectional A\* algorithm.

Availability

- Version 3.2.0
  - New proposed signature:
    - pgr\_bdAstarCost(Combinations)
- Version 3.0.0
  - Function promoted to official.
- Version 2.4.0
  - New proposed function.

Description

The pgr\_bdAstarCost function summarizes of the cost of the shortest path using the bidirectional A\* algorithm.

The main characteristics are:

- Process works for directed and undirected graphs.
- Ordering is:
  - first by start\_vid (if exists)
  - then by end\_vid
- Values are returned when there is a path.
- Let \v) and \u) be nodes on the graph:

- If there is no path from  $v$  to  $u$ :
  - no corresponding row is returned
  - $agg\_cost$  from  $v$  to  $u$  is  $\infty$
- There is no path when  $v = u$  therefore
  - no corresponding row is returned
  - $agg\_cost$  from  $v$  to  $u$  is  $0$
- When  $(x, y)$  coordinates for the same vertex identifier differ:
  - A random selection of the vertex's  $(x, y)$  coordinates is used.
- Running time:  $O((E + V) * \log V)$
- It does not return a path.
- Returns the sum of the costs of the shortest path of each pair combination of nodes requested.
- Let be the case the values returned are stored in a table, so the unique index would be the pair  $(start\_vid, end\_vid)$
- For undirected graphs, the results are symmetric.
  - The  $agg\_cost$  of  $(u, v)$  is the same as for  $(v, u)$ .
- The returned values are ordered in ascending order:
  - $start\_vid$  ascending
  - $end\_vid$  ascending

Boost Graph Inside

#### Signatures

#### Summary

`pgr_bdAstarCost(Edges SQL, start vid, end vid, [options])`  
`pgr_bdAstarCost(Edges SQL, start vid, end vids, [options])`  
`pgr_bdAstarCost(Edges SQL, start vids, end vid, [options])`  
`pgr_bdAstarCost(Edges SQL, start vids, end vids, [options])`  
`pgr_bdAstarCost(Edges SQL, Combinations SQL, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
 Returns set of (start\_vid, end\_vid, agg\_cost)  
 OR EMPTY SET

#### One to One

`pgr_bdAstarCost(Edges SQL, start vid, end vid, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
 Returns set of (start\_vid, end\_vid, agg\_cost)  
 OR EMPTY SET

#### Example:

From vertex  $6$  to vertex  $12$  on a **directed** graph with heuristic  $2$

```
SELECT * FROM pgr_bdAstarCost(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
6, 12,
directed => true, heuristic => 2
);
start_vid | end_vid | agg_cost
-----+-----+-----
6 | 12 | 3
(1 row)
```

#### One to Many

`pgr_bdAstarCost(Edges SQL, start vid, end vids, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
 Returns set of (start\_vid, end\_vid, agg\_cost)  
 OR EMPTY SET

#### Example:

From vertex  $6$  to vertices  $\{10, 12\}$  on a **directed** graph with heuristic  $3$  and factor  $3.5$

```
SELECT * FROM pgr_bdAstarCost(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
6, ARRAY[10, 12],
heuristic => 3, factor := 3.5
);
start_vid | end_vid | agg_cost
-----+-----+-----
6 | 10 | 5
6 | 12 | 3
(2 rows)
```

#### Many to One

`pgr_bdAstarCost(Edges SQL, start vids, end vid, [options])`  
**options:** [directed, heuristic, factor, epsilon]  
 Returns set of (start\_vid, end\_vid, agg\_cost)  
 OR EMPTY SET

#### Example:

From vertices  $\{6, 8\}$  to vertex  $10$  on an **undirected** graph with heuristic  $4$

```
SELECT * FROM pgr_bdAstarCost(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
ARRAY[6, 8], 10,
false, heuristic => 4
);
start_vid | end_vid | agg_cost
-----+-----+-----
6 | 10 | 1
8 | 10 | 3
(2 rows)
```

Many to Many

pgr\_bdAstarCost([Edges SQL](#), start vids, end vids, [options])  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertices \{6, 8\} to vertices \{10, 12\} on a **directed** graph with factor \{0.5\}

```
SELECT * FROM pgr_bdAstarCost(
  'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
   FROM edges',
  ARRAY[6, 8], ARRAY[10, 12],
  factor => 0.5
);
```

start_vid	end_vid	agg_cost
6	10	5
6	12	3
8	10	5
8	12	1

(4 rows)

Combinations

pgr\_bdAstarCost([Edges SQL](#), [Combinations SQL](#), [options])  
**options:** [directed, heuristic, factor, epsilon]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table on a **directed** graph with factor \{0.5\}.

The combinations table:

```
SELECT * FROM combinations;
source | target
```

source	target
5	6
5	10
6	5
6	15
6	14

(5 rows)

The query:

```
SELECT * FROM pgr_bdAstarCost(
  'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
   FROM edges',
  'SELECT * FROM combinations',
  factor => 0.5
);
```

start_vid	end_vid	agg_cost
5	6	1
5	10	6
6	5	1
6	15	4

(4 rows)

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

aStar optional parameters

Parameter	Type	Default	Description
-----------	------	---------	-------------

Parameter	Type	Default	Description
heuristic	INTEGER	5	Heuristic number. Current valid values 0~5. <ul style="list-style-type: none"><li>0: <math>\backslash(h(v) = 0\backslash)</math> (Use this value to compare with <code>pgf_dijkstra</code>)</li><li>1: <math>\backslash(h(v) = \text{abs}(\text{max}(\Delta x, \Delta y))\backslash)</math></li><li>2: <math>\backslash(h(v) = \text{abs}(\text{min}(\Delta x, \Delta y))\backslash)</math></li><li>3: <math>\backslash(h(v) = \Delta x * \Delta x + \Delta y * \Delta y\backslash)</math></li><li>4: <math>\backslash(h(v) = \text{sqrt}(\Delta x * \Delta x + \Delta y * \Delta y)\backslash)</math></li><li>5: <math>\backslash(h(v) = \text{abs}(\Delta x) + \text{abs}(\Delta y)\backslash)</math></li></ul>
factor	FLOAT	1	For units manipulation. $\backslash(\text{factor} > 0\backslash)$ .
epsilon	FLOAT	1	For less restricted results. $\backslash(\text{epsilon} \geq 1\backslash)$ .

See [heuristics](#) available and [factor](#) handling.

Inner Queries<sup>1</sup>

Edges SQL<sup>1</sup>

Parameter	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target) <ul style="list-style-type: none"><li>When negative: edge (source, target) does not exist, therefore it's not part of the graph.</li></ul>
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source), <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>
x1	ANY-NUMERICAL		X coordinate of source vertex.
y1	ANY-NUMERICAL		Y coordinate of source vertex.
x2	ANY-NUMERICAL		X coordinate of target vertex.
y2	ANY-NUMERICAL		Y coordinate of target vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL<sup>1</sup>

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns<sup>1</sup>

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.

Column	Type	Description
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

Additional Examples¶

Example 1:

Demonstration of repeated values are ignored, and result is sorted.

```
SELECT * FROM pgr_bdAstarCost(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
ARRAY[7, 10, 15, 10, 10, 15], ARRAY[7, 10, 15]);
start_vid | end_vid | agg_cost
-----+-----+-----
7 | 10 | 4
7 | 15 | 3
10 | 7 | 2
10 | 15 | 3
15 | 7 | 3
15 | 10 | 1
(6 rows)
```

Example 2:

Making **start vids** the same as **end vids**.

```
SELECT * FROM pgr_bdAstarCost(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
ARRAY[7, 10, 15], ARRAY[7, 10, 15]);
start_vid | end_vid | agg_cost
-----+-----+-----
7 | 10 | 4
7 | 15 | 3
10 | 7 | 2
10 | 15 | 3
15 | 7 | 3
15 | 10 | 1
(6 rows)
```

Example 3:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_bdAstarCost(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2
FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)');
start_vid | end_vid | agg_cost
-----+-----+-----
6 | 7 | 1
6 | 10 | 5
12 | 10 | 4
(3 rows)
```

See Also¶

- [Bidirectional A\\* - Family of functions](#)
- [Cost - Category](#)
- [Sample Data](#)
- [Boost: A\\* search](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_bdAstarCostMatrix¶

pgr\_bdAstarCostMatrix - Calculates the a cost matrix using [pgr\\_aStar](#).

Availability

- Version 3.0.0
  - Function promoted to official.
- Version 2.5.0
  - New proposed function.

Description¶

The main characteristics are:

- Using internally the [pgr\\_bdAstar](#) algorithm
- Returns a cost matrix.
- No ordering is performed
- let  $v$  and  $u$  are nodes on the graph:
  - when there is no path from  $v$  to  $u$ :
    - no corresponding row is returned
    - cost from  $v$  to  $u$  is  $\infty$
  - when  $(v = u)$  then
    - no corresponding row is returned
    - cost from  $v$  to  $u$  is  $(0)$



- When the graph is **undirected** the cost matrix is symmetric

Boost Graph Inside

Signatures

pgr\_bdAstarCostMatrix([Edges SQL](#), start vids, [options])

**options:** [directed, heuristic, factor, epsilon]

Returns set of (start\_vid, end\_vid, agg\_cost)

OR EMPTY SET

Example:

Symmetric cost matrix for vertices  $\{5, 6, 10, 15\}$  on an **undirected** graph using heuristic 2)

```
SELECT * FROM pgr_bdAstarCostMatrix(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2 FROM edges',
(SELECT array_agg(id) FROM vertices WHERE id IN (5, 6, 10, 15)),
directed => false, heuristic => 2
);
 start_vid | end_vid | agg_cost
-----+-----+-----
      5 |      6 |         1
      5 |     10 |         2
      5 |     15 |         3
      6 |      5 |         1
      6 |     10 |         1
      6 |     15 |         2
     10 |      5 |         2
     10 |      6 |         1
     10 |     15 |         1
     15 |      5 |         3
     15 |      6 |         2
     15 |     10 |         1
(12 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below

**start vids**    `ARRAY[BIGINT]`    Array of identifiers of starting vertices.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

aStar optional parameters

Parameter	Type	Default	Description
heuristic	INTEGER	5	Heuristic number. Current valid values 0~5.
			<ul style="list-style-type: none"><li>0: <math>h(v) = 0</math> (Use this value to compare with pgr_dijkstra)</li></ul>
			<ul style="list-style-type: none"><li>1: <math>h(v) = \text{abs}(\text{max}(\Delta x, \Delta y))</math></li></ul>
			<ul style="list-style-type: none"><li>2: <math>h(v) = \text{abs}(\text{min}(\Delta x, \Delta y))</math></li></ul>
			<ul style="list-style-type: none"><li>3: <math>h(v) = \Delta x * \Delta x + \Delta y * \Delta y</math></li></ul>
			<ul style="list-style-type: none"><li>4: <math>h(v) = \text{sqrt}(\Delta x * \Delta x + \Delta y * \Delta y)</math></li></ul>
			<ul style="list-style-type: none"><li>5: <math>h(v) = \text{abs}(\Delta x) + \text{abs}(\Delta y)</math></li></ul>
factor	FLOAT	1	For units manipulation. $(\text{factor} > 0)$ .
epsilon	FLOAT	1	For less restricted results. $(\text{epsilon} \geq 1)$ .

See [heuristics](#) available and [factor](#) handling.

Inner Queries

Edges SQL

Parameter	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
			<ul style="list-style-type: none"><li>When negative: edge (source, target) does not exist, therefore it's not part of the graph.</li></ul>

Parameter	Type	Default	Description
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source), <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>
x1	ANY-NUMERICAL		X coordinate of source vertex.
y1	ANY-NUMERICAL		Y coordinate of source vertex.
x2	ANY-NUMERICAL		X coordinate of target vertex.
y2	ANY-NUMERICAL		Y coordinate of target vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns¶

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

Additional Examples¶

Example:

Use with [pgr\\_TSP](#)

```
SELECT * FROM pgr_TSP(
$$
SELECT * FROM pgr_bdAstarCostMatrix(
'SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2 FROM edges',
(SELECT array_agg(id) FROM vertices WHERE id IN (5, 6, 10, 15)),
directed=> false, heuristic => 2
)
$$
);
NOTICE: pgr_TSP no longer solving with simulated annealing
HINT: Ignoring annealing parameters
seq | node | cost | agg_cost
-----+-----+-----+-----
1 | 5 | 0 | 0
2 | 6 | 1 | 1
3 | 10 | 1 | 2
4 | 15 | 1 | 3
5 | 5 | 3 | 6
(5 rows)
```

See Also¶

- [Bidirectional A\\* - Family of functions](#)
- [Cost Matrix - Category](#)
- [Traveling Sales Person - Family of functions](#)
- [Sample Data](#)
- [Boost: A\\* search](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Description¶

Based on A\* algorithm, the bidirectional search finds a shortest path from a starting vertex (start\_vid) to an ending vertex (end\_vid). It runs two simultaneous searches: one forward from the start\_vid, and one backward from the end\_vid, stopping when the two meet in the middle. This implementation can be used with a directed graph and an undirected graph.

The main Characteristics are:

- Process works for directed and undirected graphs.
- Ordering is:
  - first by start\_vid (if exists)
  - then by end\_vid
- Values are returned when there is a path.
- Let \v and \u be nodes on the graph:
  - If there is no path from \v to \u):
    - no corresponding row is returned

- $\text{agg\_cost}$  from  $v$  to  $u$  is  $\infty$
- There is no path when  $v = u$  therefore
  - no corresponding row is returned
  - $\text{agg\_cost}$  from  $v$  to  $u$  is  $0$
- When  $(x, y)$  coordinates for the same vertex identifier differ:
  - A random selection of the vertex's  $(x, y)$  coordinates is used.
- Running time:  $O((E + V) * \log V)$
- For large graphs where there is a path between the starting vertex and ending vertex:
  - It is expected to terminate faster than `pgr_aStar`

See [heuristics](#) available and [factor](#) handling.

See Also

- [A\\* - Family of functions](#)
- [Boost: A\\* search](#)
- [https://en.wikipedia.org/wiki/A\\*\\_search\\_algorithm](https://en.wikipedia.org/wiki/A*_search_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

#### Bidirectional Dijkstra - Family of functions

- [pgr\\_bdDijkstra](#) - Bidirectional Dijkstra algorithm for the shortest paths.
- [pgr\\_bdDijkstraCost](#) - Bidirectional Dijkstra to calculate the cost of the shortest paths
- [pgr\\_bdDijkstraCostMatrix](#) - Bidirectional Dijkstra algorithm to create a matrix of costs of the shortest paths.

`pgr_bdDijkstra`

`pgr_bdDijkstra` — Returns the shortest path using Bidirectional Dijkstra algorithm.

Availability:

- Version 3.2.0
  - New proposed signature:
    - `pgr_bdDijkstra(Combinations)`
- Version 3.0.0
  - Function promoted to official.
- Version 2.5.0
  - New proposed signatures:
    - `pgr_bdDijkstra(One to Many)`
    - `pgr_bdDijkstra(Many to One)`
    - `pgr_bdDijkstra(Many to Many)`
- Version 2.4.0
  - Signature change on `pgr_bdDijkstra(One to One)`
    - Old signature no longer supported
- Version 2.0.0
  - New official function.

Description

The main characteristics are:

- Process is done only on edges with positive costs.
  - A negative value on a cost column is interpreted as the edge does not exist.
- Values are returned when there is a path.
- When there is no path:
  - When the starting vertex and ending vertex are the same.
    - The **aggregate cost** of the non included values  $(v, v)$  is  $0$
  - When the starting vertex and ending vertex are the different and there is no path:
    - The **aggregate cost** the non included values  $(u, v)$  is  $\infty$
- For optimization purposes, any duplicated value in the starting vertices or on the ending vertices are ignored.
- Running time (worse case scenario):  $O((V \log V + E))$
- For large graphs where there is a path between the starting vertex and ending vertex:
  - It is expected to terminate faster than `pgr_dijkstra`

Boost Graph Inside

Signatures

Summary

pgr\_bdDijkstra([Edges SQL](#), start vid, end vid, [directed])  
pgr\_bdDijkstra([Edges SQL](#), start vid, end vids, [directed])  
pgr\_bdDijkstra([Edges SQL](#), start vids, end vid, [directed])  
pgr\_bdDijkstra([Edges SQL](#), start vids, end vids, [directed])  
pgr\_bdDijkstra([Edges SQL](#), [Combinations SQL](#), [directed])  
Returns set of (seq, path\_seq, [start\_vid], [end\_vid], node, edge, cost, agg\_cost)  
OR EMPTY SET

One to One

pgr\_bdDijkstra([Edges SQL](#), start vid, end vid, [directed])  
Returns set of (seq, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertex \10\ on a **directed** graph

```
SELECT * FROM pgr_bdDijkstra(
'select id, source, target, cost, reverse_cost from edges',
6, 10, true);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 4 | 1 | 0
2 | 2 | 7 | 8 | 1 | 1
3 | 3 | 11 | 9 | 1 | 2
4 | 4 | 16 | 16 | 1 | 3
5 | 5 | 15 | 3 | 1 | 4
6 | 6 | 10 | -1 | 0 | 5
(6 rows)
```

One to Many

pgr\_bdDijkstra([Edges SQL](#), start vid, end vids, [directed])  
Returns set of (seq, path\_seq, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertices \10, 17\ on a **directed** graph

```
SELECT * FROM pgr_bdDijkstra(
'select id, source, target, cost, reverse_cost from edges',
6, ARRAY[10, 17]);
seq | path_seq | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 10 | 6 | 4 | 1 | 0
2 | 2 | 10 | 7 | 8 | 1 | 1
3 | 3 | 10 | 11 | 9 | 1 | 2
4 | 4 | 10 | 16 | 16 | 1 | 3
5 | 5 | 10 | 15 | 3 | 1 | 4
6 | 6 | 10 | 10 | -1 | 0 | 5
7 | 1 | 17 | 6 | 4 | 1 | 0
8 | 2 | 17 | 7 | 8 | 1 | 1
9 | 3 | 17 | 11 | 11 | 1 | 2
10 | 4 | 17 | 12 | 13 | 1 | 3
11 | 5 | 17 | 17 | -1 | 0 | 4
(11 rows)
```

Many to One

pgr\_bdDijkstra([Edges SQL](#), start vids, end vid, [directed])  
Returns set of (seq, path\_seq, start\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \6, 1\ to vertex \17\ on a **directed** graph

```
SELECT * FROM pgr_bdDijkstra(
'select id, source, target, cost, reverse_cost from edges',
ARRAY[6, 1], 17);
seq | path_seq | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | 1 | 6 | 1 | 0
2 | 2 | 1 | 3 | 7 | 1 | 1
3 | 3 | 1 | 7 | 8 | 1 | 2
4 | 4 | 1 | 11 | 11 | 1 | 3
5 | 5 | 1 | 12 | 13 | 1 | 4
6 | 6 | 1 | 17 | -1 | 0 | 5
7 | 1 | 6 | 6 | 4 | 1 | 0
8 | 2 | 6 | 7 | 8 | 1 | 1
9 | 3 | 6 | 11 | 11 | 1 | 2
10 | 4 | 6 | 12 | 13 | 1 | 3
11 | 5 | 6 | 17 | -1 | 0 | 4
(11 rows)
```

Many to Many

pgr\_bdDijkstra([Edges SQL](#), start vids, end vids, [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \6, 1\ to vertices \10, 17\ on an **undirected** graph

```
SELECT * FROM pgr_bdDijkstra(
'select id, source, target, cost, reverse_cost from edges',
ARRAY[6, 1], ARRAY[10, 17],
directed => false);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 10 | 1 | 6 | 1 | 0
2 | 2 | 1 | 10 | 3 | 7 | 1 | 1
3 | 3 | 1 | 10 | 7 | 4 | 1 | 2
4 | 4 | 1 | 10 | 6 | 2 | 1 | 3
5 | 5 | 1 | 10 | 10 | -1 | 0 | 4
6 | 1 | 1 | 17 | 1 | 6 | 1 | 0
7 | 2 | 1 | 17 | 3 | 7 | 1 | 1
8 | 3 | 1 | 17 | 7 | 8 | 1 | 2
9 | 4 | 1 | 17 | 11 | 11 | 1 | 3
10 | 5 | 1 | 17 | 12 | 13 | 1 | 4
11 | 6 | 1 | 17 | 17 | -1 | 0 | 5
12 | 1 | 6 | 10 | 6 | 2 | 1 | 0
13 | 2 | 6 | 10 | 10 | -1 | 0 | 1
14 | 1 | 6 | 17 | 6 | 2 | 1 | 0
15 | 2 | 6 | 17 | 10 | 5 | 1 | 1
16 | 3 | 6 | 17 | 11 | 11 | 1 | 2
```

17	4	6	17	12	13	1	3
18	5	6	17	17	-1	0	4
(18 rows)							

Combinations1

pgr\_bdDijkstra([Edges SQL](#), [Combinations SQL](#), [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:  
Using a combinations table on an **undirected** graph

The combinations table:

```
SELECT source, target FROM combinations;
```

source	target
5	6
5	10
6	5
6	15
6	14

(5 rows)

The query:

```
SELECT * FROM pgr_bdDijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  'SELECT source, target FROM combinations',
  false);
```

seq	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	5	6	5	1	1	0
2	2	5	6	6	-1	0	1
3	1	5	10	5	1	1	0
4	2	5	10	6	2	1	1
5	3	5	10	10	-1	0	2
6	1	6	5	6	1	1	0
7	2	6	5	5	-1	0	1
8	1	6	15	6	2	1	0
9	2	6	15	10	3	1	1
10	3	6	15	15	-1	0	2

(10 rows)

Parameters1

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters1

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries1

Edges SQL1

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:  
ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Returns set of (seq, path\_seq [, start\_vid] [, end\_vid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex. Returned when multiple starting vetrices are in the query. <ul style="list-style-type: none"><li>Many to One</li><li>Many to Many</li></ul>
end_vid	BIGINT	Identifier of the ending vertex. Returned when multiple ending vertices are in the query. <ul style="list-style-type: none"><li>One to Many</li><li>Many to Many</li></ul>
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Additional Examples

Example 1:

Demonstration of repeated values are ignored, and result is sorted.

```
SELECT * FROM pgr_bdDijkstra(
'select id, source, target, cost, reverse_cost from edges',
ARRAY[7, 10, 15, 10, 10, 15], ARRAY[10, 7, 10, 15]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 7 | 10 | 7 | 8 | 1 | 0
2 | 2 | 7 | 10 | 11 | 9 | 1 | 1
3 | 3 | 7 | 10 | 16 | 16 | 1 | 2
4 | 4 | 7 | 10 | 15 | 3 | 1 | 3
5 | 5 | 7 | 10 | 10 | -1 | 0 | 4
6 | 1 | 7 | 15 | 7 | 8 | 1 | 0
7 | 2 | 7 | 15 | 11 | 9 | 1 | 1
8 | 3 | 7 | 15 | 16 | 16 | 1 | 2
9 | 4 | 7 | 15 | 15 | -1 | 0 | 3
10 | 1 | 10 | 7 | 10 | 2 | 1 | 0
11 | 2 | 10 | 7 | 6 | 4 | 1 | 1
12 | 3 | 10 | 7 | 7 | -1 | 0 | 2
13 | 1 | 10 | 15 | 10 | 5 | 1 | 0
14 | 2 | 10 | 15 | 11 | 9 | 1 | 1
15 | 3 | 10 | 15 | 16 | 16 | 1 | 2
16 | 4 | 10 | 15 | 15 | -1 | 0 | 3
17 | 1 | 15 | 7 | 15 | 3 | 1 | 0
18 | 2 | 15 | 7 | 10 | 2 | 1 | 1
19 | 3 | 15 | 7 | 6 | 4 | 1 | 2
20 | 4 | 15 | 7 | 7 | -1 | 0 | 3
21 | 1 | 15 | 10 | 15 | 3 | 1 | 0
22 | 2 | 15 | 10 | 10 | -1 | 0 | 1
(22 rows)
```

Example 2:

Making start vids the same as end vids.

```
SELECT * FROM pgr_bdDijkstra(
'select id, source, target, cost, reverse_cost from edges',
ARRAY[7, 10, 15], ARRAY[7, 10, 15]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 7 | 10 | 7 | 8 | 1 | 0
2 | 2 | 7 | 10 | 11 | 9 | 1 | 1
3 | 3 | 7 | 10 | 16 | 16 | 1 | 2
4 | 4 | 7 | 10 | 15 | 3 | 1 | 3
5 | 5 | 7 | 10 | 10 | -1 | 0 | 4
6 | 1 | 7 | 15 | 7 | 8 | 1 | 0
7 | 2 | 7 | 15 | 11 | 9 | 1 | 1
```

8	3	7	15	16	16	1	2
9	4	7	15	15	-1	0	3
10	1	10	7	10	2	1	0
11	2	10	7	6	4	1	1
12	3	10	7	7	-1	0	2
13	1	10	15	10	5	1	0
14	2	10	15	11	9	1	1
15	3	10	15	16	16	1	2
16	4	10	15	15	-1	0	3
17	1	15	7	15	3	1	0
18	2	15	7	10	2	1	1
19	3	15	7	6	4	1	2
20	4	15	7	7	-1	0	3
21	1	15	10	15	3	1	0
22	2	15	10	10	-1	0	1

(22 rows)

Example 3:

Manually assigned vertex combinations.

```

SELECT * FROM pgr_bdDijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)');
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----

```

1	1	6	7	6	4	1	0
2	2	6	7	7	-1	0	1
3	1	6	10	6	4	1	0
4	2	6	10	7	8	1	1
5	3	6	10	11	9	1	2
6	4	6	10	16	16	1	3
7	5	6	10	15	3	1	4
8	6	6	10	10	-1	0	5
9	1	12	10	12	13	1	0
10	2	12	10	17	15	1	1
11	3	12	10	16	16	1	2
12	4	12	10	15	3	1	3
13	5	12	10	10	-1	0	4

(13 rows)

See Also

- [Bidirectional Dijkstra - Family of functions](#)
- [Sample Data](#)
- [https://en.wikipedia.org/wiki/Bidirectional\\_search](https://en.wikipedia.org/wiki/Bidirectional_search)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_bdDijkstraCost

pgr\_bdDijkstraCost — Returns the shortest path’s cost using Bidirectional Dijkstra algorithm.

Availability

- Version 3.2.0
  - New proposed signature:
    - pgr\_bdDijkstraCost(Combinations)
- Version 3.0.0
  - Function promoted to official.
- Version 2.5.0
  - New proposed function.

Description

The pgr\_bdDijkstraCost function summarizes of the cost of the shortest path using the bidirectional Dijkstra Algorithm.

- Process is done only on edges with positive costs.
  - A negative value on a cost column is interpreted as the edge does not exist.
- Values are returned when there is a path.
- When there is no path:
  - When the starting vertex and ending vertex are the same.
    - The **aggregate cost** of the non included values  $\setminus((v, v))$  is  $\setminus(0)$
  - When the starting vertex and ending vertex are the different and there is no path:
    - The **aggregate cost** the non included values  $\setminus((u, v))$  is  $\setminus(\infty)$
- For optimization purposes, any duplicated value in the starting vertices or on the ending vertices are ignored.
- Running time (worse case scenario):  $\setminus(O((V \setminus \log V + E)))$
- For large graphs where there is a path bewtween the starting vertex and ending vertex:
  - It is expected to terminate faster than pgr\_dijkstra
- It does not return a path.
- Returns the sum of the costs of the shortest path of each pair combination of nodes requested.
- Let be the case the values returned are stored in a table, so the unique index would be the pair(start\_vid, end\_vid).
- Depending on the function and its parameters, the results can be symmetric.
  - The **aggregate cost** of  $\setminus((u, v))$  is the same as for  $\setminus((v, u))$ .
- Any duplicated value in the start or end vertex identifiers are ignored.
- The returned values are ordered:

- start\_vid ascending
- end\_vid ascending

Boost Graph Inside

Signatures¶

Summary

pgr\_bdDijkstraCost([Edges SQL](#), start\_vid, end\_vid , [directed])  
pgr\_bdDijkstraCost([Edges SQL](#), start\_vid, end\_vids, [directed])  
pgr\_bdDijkstraCost([Edges SQL](#), start\_vids, end\_vid , [directed])  
pgr\_bdDijkstraCost([Edges SQL](#), start\_vids, end\_vids, [directed])  
pgr\_bdDijkstraCost([Edges SQL](#), [Combinations SQL](#), [directed])  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

One to One¶

pgr\_bdDijkstraCost([Edges SQL](#), start\_vid, end\_vid , [directed])  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertex \10\ on a **directed** graph

```
SELECT * FROM pgr_bdDijkstraCost(
'SELECT id, source, target, cost, reverse_cost FROM edges',
6, 10, true);
start_vid | end_vid | agg_cost
-----+-----+-----
6 | 10 | 5
(1 row)
```

One to Many¶

pgr\_bdDijkstraCost([Edges SQL](#), start\_vid, end\_vids, [directed])  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertices \10, 17\ on a **directed** graph

```
SELECT * FROM pgr_bdDijkstraCost(
'SELECT id, source, target, cost, reverse_cost FROM edges',
6, ARRAY[10, 17]);
start_vid | end_vid | agg_cost
-----+-----+-----
6 | 10 | 5
6 | 17 | 4
(2 rows)
```

Many to One¶

pgr\_bdDijkstraCost([Edges SQL](#), start\_vids, end\_vid , [directed])  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertices \6, 1\ to vertex \17\ on a **directed** graph

```
SELECT * FROM pgr_bdDijkstraCost(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[6, 1], 17);
start_vid | end_vid | agg_cost
-----+-----+-----
1 | 17 | 5
6 | 17 | 4
(2 rows)
```

Many to Many¶

pgr\_bdDijkstraCost([Edges SQL](#), start\_vids, end\_vids, [directed])  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertices \6, 1\ to vertices \10, 17\ on an **undirected** graph

```
SELECT * FROM pgr_bdDijkstraCost(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[6, 1], ARRAY[10, 17],
directed => false);
start_vid | end_vid | agg_cost
-----+-----+-----
1 | 10 | 4
1 | 17 | 5
6 | 10 | 1
6 | 17 | 4
(4 rows)
```

Combinations¶

pgr\_bdDijkstraCost([Edges SQL](#), [Combinations SQL](#), [directed])  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table on an **undirected** graph

The combinations table:

```
SELECT source, target FROM combinations;
source | target
-----+-----
5 | 6
5 | 10
6 | 5
```



```
6 | 15
6 | 14
(5 rows)
```

The query:

```
SELECT * FROM pgr_bdDijkstraCost(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT source, target FROM combinations',
false);
start_vid | end_vid | agg_cost
-----+-----+-----
5 | 6 | 1
5 | 10 | 2
6 | 5 | 1
6 | 15 | 2
(4 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

Additional Examples¶

Example 1:

Demonstration of repeated values are ignored, and result is sorted.

```
SELECT * FROM pgr_bdDijkstraCost(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[7, 10, 15, 10, 10, 15], ARRAY[10, 7, 10, 15]);
start_vid | end_vid | agg_cost
-----+-----+-----
7 | 10 | 4
7 | 15 | 3
10 | 7 | 2
10 | 15 | 3
15 | 7 | 3
15 | 10 | 1
(6 rows)
```

Example 2:

Making **start vids** the same as **end vids**.

```
SELECT * FROM pgr_bdDijkstraCost(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[7, 10, 15], ARRAY[7, 10, 15]);
start_vid | end_vid | agg_cost
-----+-----+-----
7 | 10 | 4
7 | 15 | 3
10 | 7 | 2
10 | 15 | 3
15 | 7 | 3
15 | 10 | 1
(6 rows)
```

Example 3:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_bdDijkstraCost(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)');
start_vid | end_vid | agg_cost
-----+-----+-----
6 | 7 | 1
6 | 10 | 5
12 | 10 | 4
(3 rows)
```

See Also¶

- [Bidirectional Dijkstra - Family of functions](#)
- [Sample Data](#)
- [https://en.wikipedia.org/wiki/Bidirectional\\_search](https://en.wikipedia.org/wiki/Bidirectional_search)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_bdDijkstraCostMatrix¶

pgr\_bdDijkstraCostMatrix - Calculates a cost matrix using [pgr\\_bdDijkstra](#).

Availability

- Version 3.0.0
  - Function promoted to official.
- Version 2.5.0
  - New proposed function.

Description¶

Using bidirectional Dijkstra algorithm, calculate and return a cost matrix.

- Process is done only on edges with positive costs.
  - A negative value on a cost column is interpreted as the edge does not exist.
- Values are returned when there is a path.
- When there is no path:
  - When the starting vertex and ending vertex are the same.
    - The **aggregate cost** of the non included values  $\setminus((v, v))$  is  $\setminus(0)$
  - When the starting vertex and ending vertex are the different and there is no path:
    - The **aggregate cost** the non included values  $\setminus((u, v))$  is  $\setminus(\infty)$
- For optimization purposes, any duplicated value in the starting vertices or on the ending vertices are ignored.

- Running time (worse case scenario):  $O((V \log V + E))$
- For large graphs where there is a path between the starting vertex and ending vertex:
  - It is expected to terminate faster than `pgr_dijkstra`

The main Characteristics are:

- Can be used as input to [pgr\\_TSP](#).
  - Use directly when the resulting matrix is symmetric and there is no  $\infty$  value.
  - It will be the users responsibility to make the matrix symmetric.
    - By using geometric or harmonic average of the non symmetric values.
    - By using max or min the non symmetric values.
    - By setting the upper triangle to be the mirror image of the lower triangle.
    - By setting the lower triangle to be the mirror image of the upper triangle.
  - It is also the users responsibility to fix an  $\infty$  value.
- Each function works as part of the family it belongs to.
- It does not return a path.
- Returns the sum of the costs of the shortest path for pair combination of nodes in the graph.
- Process is done only on edges with positive costs.
- Values are returned when there is a path.
  - When the starting vertex and ending vertex are the same, there is no path.
    - The aggregate cost in the non included values  $(v, v)$  is 0.
  - When the starting vertex and ending vertex are the different and there is no path.
    - The aggregate cost in the non included values  $(u, v)$  is  $\infty$ .
- Let be the case the values returned are stored in a table:
  - The unique index would be the pair:  $(start\_vid, end\_vid)$ .
- Depending on the function and its parameters, the results can be symmetric.
  - The aggregate cost of  $(u, v)$  is the same as for  $(v, u)$ .
- Any duplicated value in the **start vids** are ignored.
- The returned values are ordered:
  - `start_vid` ascending
  - `end_vid` ascending

 Boost Graph Inside

Signatures

Summary

`pgr_bdDijkstraCostMatrix(Edges SQL, start vids, [directed])`  
Returns set of  $(start\_vid, end\_vid, agg\_cost)$   
OR EMPTY SET

Example:

Symmetric cost matrix for vertices  $\{5, 6, 10, 15\}$  on an **undirected** graph

```
SELECT * FROM pgr_bdDijkstraCostMatrix(
'SELECT id, source, target, cost, reverse_cost FROM edges',
(SELECT array_agg(id)
 FROM vertices
 WHERE id IN (5, 6, 10, 15)),
false);
start_vid | end_vid | agg_cost
-----+-----+-----
5 | 6 | 1
5 | 10 | 2
5 | 15 | 3
6 | 5 | 1
6 | 10 | 1
6 | 15 | 2
10 | 5 | 2
10 | 6 | 1
10 | 15 | 1
15 | 5 | 3
15 | 6 | 2
15 | 10 | 1
(12 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<b>start vids</b>	ARRAY[BIGINT]	Array of identifiers of starting vertices.

Optional parameters

Column	Type	Default	Description
--------	------	---------	-------------

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

Additional Examples

Example:

Use with [pgr\\_TSP](#).

```
SELECT * FROM pgr_TSP(
$$
SELECT * FROM pgr_bdDijkstraCostMatrix(
'SELECT id, source, target, cost, reverse_cost FROM edges',
(SELECT array_agg(id)
FROM vertices
WHERE id IN (5, 6, 10, 15)),
false)
$$);
NOTICE: pgr_TSP no longer solving with simulated annealing
HINT: Ignoring annealing parameters
seq | node | cost | agg_cost
-----+-----+-----+-----
1 | 5 | 0 | 0
2 | 6 | 1 | 1
3 | 10 | 1 | 2
4 | 15 | 1 | 3
5 | 5 | 3 | 6
(5 rows)
```

See Also

- [Bidirectional Dijkstra - Family of functions](#)
- [Cost Matrix - Category](#)
- [Traveling Sales Person - Family of functions](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Synopsis

Based on Dijkstra's algorithm, the bidirectional search finds a shortest path a starting vertex to an ending vertex.

It runs two simultaneous searches: one forward from the source, and one backward from the target, stopping when the two meet in the middle.

This implementation can be used with a directed graph and an undirected graph.

Characteristics

The main Characteristics are:

- Process is done only on edges with positive costs.
  - A negative value on a cost column is interpreted as the edge does not exist.
- Values are returned when there is a path.
- When there is no path:
  - When the starting vertex and ending vertex are the same.
    - The **aggregate cost** of the non included values  $\backslash((v, v)\backslash)$  is  $\backslash(0\backslash)$
  - When the starting vertex and ending vertex are the different and there is no path:
    - The **aggregate cost** the non included values  $\backslash((u, v)\backslash)$  is  $\backslash(\infty\backslash)$
- For optimization purposes, any duplicated value in the starting vertices or on the ending vertices are ignored.
- Running time (worse case scenario):  $\backslash(O((V \log V + E))\backslash)$
- For large graphs where there is a path bewtween the starting vertex and ending vertex:
  - It is expected to terminate faster than `pgr_dijkstra`

See Also

Indices and tables

- [Index](#)
- [Search Page](#)

Components - Family of functions

- [pgr\\_connectedComponents](#) - Connected components of an undirected graph.
- [pgr\\_strongComponents](#) - Strongly connected components of a directed graph.
- [pgr\\_biconnectedComponents](#) - Biconnected components of an undirected graph.
- [pgr\\_articulationPoints](#) - Articulation points of an undirected graph.
- [pgr\\_bridges](#) - Bridges of an undirected graph.

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting
- [pgr\\_makeConnected - Experimental](#) - Details of edges to make graph connected.

`pgr_connectedComponents`

`pgr_connectedComponents` — Connected components of an undirected graph using a DFS-based approach.

Availability

- Version 3.0.0
  - Result columns change:
    - `n_seq` is removed
    - `seq` changed type to BIGINT
  - Function promoted to official.
- Version 2.5.0
  - New experimental function.

Description

A connected component of an undirected graph is a set of vertices that are all reachable from each other.

The main characteristics are:

- Works for **undirected** graphs.
- Components are described by vertices
- The returned values are ordered:
  - component ascending
  - node ascending
- Running time:  $\backslash(O(V + E)\backslash)$

Boost Graph Inside

Signatures

pgr\_connectedComponents([Edges SQL](#))  
Returns set of (seq, component, node)  
OR EMPTY SET

Example:

The connected components of the graph

```
SELECT * FROM pgr_connectedComponents(
'SELECT id, source, target, cost, reverse_cost FROM edges'
);
seq | component | node
-----+-----+-----
1 | 1 | 1
2 | 1 | 3
3 | 1 | 5
4 | 1 | 6
5 | 1 | 7
6 | 1 | 8
7 | 1 | 9
8 | 1 | 10
9 | 1 | 11
10 | 1 | 12
11 | 1 | 15
12 | 1 | 16
13 | 1 | 17
14 | 2 | 2
15 | 2 | 4
16 | 13 | 13
17 | 13 | 14
(17 rows)
```

[images/cc\\_sampledata.png](#)

Parameters

Parameter Type	Description
----------------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, component, node)

Column	Type	Description
--------	------	-------------

Column	Type	Description
seq	BIGINT	Sequential value starting from 1.
Component identifier.		
component	BIGINT	<ul style="list-style-type: none"><li>Has the value of the minimum node identifier in the component.</li></ul>
node	BIGINT	Identifier of the vertex that belongs to the component.

Additional Examples

Connecting disconnected components

To get the graph connectivity:

```
SELECT * FROM pgr_connectedComponents(
'SELECT id, source, target, cost, reverse_cost FROM edges'
);
seq | component | node
-----+-----+-----
1 | 1 | 1
2 | 1 | 3
3 | 1 | 5
4 | 1 | 6
5 | 1 | 7
6 | 1 | 8
7 | 1 | 9
8 | 1 | 10
9 | 1 | 11
10 | 1 | 12
11 | 1 | 15
12 | 1 | 16
13 | 1 | 17
14 | 2 | 2
15 | 2 | 4
16 | 13 | 13
17 | 13 | 14
(17 rows)
```

There are three basic ways to connect components:

- From the vertex to the starting point of the edge
- From the vertex to the ending point of the edge
- From the vertex to the closest vertex on the edge
  - This solution requires the edge to be split.

In this example [pgr\\_separateCrossing](#) and [pgr\\_separateTouching](#) will be used.

Get the connectivity

```
SELECT * FROM pgr_connectedComponents(
'SELECT id, source, target, cost, reverse_cost FROM edges'
);
seq | component | node
-----+-----+-----
1 | 1 | 1
2 | 1 | 3
3 | 1 | 5
4 | 1 | 6
5 | 1 | 7
6 | 1 | 8
7 | 1 | 9
8 | 1 | 10
9 | 1 | 11
10 | 1 | 12
11 | 1 | 15
12 | 1 | 16
13 | 1 | 17
14 | 2 | 2
15 | 2 | 4
16 | 13 | 13
17 | 13 | 14
(17 rows)
```

Prepare tables

In this example: the edges table will need an additional column and the vertex table will be rebuilt completely.

```
ALTER TABLE edges ADD old_id BIGINT;
ALTER TABLE
DROP TABLE vertices;
DROP TABLE
```

Insert new edges

Using [pgr\\_separateCrossing](#) and [pgr\\_separateTouching](#) insert the results into the edges table.

```
INSERT INTO edges (old_id, geom)
SELECT id, geom FROM pgr_separateCrossing('SELECT * FROM edges')
UNION
SELECT id, geom FROM pgr_separateTouching('SELECT * FROM edges');
INSERT 0 6
```

Create the vertices table

Using [pgr\\_extractVertices](#) create the table.

```
CREATE TABLE vertices AS
SELECT * FROM pgr_extractVertices('SELECT id, geom FROM edges');
SELECT 18
```

Update the topology

```
/* -- set the source information */
UPDATE edges AS e
SET source = v.id, x1 = x, y1 = y
FROM vertices AS v
WHERE ST_StartPoint(e.geom) = v.geom;
UPDATE 24
/* -- set the target information */
UPDATE edges AS e
SET target = v.id, x2 = x, y2 = y
```

```
FROM vertices AS v
WHERE ST_EndPoint(e.geom) = v.geom;
UPDATE 24
```

Update other values

In this example only cost and reverse\_cost are updated, where they are based on the length of the geometry and the directionality is kept using the sign function.

```
UPDATE edges e
SET cost      = ST_length(e.geom)*sign(e1.cost),
    reverse_cost = ST_length(e.geom)*sign(e1.reverse_cost)
FROM edges e1
WHERE e.cost IS NULL AND e1.id = e.old_id;
UPDATE 6
```

```
SELECT * FROM pgr_connectedComponents(
'SELECT id, source, target, cost, reverse_cost FROM edges'
);
```

seq | component | node

1	1	1
2	1	2
3	1	3
4	1	4
5	1	5
6	1	6
7	1	7
8	1	8
9	1	9
10	1	10
11	1	11
12	1	12
13	1	13
14	1	14
15	1	15
16	1	16
17	1	17
18	1	18

(18 rows)

See Also

- [Components - Family of functions](#)
- [Sample Data](#)
- [Boost: Connected components](#)
- wikipedia: [Connected component](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_strongComponents

pgr\_strongComponents — Strongly connected components of a directed graph using Tarjan's algorithm based on DFS.

Availability

- Version 3.0.0
  - Result columns change:
    - n\_seq is removed
    - seq changed type to BIGINT
  - Function promoted to official.
- Version 2.5.0
  - New experimental function.

Description

A strongly connected component of a directed graph is a set of vertices that are all reachable from each other.

The main characteristics are:

- Works for **directed** graphs.
- Components are described by vertices identifiers.
- The returned values are ordered:
  - component ascending
  - node ascending
- Running time:  $\mathcal{O}(V + E)$

Boost Graph Inside

Signatures

pgr\_strongComponents([Edges SQL](#))  
Returns set of (seq, component, node)  
OR EMPTY SET

Example:

The strong components of the graph

```
SELECT * FROM pgr_strongComponents(
'SELECT id, source, target, cost, reverse_cost FROM edges'
);
```

seq | component | node

1	1	1
2	1	3
3	1	5
4	1	6
5	1	7
6	1	8
7	1	9



8		1		10
9		1		11
10		1		12
11		1		15
12		1		16
13		1		17
14		2		2
15		2		4
16		13		13
17		13		14

(17 rows)



Parameters

Parameter	Type	Description
-----------	------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, component, node)

Column	Type	Description
seq	BIGINT	Sequential value starting from 1.
component	BIGINT	Component identifier. <ul style="list-style-type: none"><li>Has the value of the minimum node identifier in the component.</li></ul>
node	BIGINT	Identifier of the vertex that belongs to the component.

See Also

- [Components - Family of functions](#)
- [Sample Data](#)
- [Boost: Strong components](#)
- wikipedia: [Strongly connected component](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_biconnectedComponents

pgr\_biconnectedComponents — Biconnected components of an undirected graph.

Availability

- Version 3.0.0
  - Result columns change:
    - n\_seq is removed
    - seq changed type to BIGINT
  - Function promoted to official.
- Version 2.5.0
  - New experimental function.

Description

The biconnected components of an undirected graph are the maximal subsets of vertices such that the removal of a vertex from particular component will not disconnect the component. Unlike connected components, vertices may belong to multiple biconnected components. Vertices can be present in multiple biconnected components, but each edge can only be contained in a single biconnected component.

The main characteristics are:

- Works for **undirected** graphs.
- Components are described by edges.
- The returned values are ordered:
  - component ascending.
  - edge ascending.
- Running time:  $\backslash(O(V + E)\backslash)$

 Boost Graph Inside

Signatures

pgr\_biconnectedComponents([Edges SQL](#))  
Returns set of (seq, component, edge)  
OR EMPTY SET

Example:

The biconnected components of the graph

```
SELECT * FROM pgr_biconnectedComponents(
'SELECT id, source, target, cost, reverse_cost FROM edges'
);
seq | component | edge
-----+-----+-----
1 | 1 | 1
2 | 2 | 2
3 | 2 | 3
4 | 2 | 4
5 | 2 | 5
6 | 2 | 8
7 | 2 | 9
8 | 2 | 10
9 | 2 | 11
10 | 2 | 12
11 | 2 | 13
12 | 2 | 15
13 | 2 | 16
14 | 6 | 6
15 | 7 | 7
16 | 14 | 14
17 | 17 | 17
18 | 18 | 18
(18 rows)
```



Parameters

Parameter Type	Description
<a href="#">Edges SQL</a> TEXT	<a href="#">Edges SQL</a> as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)

Column	Type	Default	Description
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, component, edge)

Column	Type	Description
seq	BIGINT	Sequential value starting from 1.
component	BIGINT	Component identifier. <ul style="list-style-type: none"> <li>Has the value of the minimum edge identifier in the component.</li> </ul>
edge	BIGINT	Identifier of the edge that belongs to the component.

See Also

- Components - Family of functions
- Sample Data
- Boost: Biconnected components & articulation points
- wikipedia: Biconnected component

Indices and tables

- Index
- Search Page

pgr\_articulationPoints

pgr\_articulationPoints - Return the articulation points of an undirected graph.

Availability

- Version 3.0.0
  - Result columns change: seq is removed
  - Function promoted to official.
- Version 2.5.0
  - New experimental function.

Description

Those vertices that belong to more than one biconnected component are called articulation points or, equivalently, cut vertices. Articulation points are vertices whose removal would increase the number of connected components in the graph. This implementation can only be used with an undirected graph.

The main characteristics are:

- Works for **undirected** graphs.
- The returned values are ordered:
  - node ascending
- Running time:  $\backslash(O(V + E)\backslash)$

Boost Graph Inside

Signatures

pgr\_articulationPoints([Edges SQL](#))

Returns set of (node)

OR EMPTY SET

Example:

The articulation points of the graph

```
SELECT * FROM pgr_articulationPoints(
'SELECT id, source, target, cost, reverse_cost FROM edges'
);
node
-----
 3
 6
 7
 8
(4 rows)
```

Nodes in red are the articulation points.



Parameters

Parameter	Type	Description
-----------	------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (node)

Column	Type	Description
--------	------	-------------

node	BIGINT	Identifier of the vertex.
------	--------	---------------------------

See Also

- [Components - Family of functions](#)
- [Sample Data](#)
- [Boost: Biconnected components & articulation points](#)
- wikipedia: [Biconnected component](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_bridges

pgr\_bridges - Return the bridges of an undirected graph.

Availability

- Version 3.0.0
  - Result columns change: seq is removed
  - Function promoted to official.
- Version 2.5.0
  - New experimental function.

Description

A bridge is an edge of an undirected graph whose deletion increases its number of connected components. This implementation can only be used with an undirected graph.

The main characteristics are:

- Works for **undirected** graphs.

- The returned values are ordered:
  - edge ascending
- Running time:  $\mathcal{O}(E * (V + E))$

 Boost Graph Inside

Signatures

pgr\_bridges([Edges SQL](#))  
Returns set of (edge)  
OR EMPTY SET

Example:

The bridges of the graph

```
SELECT * FROM pgr_bridges(
'SELECT id, source, target, cost, reverse_cost FROM edges'
);
edge
-----
1
6
7
14
17
18
(6 rows)
```



Parameters

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (edge)

Column	Type	Description
edge	BIGINT	Identifier of the edge that is a bridge.

See Also

- [https://en.wikipedia.org/wiki/Bridge\\_%28graph\\_theory%29](https://en.wikipedia.org/wiki/Bridge_%28graph_theory%29)
- [Sample Data](#)
- [Boost: Connected components](#)

Indices and tables

- [Index](#)

- [Search Page](#)

pg\_r\_makeConnected - Experimental

pg\_r\_makeConnected — Set of edges that will connect the graph.

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.2.0
  - New experimental function.

Description

Adds the minimum number of edges needed to make the input graph connected. The algorithm first identifies all of the connected components in the graph, then adds edges to connect those components together in a path. For example, if a graph contains three connected components A, B, and C, make\_connected will add two edges. The two edges added might consist of one connecting a vertex in A with a vertex in B and one connecting a vertex in B with a vertex in C.

The main characteristics are:

- Works for **undirected** graphs.
- It will give a minimum list of all edges which are needed in the graph to make connect it.
- The algorithm does not considers traversal costs in the calculations.
- The algorithm does not considers geometric topology in the calculations.
- Running time:  $\backslash(O(V + E))$

Boost Graph Inside

Signatures

pg\_r\_makeConnected([Edges SQL](#))

Returns set of (seq, start\_vid, end\_vid)

OR EMPTY SET

Example:

List of edges that are needed to connect the graph.

```
SELECT * FROM pg_r_makeConnected(
'SELECT id, source, target, cost, reverse_cost FROM edges'
);
seq | start_vid | end_vid
-----+-----+-----
1 | 5 | 2
2 | 4 | 13
(2 rows)
```

Parameters

Parameter	Type	Description
-----------	------	-------------

<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below.
---------------------------	------	---

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.

Column	Type	Default	Description
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns<sup>1</sup>

Returns set of (seq, start\_vid, end\_vid)

Column	Type	Description
seq	BIGINT	Sequential value starting from 1.
start_vid	BIGINT	Identifier of the first end point vertex of the edge.
end_vid	BIGINT	Identifier of the second end point vertex of the edge.

See Also<sup>1</sup>

- [Boost: make connected](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

See Also<sup>1</sup>

Indices and tables

- [Index](#)
- [Search Page](#)

Contraction - Family of functions<sup>1</sup>

- [pgr\\_contraction](#)

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.
- [pgr\\_contractionDeadEnd - Proposed](#)
- [pgr\\_contractionLinear - Proposed](#)

☐ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:

- The functions might not make use of ANY-INTEGGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting
- [pgr\\_contractionHierarchies - Experimental](#)

pgr\_contraction

pgr\_contraction — Performs graph contraction and returns the contracted vertices and edges.

Availability

Version 3.8.0

- New signature:
  - Previously compulsory parameter **Contraction order** is now optional with name `methods`.
  - New name and order of optional parameters.
- Deprecated signature `pgr_contraction(text,bigint[],integer,bigint[],boolean)`

Version 3.0.0

- Result columns change: seq is removed
- Name change from `pgr_contractGraph`
- Bug fixes
- Function promoted to official.

Version 2.3.0

- New experimental function.

Description

Contraction reduces the size of the graph by removing some of the vertices and edges and, for example, might add edges that represent a sequence of original edges decreasing the total time and space used in graph algorithms.

The main Characteristics are:

- Process is done only on edges with positive costs.
- Does not return the full contracted graph.
  - Only changes on the graph are returned.
- The returned values include:
  - The new edges generated by linear contraction.
  - The modified vertices generated by dead end contraction.
- The returned values are ordered as follows:
  - column `id` ascending when its a modified vertex.
  - column `id` with negative numbers descending when its a new edge.
- Currently there are two types of contraction methods included in this function:
  - Dead End Contraction. See [pgr\\_contractionDeadEnd - Proposed](#)
  - Linear Contraction. See [pgr\\_contractionLinear - Proposed](#).

Boost Graph Inside

Signatures

pgr\_contraction([Edges SQL](#), **[options]**)

**options:** [directed, methods, cycles, forbidden]

Returns set of (type, id, contracted\_vertices, source, target, cost)

Example:

Dead end and linear contraction in that order on an undirected graph.

```
SELECT * FROM pgr_contraction(
  'SELECT id, source, target, cost, reverse_cost FROM edges', false);
type | id | contracted_vertices | source | target | cost
```

v	4	{2}		-1	-1	-1
v	7	{1,3}		-1	-1	-1
v	14	{13}		-1	-1	-1
e	-1	{5,6}	7	10	2	
e	-2	{8,9}	7	12	2	
e	-3	{17}	12	16	2	
e	-4	{15}	10	16	2	
(7 rows)						

(7 rows)

Parameters



Parameter Type	Description
<a href="#">Edges SQL</a> TEXT	<a href="#">Edges SQL</a> as described below.

Optional parameters¶

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Contraction optional parameters¶

Column	Type	Default	Description
			Ordered contraction operations.
methods	INTEGER[]	ARRAY[1,2]	<ul style="list-style-type: none"><li>1 = Dead end contraction</li><li>2 = Linear contraction</li></ul>
cycles	INTEGER	\(1\)	Number of times the contraction methods will be performed.
forbidden	BIGINT[]	ARRAY[]::BIGINT[]	Identifiers of vertices forbidden for contraction.

Inner Queries¶

Edges SQL¶

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns¶

Returns set of (type, id, contracted\_vertices, source, target, cost)

The function returns a single row. The columns of the row are:

Column	Type	Description
		Type of the row. <ul style="list-style-type: none"><li>v when the row is a vertex.</li></ul>
type	TEXT	<ul style="list-style-type: none"><li>Column id has a positive value.</li><li>e when the row is an edge.<ul style="list-style-type: none"><li>Column id has a negative value.</li></ul></li></ul>
		All numbers on this column are DISTINCT <ul style="list-style-type: none"><li>When type = 'v'.<ul style="list-style-type: none"><li>Identifier of the modified vertex.</li></ul></li></ul>
id	BIGINT	<ul style="list-style-type: none"><li>When type = 'e'.<ul style="list-style-type: none"><li>Decreasing sequence starting from -1.</li><li>Representing a pseudo id as is not incorporated in the set of original edges.</li></ul></li></ul>

Column	Type	Description
--------	------	-------------

contracted\_vertices ARRAY[BIGINT] Array of contracted vertex identifiers.

source	BIGINT	<ul style="list-style-type: none"> <li>When type = 'v': \(-1\)</li> <li>When type = 'e': Identifier of the source vertex of the current edge <math>\langle</math>source, target<math>\rangle</math>.</li> </ul>
target	BIGINT	<ul style="list-style-type: none"> <li>When type = 'v': \(-1\)</li> <li>When type = 'e': Identifier of the target vertex of the current edge <math>\langle</math>source, target<math>\rangle</math>.</li> </ul>
cost	FLOAT	<ul style="list-style-type: none"> <li>When type = 'v': \(-1\)</li> <li>When type = 'e': Weight of the current edge <math>\langle</math>source, target<math>\rangle</math>.</li> </ul>

Additional Examples¶

- [Only dead end contraction](#)
- [Only linear contraction](#)
- [The cycle](#)

Only dead end contraction¶

```
SELECT type, id, contracted_vertices FROM pgr_contraction(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  methods => ARRAY[1]);
type | id | contracted_vertices
-----+-----+-----
v    | 4  | {2}
v    | 6  | {5}
v    | 7  | {1,3}
v    | 8  | {9}
v    | 14 | {13}
(5 rows)
```

Only linear contraction¶

```
SELECT * FROM pgr_contraction(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  methods => ARRAY[2]);
type | id | contracted_vertices | source | target | cost
-----+-----+-----+-----+-----+-----
e    | -1 | {3}                | 1      | 7      | 2
e    | -2 | {3}                | 7      | 1      | 2
(2 rows)
```

The cycle¶

Contracting a graph can be done with more than one operation. The order of the operations affect the resulting contracted graph, after applying one operation, the set of vertices that can be contracted by another operation changes.

This implementation cycles  $cycles$  times through the  $methods$  .

```
<input>
do max_cycles times {
  for (operation in operations_order)
  { do operation }
}
<output>
```

Contracting sample data¶

In this section, building and using a contracted graph will be shown by example.

- The [Sample Data](#) for an undirected graph is used
- a dead end operation first followed by a linear operation.
- [Construction of the graph in the database](#)

Construction of the graph in the database¶

The original graph:



The results do not represent the contracted graph. They represent the changes that need to be done to the graph after applying the contraction methods.

Observe that vertices, for example, \{6\} do not appear in the results because it was not affected by the contraction algorithm.

```
SELECT * FROM pgr_contraction(
  'SELECT id, source, target, cost, reverse_cost FROM edges', false);
type | id | contracted_vertices | source | target | cost
-----+-----+-----+-----+-----+-----
v    | 4  | {2}                | -1     | -1     | -1
v    | 7  | {1,3}              | -1     | -1     | -1
v    | 14 | {13}               | -1     | -1     | -1
v    | -1 | {5,6}              | 7      | 10     | 2
e    | -2 | {8,9}              | 7      | 12     | 2
e    | -3 | {17}               | 12     | 16     | 2
e    | -4 | {15}               | 10     | 16     | 2
(7 rows)
```

After doing the dead end contraction operation:



After doing the linear contraction operation to the graph above:



The process to create the contraction graph on the database [¶](#)

- [Add additional columns](#)
- [Store contraction information](#)
- [Update the edges and vertices tables](#)
- [The contracted graph](#)

[Add additional columns¶](#)

Adding extra columns to the edges and vertices tables. In this documentation the following will be used:

Column.	Description
contracted_vertices	The vertices set belonging to the vertex/edge
On the vertex table	
is_contracted	<ul style="list-style-type: none"><li>• when true the vertex is contracted, its not part of the contracted graph.</li><li>• when false the vertex is not contracted, its part of the contracted graph.</li></ul>
On the edge table	
is_new	<ul style="list-style-type: none"><li>• when true the edge was generated by the contraction algorithm. its part of the contracted graph.</li><li>• when false the edge is an original edge, might be or not part of the contracted graph.</li></ul>

```
ALTER TABLE vertices
ADD is_contracted BOOLEAN DEFAULT false,
ADD contracted_vertices BIGINT[];
ALTER TABLE
ALTER TABLE edges
ADD is_new BOOLEAN DEFAULT false,
ADD contracted_vertices BIGINT[];
ALTER TABLE
```

[Store contraction information¶](#)

Store the contraction results in a table.

```
SELECT * INTO contraction_results
FROM pgr_contraction(
'SELECT id, source, target, cost, reverse_cost FROM edges', false);
SELECT 7
```

[Update the edges and vertices tables¶](#)

Use is\_contracted column to indicate the vertices that are contracted.

```
UPDATE vertices
SET is_contracted = true
WHERE id IN (SELECT unnest(contractd_vertices) FROM contraction_results);
UPDATE 10
```

Fill contracted\_vertices with the information from the results that belong to the vertices.

```
UPDATE vertices
SET contracted_vertices = contraction_results.contractd_vertices
FROM contraction_results
WHERE type = 'v' AND vertices.id = contraction_results.id;
UPDATE 3
```

Insert the new edges generated by pgr\_contraction.

```
INSERT INTO edges(source, target, cost, reverse_cost, contracted_vertices, is_new)
SELECT source, target, cost, -1, contracted_vertices, true
FROM contraction_results
WHERE type = 'e';
INSERT 0 4
```

[The contracted graph¶](#)

Vertices that belong to the contracted graph.

```
SELECT id FROM vertices WHERE is_contracted = false ORDER BY id;
id
----
 4
 7
10
11
12
14
```

Edges that belong to the contracted graph.

```
WITH
vertices_in_graph AS (SELECT id FROM vertices WHERE is_contracted = false)
SELECT id, source, target, cost, reverse_cost, contracted_vertices
FROM edges
WHERE
  EXISTS (SELECT id FROM vertices AS v WHERE NOT is_contracted AND v.id = edges.source)
  AND
  EXISTS (SELECT id FROM vertices AS v WHERE NOT is_contracted AND v.id = edges.target)
ORDER BY id;
id | source | target | cost | reverse_cost | contracted_vertices
-----+-----+-----+-----+-----+-----
5 | 10 | 11 | 1 | -1 | 
8 | 7 | 11 | 1 | 1 | 
9 | 11 | 16 | 1 | 1 | 
11 | 11 | 12 | 1 | -1 | 
19 | 7 | 10 | 2 | -1 | {5,6}
20 | 7 | 12 | 2 | -1 | {8,9}
21 | 12 | 16 | 2 | -1 | {17}
22 | 10 | 16 | 2 | -1 | {15}
(8 rows)
```

Visually:



Using the contracted graph

Depending on the final application the graph is to be prepared. In this example the final application will be to calculate the cost from two vertices in the original graph by using the contracted graph with pgr\_dijkstraCost

There are three cases when calculating the shortest path between a given source and target in a contracted graph:

- Case 1: Both source and target belong to the contracted graph.
- Case 2: Source and/or target belong to an edge subgraph.
- Case 3: Source and/or target belong to a vertex.

The final application should consider all of those cases.

Create a view (or table) of the contracted graph:

```
DROP VIEW IF EXISTS contracted_graph;
NOTICE: view "contracted_graph" does not exist, skipping
DROP VIEW
CREATE VIEW contracted_graph AS
SELECT id,source, target, cost, reverse_cost, contracted_vertices FROM edges
WHERE
  EXISTS (SELECT id FROM vertices AS v WHERE NOT is_contracted AND v.id = edges.source)
  AND
  EXISTS (SELECT id FROM vertices AS v WHERE NOT is_contracted AND v.id = edges.target);
CREATE VIEW
```

Create the function that will use the contracted graph.

```
CREATE OR REPLACE FUNCTION path_cost(source BIGINT, target BIGINT)
RETURNS SETOF FLOAT AS
$BODY$

SELECT agg_cost FROM pgr_dijkstraCost(
/* The inner query */
'WITH
cul_de_sac AS (
  SELECT contracted_vertices || id as v
  FROM vertices WHERE ' || $1 || ' = ANY(contracted_vertices)
  OR ' || $2 || ' = ANY(contracted_vertices)),
linears_to_expand AS (
  SELECT id, contracted_vertices
  FROM edges WHERE is_new AND (' || $1 || ' = ANY(contracted_vertices)
  OR ' || $2 || ' = ANY(contracted_vertices))
),
additional_vertices AS (
  SELECT * FROM cul_de_sac UNION SELECT contracted_vertices FROM linears_to_expand)
SELECT id, source, target, cost, reverse_cost
FROM edges, additional_vertices WHERE source = ANY(v) OR target = ANY(v)

UNION

SELECT id, source, target, cost, reverse_cost
FROM contracted_graph LEFT JOIN linears_to_expand c USING (id) WHERE c.id IS NULL',

source, target, false);

$BODY$ LANGUAGE SQL;
CREATE FUNCTION
```

Case 1: Both source and target belong to the contracted graph.

```
SELECT * FROM path_cost(10, 12);
path_cost
-----
2
(1 row)
```

Case 2: Source and/or target belong to an edge that has contracted vertices.

```
SELECT * FROM path_cost(15, 12);
path_cost
-----
3
(1 row)
```

Case 3: Source and/or target belong to a vertex that has been contracted.

```
SELECT * FROM path_cost(15, 1);
path_cost
-----
```

See Also

- [Contraction - Family of functions](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_contractionDeadEnd - Proposed

pgr\_contractionDeadEnd — Performs graph contraction and returns the contracted vertices and edges.

Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Availability

- Version 3.8.0
  - New proposed function.

Description

Contraction reduces the size of the graph by removing some of the vertices and edges and, for example, might add edges that represent a sequence of original edges decreasing the total time and space used in graph algorithms.

The main Characteristics are:

- Process is done only on edges with positive costs.
- Does not return the full contracted graph.
  - Only changes on the graph are returned.
- The returned values include:
  - The new edges generated by linear contraction.
  - The modified vertices generated by dead end contraction.
- The returned values are ordered as follows:
  - column id ascending when its a modified vertex.
  - column id with negative numbers descending when its a new edge.

A node is considered a dead end node when:

- On undirected graphs:
  - The number of adjacent vertices is 1.
- On directed graphs:
  - When there is only one adjacent vertex or
  - When all edges are incoming regardless of the number of adjacent vertices.

Boost Graph Inside

Signatures

pgr\_contractionDeadEnd([Edges SQL](#), [options](#))

**options:** [directed, forbidden]

Returns set of (type, id, contracted\_vertices, source, target, cost)

Example:

Dead end contraction on an undirected graph.

```
SELECT * FROM pgr_contractionDeadEnd(
'SELECT id, source, target, cost, reverse_cost FROM edges',
directed => false);
```

	type	id	contracted_vertices	source	target	cost
v	4	{2}		-1	-1	-1
v	6	{5}		-1	-1	-1
v	7	{1,3}		-1	-1	-1
v	8	{9}		-1	-1	-1
v	14	{13}		-1	-1	-1

(5 rows)

- The green nodes are dead end nodes.
  - Node \3\ is a dead end node after node \1\ is contracted.

Parameters

Parameter Type	Description
<a href="#">Edges SQL</a> TEXT	<a href="#">Edges SQL</a> as described below.

Optional parameters¶

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true the graph is considered <i>Directed</i></li> <li>When false the graph is considered as <i>Undirected</i>.</li> </ul>

Contraction optional parameters¶

Column	Type	Default	Description
forbidden	ARRAY[ ANY-INTEGER ]	Empty	Identifiers of vertices forbidden for contraction.

Inner Queries¶

Edges SQL ¶

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns¶

Returns set of (type, id, contracted\_vertices, source, target, cost)

The function returns a single row. The columns of the row are:

Column	Type	Description
type	TEXT	Value = e indicating the row is an edge.
id	BIGINT	A pseudo <i>id</i> of the edge. <ul style="list-style-type: none"> <li>All numbers on this column are DISTINCT</li> <li>Decreasing sequence starting from -1.</li> </ul>
contracted_vertices	ARRAY[BIGINT]	Array of contracted vertex identifiers.
source	BIGINT	Identifier of the source vertex of the current edge.
target	BIGINT	Identifier of the target vertex of the current edge.
cost	FLOAT	Weight of the current edge.

Additional Examples¶

- [Dead end vertex on undirected graph](#)
- [Dead end vertex on directed graph](#)
- [Step by step dead end contraction](#)
- [Creating the contracted graph](#)
  - [Steps for the creation of the contracted graph](#)
  - [The contracted graph](#)

- [Using when departure and destination are in the contracted graph](#)
- [Using when departure/destination is not in the contracted graph](#)
- [Using when departure and destination are not in the contracted graph](#)

Dead end vertex on undirected graph¶

The green nodes are dead end nodes.

- They have only one adjacent node.

```
SELECT * FROM pgr_contractionDeadEnd(
$$SELECT * FROM (VALUES
(1, 1, 2, 1, -1),
(2, 3, 4, 1, -1),
(3, 2, 5, 1, 1), (4, 2, 6, 1, 1),
(5, 3, 5, 1, 1), (5, 3, 6, 1, 1))
AS edges(id,source,target,cost,reverse_cost))$,
directed => true);
type | id | contracted_vertices | source | target | cost
-----+-----+-----+-----+-----+-----
v    | 2 | {1}                | -1    | -1    | -1
v    | 3 | {4}                | -1    | -1    | -1
(2 rows)
```

Dead end vertex on directed graph¶

- The green nodes are dead end nodes
- The blue nodes have an unlimited number of incoming and/or outgoing edges.

Node	Adjacent nodes	Dead end	Reason
\(6\)	\(\{1\}\)	Yes	Has only one adjacent node.
\(7\)	\(\{2\}\)	Yes	Has only one adjacent node.
\(8\)	\(\{2, 3\}\)	Yes	Has more than one adjacent node and all edges are incoming.
\(9\)	\(\{4\}\)	Yes	Has only one adjacent node.
\(10\)	\(\{4, 5\}\)	No	Has more than one adjacent node and all edges are outgoing.
\(1,2,3,4,5\)	Many adjacent nodes.	No	Has more than one adjacent node and some edges are incoming and some are outgoing.

From above, nodes \(\{6, 7, 9\}\) are dead ends because the total number of adjacent vertices is one.

When there are more than one adjacent vertex, all edges need to be all incoming edges otherwise it is not a dead end.

```
SELECT * FROM pgr_contractionDeadEnd(
$$SELECT * FROM (VALUES
(1, 1, 6, 1, 1),
(2, 2, 7, 1, -1),
(3, 2, 8, 1, -1),
(4, 3, 8, 1, -1),
(5, 9, 4, 1, -1),
(6, 10, 4, 1, 1),
(7, 10, 5, 1, 1),
/* Rest of the graph */
(8, 1, 25, 1, 1), (9, 1, 26, 1, 1),
(10, 2, 25, 1, 1), (11, 2, 26, 1, 1),
(12, 3, 25, 1, 1), (13, 3, 26, 1, 1),
(14, 4, 25, 1, 1), (15, 4, 26, 1, 1),
(16, 5, 25, 1, 1), (17, 5, 26, 1, 1))
AS edges(id,source,target,cost,reverse_cost))$,
directed => true);
type | id | contracted_vertices | source | target | cost
-----+-----+-----+-----+-----+-----
v    | 1 | {6}                | -1    | -1    | -1
v    | 2 | {7,8}              | -1    | -1    | -1
v    | 3 | {8}                | -1    | -1    | -1
v    | 4 | {9}                | -1    | -1    | -1
(4 rows)
```

Step by step dead end contraction¶

The dead end contraction will stop until there are no more dead end nodes. For example, from the following graph where\(\{3\}\) is the dead end node:

After contracting \(\{3\}\), node \(\{2\}\) is now a dead end node and is contracted:

After contracting \(\{2\}\), stop. Node \(\{1\}\) has the information of nodes that were contracted.

```
SELECT * FROM pgr_contractionDeadEnd(
$$SELECT * FROM (VALUES
(1, 1, 2, 1, -1),
(2, 2, 3, 1, -1),
/* Rest of the graph */
(3, 1, 25, 1, 1), (4, 1, 26, 1, 1),
(5, 25, 25, 1, 1), (6, 25, 26, 1, 1))
AS edges(id,source,target,cost,reverse_cost))$,
directed => true);
type | id | contracted_vertices | source | target | cost
-----+-----+-----+-----+-----+-----
v    | 1 | {2,3}              | -1    | -1    | -1
(1 row)
```

Creating the contracted graph¶

- [Steps for the creation of the contracted graph](#)
- [The contracted graph](#)

Steps for the creation of the contracted graph¶

Add additional columns.

```
ALTER TABLE vertices ADD is_contracted BOOLEAN DEFAULT false;
ALTER TABLE
ALTER TABLE vertices ADD contracted_vertices BIGINT[];
ALTER TABLE
```

Save results into a table.

```
SELECT * INTO contraction_results
FROM pgr_contractionDeadEnd(
'SELECT id, source, target, cost, reverse_cost FROM edges',
directed => false);
SELECT 5
```

Use is\_contracted column to indicate the vertices that are contracted.

```
UPDATE vertices
SET is_contracted = true
WHERE id IN (SELECT unnest(contracted_vertices) FROM contraction_results);
UPDATE 6
```

Fill contracted\_vertices with the information from the results that belong to the vertices.

```
UPDATE vertices
SET contracted_vertices = contraction_results.contracted_vertices
FROM contraction_results
WHERE type = 'v' AND vertices.id = contraction_results.id;
UPDATE 5
```

The contracted vertices are not part of the contracted graph.

```
SELECT id, is_contracted
FROM vertices WHERE is_contracted ORDER BY id;
id | is_contracted
-----+-----
1 | t
2 | t
3 | t
5 | t
9 | t
13 | t
(6 rows)
```

The contracted graph¶

```
DROP VIEW IF EXISTS contracted_graph;
NOTICE: view "contracted_graph" does not exist, skipping
DROP VIEW
CREATE VIEW contracted_graph AS
WITH
vertices_in_graph AS (
SELECT id FROM vertices WHERE is_contracted = false
)
SELECT id, source, target, cost, reverse_cost
FROM edges
WHERE source IN (SELECT * FROM vertices_in_graph)
AND target IN (SELECT * FROM vertices_in_graph)
ORDER BY id;
CREATE VIEW
```

Using when departure and destination are in the contracted graph¶

```
SELECT *
FROM pgr_dijkstra("SELECT * FROM contracted_graph", 6, 17);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 17 | 6 | 4 | 1 | 0
2 | 2 | 6 | 17 | 7 | 8 | 1 | 1
3 | 3 | 6 | 17 | 11 | 11 | 1 | 2
4 | 4 | 6 | 17 | 12 | 13 | 1 | 3
5 | 5 | 6 | 17 | 17 | -1 | 0 | 4
(5 rows)
```

Using when departure/destination is not in the contracted graph¶

```
SELECT * FROM pgr_dijkstra(
'WITH cul_de_sac AS (
SELECT contracted_vertices || id as v
FROM vertices WHERE 1 = ANY(contracted_vertices)
SELECT id, source, target, cost, reverse_cost FROM edges, cul_de_sac
WHERE source = ANY(v) AND target = ANY(v)

UNION

SELECT id, source, target, cost, reverse_cost FROM contracted_graph',
1, 17);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 17 | 1 | 6 | 1 | 0
2 | 2 | 1 | 17 | 3 | 7 | 1 | 1
3 | 3 | 1 | 17 | 7 | 8 | 1 | 2
4 | 4 | 1 | 17 | 11 | 9 | 1 | 3
5 | 5 | 1 | 17 | 16 | 15 | 1 | 4
6 | 6 | 1 | 17 | 17 | -1 | 0 | 5
(6 rows)
```

Using when departure and destination are not in the contracted graph¶

```
SELECT * FROM pgr_dijkstra(
'WITH cul_de_sac AS (
SELECT contracted_vertices || id as v
FROM vertices WHERE 1 = ANY(contracted_vertices) OR 9 = ANY(contracted_vertices)
SELECT id, source, target, cost, reverse_cost FROM edges, cul_de_sac WHERE source = ANY(v) AND target = ANY(v)

UNION

SELECT id, source, target, cost, reverse_cost FROM contracted_graph',
1, 9);
```



seq	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	1	9	1	6	1	0
2	2	1	9	3	7	1	1
3	3	1	9	7	10	1	2
4	4	1	9	8	14	1	3
5	5	1	9	9	-1	0	4

(5 rows)

See Also

- [Contraction - Family of functions](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_contractionLinear - Proposed

pgr\_contractionLinear — Performs graph contraction and returns the contracted vertices and edges.

Availability

- Version 3.8.0
  - New proposed function.

Description

Contraction reduces the size of the graph by removing some of the vertices and edges and, for example, might add edges that represent a sequence of original edges decreasing the total time and space used in graph algorithms.

The main Characteristics are:

- Process is done only on edges with positive costs.
- Does not return the full contracted graph.
  - Only changes on the graph are returned.
- The returned values include:
  - The new edges generated by linear contraction.
  - The modified vertices generated by dead end contraction.
- The returned values are ordered as follows:
  - column id ascending when its a modified vertex.
  - column id with negative numbers descending when its a new edge.

Boost Graph Inside

Signatures

pgr\_contractionLinear([Edges SQL](#), [options](#))

**options:** [directed, forbidden]

Returns set of (type, id, contracted\_vertices, source, target, cost)

Example:

Linear contraction on an undirected graph.

```
SELECT * FROM pgr_contractionLinear(
'SELECT id, source, target, cost, reverse_cost FROM edges',
directed => false);
type | id | contracted_vertices | source | target | cost
-----+-----+-----+-----+-----+-----
e    |-1 | {3}                | 1      | 7      | 2
e    |-2 | {17}               | 12     | 16     | 2
e    |-3 | {15}               | 10     | 16     | 2
(3 rows)
```

- The green nodes are linear nodes and will not be part of the contracted graph.
  - All edges adjacent will not be part of the contracted graph.
- The red lines will be new edges of the contracted graph.

Parameters

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below.
<b>contraction Order</b>	ARRAY[ <b>ANY-INTEGER</b> ]	Ordered contraction operations. <ul style="list-style-type: none"><li>1 = Dead end contraction</li><li>2 = Linear contraction</li></ul>

Optional parameters

Column	Type	Default	Description
--------	------	---------	-------------

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true the graph is considered <i>Directed</i></li> <li>When false the graph is considered as <i>Undirected</i>.</li> </ul>

Contraction optional parameters

Column	Type	Default	Description
forbidden	ARRAY[ <b>ANY-INTEGER</b> ]	<b>Empty</b>	Identifiers of vertices forbidden for contraction.
cycles	INTEGER	\(1\)	Number of times the contraction operations on <code>contraction_order</code> will be performed.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	<b>ANY-INTEGER</b>		Identifier of the edge.
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
target	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
cost	<b>ANY-NUMERICAL</b>		Weight of the edge (source, target)
reverse_cost	<b>ANY-NUMERICAL</b>	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (type, id, contracted\_vertices, source, target, cost)

The function returns a single row. The columns of the row are:

Column	Type	Description
type	TEXT	Value = e indicating the row is an edge.
		A pseudo <i>id</i> of the edge.
id	BIGINT	<ul style="list-style-type: none"> <li>All numbers on this column are <b>DISTINCT</b></li> <li>Decreasing sequence starting from -1.</li> </ul>
contracted_vertices	ARRAY[BIGINT]	Array of contracted vertex identifiers.
source	BIGINT	Identifier of the source vertex of the current edge.
target	BIGINT	Identifier of the target vertex of the current edge.
cost	FLOAT	Weight of the current edge.

Additional Examples

- Linear edges
- Linearity is not symmetrical
- Linearity is symmetrical
- Step by step linear contraction
- Creating the contracted graph
  - Steps for the creation of the contracted graph
  - The contracted graph
- Using when departure and destination are in the contracted graph
- Using when departure/destination is not in the contracted graph

Linear edges

Undirected graph

A node connects two (or more) *linear* edges when

- The number of adjacent vertices is 2.

In case of a directed graph, a node is considered *alinear* node when

- The number of adjacent vertices is 2.
- Linearity is symmetrical.

[Linearity is not symmetrical¶](#)

Directed graph

Graph where linearity is not symmetrical.

When the graph is processed as a directed graph, linearity is not symmetrical, therefore the graph can not be contracted.

```
SELECT * FROM pgr_contractionLinear(
$$SELECT * FROM (VALUES
  (1, 1, 2, 1, -1),
  (2, 2, 3, 3, 4))
AS edges(id,source,target,cost,reverse_cost)$$,
directed => true);
type | id | contracted_vertices | source | target | cost
-----+-----+-----+-----+-----+-----
(0 rows)
```

Undirected graph

When the same graph is processed as an undirected graph, linearity is symmetrical, therefore the graph can be contracted.

```
SELECT * FROM pgr_contractionLinear(
$$SELECT * FROM (VALUES
  (1, 1, 2, 1, -1),
  (2, 2, 3, 3, 4))
AS edges(id,source,target,cost,reverse_cost)$$,
directed => false);
type | id | contracted_vertices | source | target | cost
-----+-----+-----+-----+-----+-----
e | -1 | {2} | 1 | 3 | 4
(1 row)
```

The three edges can be replaced by one undirected edge

- Edge  $\setminus(1 - 3\setminus)$ .
  - With cost:  $\setminus(4\setminus)$ .
  - Contracted vertices in the edge:  $\setminus(\setminus 2\setminus)\setminus$ .

[Linearity is symmetrical¶](#)

Directed graph

Graph where linearity is symmetrical.

When the graph is processed as a directed graph, linearity is not symmetrical, therefore the graph can not be contracted.

```
SELECT * FROM pgr_contractionLinear(
$$SELECT * FROM (VALUES
  (1, 1, 2, 1, 2),
  (2, 2, 3, 3, 4))
AS edges(id,source,target,cost,reverse_cost)$$,
directed => true);
type | id | contracted_vertices | source | target | cost
-----+-----+-----+-----+-----+-----
e | -1 | {2} | 1 | 3 | 4
e | -2 | {2} | 3 | 1 | 6
(2 rows)
```

The four edges can be replaced by two directed edges.

- Edge  $\setminus(1 - 3\setminus)$ .
  - With cost:  $\setminus(4\setminus)$ .
  - Contracted vertices in the edge:  $\setminus(\setminus 2\setminus)\setminus$ .
- Edge  $\setminus(3 - 1\setminus)$ .
  - With cost:  $\setminus(6\setminus)$ .
  - Contracted vertices in the edge:  $\setminus(\setminus 2\setminus)\setminus$ .

Undirected graph

When the same graph is processed as an undirected graph, linearity is symmetrical, therefore the graph can be contracted.

```
SELECT * FROM pgr_contractionLinear(
$$SELECT * FROM (VALUES
  (1, 1, 2, 1, 2),
  (2, 2, 3, 3, 4))
AS edges(id,source,target,cost,reverse_cost)$$,
directed => false);
type | id | contracted_vertices | source | target | cost
```

e	-1	[2]		1	3	4
(1 row)						

The four edges can be replaced by one undirected edge.

- Edge \{(1 - 3)\}.
  - With cost: \{4\}.
  - Contracted vertices in the edge:\{\{2\}\}.

[Step by step linear contraction¶](#)

The linear contraction will stop when there are no more linear edges. For example from the following graph there are linear edges

Contracting vertex \{3\},

- The vertex \{3\} is removed from the graph
- The edges \{2 \rightarrow 3\} and \{w \rightarrow z\} are removed from the graph.
- A new edge \{2 \rightarrow 4\} is inserted represented with red color.

Contracting vertex \{2\}:

- The vertex \{2\} is removed from the graph
- The edges \{1 \rightarrow 2\} and \{2 \rightarrow 3\} are removed from the graph.
- A new edge \{1 \rightarrow 3\} is inserted represented with red color.

Edge \{1 \rightarrow 3\} has the information of cost and the nodes that were contracted.

```

SELECT * FROM pgr_contractionLinear(
$$SELECT * FROM (VALUES
  (1, 1, 2, 1),
  (2, 2, 3, 1),
  (2, 3, 4, 1))
AS edges(id,source,target,cost)$$);
type | id | contracted_vertices | source | target | cost
-----+-----+-----+-----+-----+-----
e    | -1 | [2,3]              | 1      | 4      | 3
(1 row)

```

[Creating the contracted graph¶](#)

- [Steps for the creation of the contracted graph](#)
- [The contracted graph](#)

[Steps for the creation of the contracted graph¶](#)

Add additional columns.

```

ALTER TABLE vertices ADD is_contracted BOOLEAN DEFAULT false;
ALTER TABLE
ALTER TABLE edges ADD is_new BOOLEAN DEFAULT false;
ALTER TABLE
ALTER TABLE edges ADD contracted_vertices BIGINT[];
ALTER TABLE

```

Save results into a table.

```

SELECT * INTO contraction_results
FROM pgr_contractionLinear(
'SELECT id, source, target, cost, reverse_cost FROM edges',
directed => false);
SELECT 3

```

Use is\_contracted column to indicate the vertices that are contracted.

```

UPDATE vertices
SET is_contracted = true
WHERE id IN (SELECT unnest(contracted_vertices) FROM contraction_results);
UPDATE 3

```

The contracted vertices are not part of the contracted graph.

```

SELECT id, is_contracted
FROM vertices WHERE is_contracted ORDER BY id;
id | is_contracted
-----+-----
3  | t
15 | t
17 | t
(3 rows)

```

Insert the new edges generated by pgr\_contraction.

```

INSERT INTO edges(source, target, cost, reverse_cost, contracted_vertices, is_new)
SELECT source, target, cost, -1, contracted_vertices, true
FROM contraction_results;
INSERT 0 3

```

Create the contracted graph.

```

CREATE VIEW contracted_graph AS
WITH
vertices_in_graph AS (
  SELECT id FROM vertices WHERE NOT is_contracted
)
SELECT id, source, target, cost, reverse_cost
FROM edges
WHERE source IN (SELECT * FROM vertices_in_graph)
AND target IN (SELECT * FROM vertices_in_graph)
ORDER BY id;
CREATE VIEW

```

[The contracted graph¶](#)

```
SELECT * FROM contracted_graph ORDER by id;
id | source | target | cost | reverse_cost
-----
1 | 5 | 6 | 1 | 1
2 | 6 | 10 | -1 | 1
4 | 6 | 7 | 1 | 1
5 | 10 | 11 | 1 | -1
8 | 7 | 11 | 1 | 1
9 | 11 | 16 | 1 | 1
10 | 7 | 8 | 1 | 1
11 | 11 | 12 | 1 | -1
12 | 8 | 12 | 1 | -1
14 | 8 | 9 | 1 | 1
17 | 2 | 4 | 1 | 1
18 | 13 | 14 | 1 | 1
19 | 1 | 7 | 2 | -1
20 | 12 | 16 | 2 | -1
21 | 10 | 16 | 2 | -1
(15 rows)
```

Using when departure and destination are in the contracted graph

```
SELECT *
FROM pgr_dijkstra('SELECT * FROM contracted_graph', 7, 16);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----
1 | 1 | 7 | 16 | 7 | 8 | 1 | 0
2 | 2 | 7 | 16 | 11 | 9 | 1 | 1
3 | 3 | 7 | 16 | 16 | -1 | 0 | 2
(3 rows)
```

Using when departure/destination is not in the contracted graph

```
SELECT * FROM pgr_dijkstra(
  'WITH in_line AS (SELECT contracted_vertices FROM edges WHERE 17 = ANY(contracted_vertices))
  SELECT id, source, target, cost, reverse_cost
  FROM edges, in_line
  WHERE source = ANY(in_line.contracted_vertices) OR target = ANY(in_line.contracted_vertices)

  UNION

  SELECT id, source, target, cost, reverse_cost FROM contracted_graph',
  1, 17);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----
1 | 1 | 1 | 17 | 1 | 19 | 2 | 0
2 | 2 | 1 | 17 | 7 | 8 | 1 | 2
3 | 3 | 1 | 17 | 11 | 9 | 1 | 3
4 | 4 | 1 | 17 | 16 | 15 | 1 | 4
5 | 5 | 1 | 17 | 17 | -1 | 0 | 5
(5 rows)
```

See Also

- Contraction - Family of functions

Indices and tables

- Index
- Search Page

pgr\_contractionHierarchies - Experimental

pgr\_contractionHierarchies — Performs graph contraction according to the contraction hierarchies method and returns the contracted vertices and shortcut edges created.

Availability

- Version 3.8.0
  - New **experimental** function

Description

The contraction hierarchies method builds, from an initial order of the vertices, a hierarchical order, giving priority to some vertices during the processing of label fixing of shortest paths algorithms. Furthermore, the contraction hierarchies algorithm adds shortcut edges in the graph, that helps the shortest paths algorithm to follow the created hierarchical graph structure.

The idea of the hierarchy is to put at a high priority level vertices that belong to the long distance network (highways for example in a road network) and to a low level of priority nodes that belong to the short distance network (arterials or secondary roads for example in road networks).

The contraction hierarchies algorithm makes the assumption that there is already a valuable vertices order that is used to initialize the contraction process. As in most cases there is no valuable initial node ordering, we use the order given by vertices ID. Then, the contraction process is made on the basis of this first order to give the final hierarchy.

The basic idea is to keep the vertices in a priority queue sorted by some estimate of how attractive is their contraction. The implemented case uses the metric called *edge difference*, which corresponds to the difference between the number of shortcuts produced by a vertex contraction and the number of incident edges in the graph before contraction (*#shortcuts - #incident edges*).

Finally, the aim is to reduce the explored part of the graph, when using a bidirectional Dijkstra-like algorithm. The vertices order is used to feed the oriented search. The search is made without losing optimality.

Finding an optimal vertices ordering for contraction is a difficult problem. Nevertheless, very simple local heuristics work quite well, according to Geisberger et al. [2]. The principle here is to a priori estimate the value of the *edge difference* and to contract the node at the top of the queue only if the new value of the metric keeps it at the top of the queue. Otherwise, it is reinserted in the queue, at its right place corresponding to the new metric value.

The process is done on graphs having only edges with positive costs.

It is necessary to remember that there are no deleted vertices with this function. At the end, the graph keeps every vertex it had, but has some added edges, corresponding to shortcuts. The vertices which have been contracted, to build the shortcut edges, are kept and hierarchically ordered.

As for the other contraction methods, it does not return the full contracted graph, only the changes. They are here of two types:

- added shortcut edges, with negative identifiers;
- contracted nodes with an order.

Boost Graph Inside

Signatures

Summary

The `pgr_contractionHierarchies` function has the following signature:

```
pgr_contractionHierarchies(Edges SQL, [options])
options: [directed, forbidden]
Returns set of (type, id, contracted_vertices, source, target, cost, metric, vertex_order)
```

Parameters

Parameter	Type	Default	Description
<a href="#">Edges SQL</a>	TEXT		<a href="#">Edges SQL</a> as described below.
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Contraction hierarchies optional parameters

Column	Type	Default	Description
forbidden	ARRAY[ <b>ANY-INTEGER</b> ]	Empty	Identifiers of vertices forbidden for contraction.
directed	BOOLEAN	\(1\)	True if the graph is directed, False otherwise.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	<b>ANY-INTEGER</b>		Identifier of the edge.
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
target	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
cost	<b>ANY-NUMERICAL</b>		Weight of the edge (source, target)
reverse_cost	<b>ANY-NUMERICAL</b>	-1	<div>Weight of the edge (target, source)</div> <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (type, id, contracted\_vertices, source, target, cost, metric, vertex\_order)

The function returns many rows (one per vertex and one per shortcut edge created). The columns of the rows are:

Column	Type	Description
		Type of the id.
type	TEXT	<ul style="list-style-type: none"><li>v when the row is a vertex.<ul style="list-style-type: none"><li>Column id has a positive value</li></ul></li><li>e when the row is an edge.<ul style="list-style-type: none"><li>Column id has a negative value</li></ul></li></ul>

Column	Type	Description
All numbers on this column are DISTINCT		
id	BIGINT	<ul style="list-style-type: none"><li>When type = 'v'.<ul style="list-style-type: none"><li>Identifier of the modified vertex.</li></ul></li></ul>
		<ul style="list-style-type: none"><li>When type = 'e'.<ul style="list-style-type: none"><li>Decreasing sequence starting from -1.</li><li>Representing a pseudo <i>id</i> as is not incorporated in the set of original edges.</li></ul></li></ul>
contracted_vertices ARRAY[BIGINT] Array of contracted vertex identifiers.		
source	BIGINT	<ul style="list-style-type: none"><li>When type = 'v': \(-1\)</li><li>When type = 'e': Identifier of the source vertex of the current edge <math>\langle</math>source, target<math>\rangle</math>.</li></ul>
target	BIGINT	<ul style="list-style-type: none"><li>When type = 'v': \(-1\)</li><li>When type = 'e': Identifier of the target vertex of the current edge <math>\langle</math>source, target<math>\rangle</math>.</li></ul>
cost	FLOAT	<ul style="list-style-type: none"><li>When type = 'v': \(-1\)</li><li>When type = 'e': Weight of the current edge <math>\langle</math>source, target<math>\rangle</math>.</li></ul>
metric	BIGINT	<ul style="list-style-type: none"><li>When type = 'v': \(-1\)</li><li>When type = 'e': Weight of the current edge <math>\langle</math>source, target<math>\rangle</math>.</li></ul>
vertex_order	BIGINT	<ul style="list-style-type: none"><li>When type = 'v': \(-1\)</li><li>When type = 'e': Weight of the current edge <math>\langle</math>source, target<math>\rangle</math>.</li></ul>

Examples

On an undirected graph

The following query shows the original data involved in the contraction operation on an undirected graph.

```
SELECT id, source, target, cost FROM edges ORDER BY id;
```

id	source	target	cost
1	5	6	1
2	6	10	-1
3	10	15	-1
4	6	7	1
5	10	11	1
6	1	3	1
7	3	7	1
8	7	11	1
9	11	16	1
10	7	8	1
11	11	12	1
12	8	12	1
13	12	17	1
14	8	9	1
15	16	17	1
16	15	16	1
17	2	4	1
18	13	14	1

(18 rows)

The original graph:

[images/sample\\_graph.png](#)

Example:

building contraction hierarchies on the whole graph

```
SELECT * FROM pgr_contractionHierarchies(
'SELECT id, source, target, cost FROM edges',
directed => false);
```


type	id	contracted_vertices	source	target	cost	metric	vertex_order
v	1	{}	-1	-1	-1	-1	3
v	2	{}	-1	-1	-1	-1	10
v	3	{}	-1	-1	-1	-1	4
v	4	{}	-1	-1	-1	0	15
v	5	{}	-1	-1	-1	-1	14
v	6	{}	-1	-1	-1	-1	6
v	7	{}	-1	-1	-1	-1	5
v	8	{}	-1	-1	-1	-1	9
v	9	{}	-1	-1	-1	-2	2
v	10	{}	-1	-1	-1	-1	7
v	11	{}	-1	-1	-1	-1	8
v	12	{}	-1	-1	-1	-2	1
v	13	{}	-1	-1	-1	-1	11
v	14	{}	-1	-1	-1	0	16
v	15	{}	-1	-1	-1	-1	13
v	16	{}	-1	-1	-1	-1	12
v	17	{}	-1	-1	-1	0	17
e	-1	{7}	11	8	2	-1	-1
e	-2	{7,8}	11	9	3	-1	-1
e	-3	{8}	12	9	2	-1	-1
e	-4	{11}	12	16	2	-1	-1

(21 rows)

The results do not represent the contracted graph. They represent the changes done to the graph after applying the contraction algorithm and give the vertex order built by the algorithm, by ordering vertices according to the *edge difference* metric. As a consequence, vertices are all represented in the result (except of course forbidden ones). Only shortcut built by the algorithm are represented in the result.

After computing the contraction hierarchies, an order is now given to the vertices,  
in order to be used with a specific Dijkstra algorithm (implementation coming in a future version), which speeds up the search.

We obtain the contracted graph above:



We can see without surprise that the vertices belonging to the shortcuts have a tendency to have a high priority level in the resulting vertices order.

On an undirected graph with forbidden vertices

Example:  
building contraction with a set of forbidden vertices

```
SELECT * FROM pgr_contractionHierarchies(
'SELECT id, source, target, cost FROM edges',
directed => false,
forbidden => ARRAY[6]);
```

type	id	contracted_vertices	source	target	cost	metric	vertex_order
v	1	{}	-1	-1	-1	-1	4
v	2	{}	-1	-1	-1	-1	8
v	3	{}	-1	-1	-1	-1	5
v	4	{}	-1	-1	-1	0	15
v	5	{}	-1	-1	-1	-1	12
v	7	{}	-1	-1	-1	0	13
v	8	{}	-1	-1	-1	-1	7
v	9	{}	-1	-1	-1	-3	1
v	10	{}	-1	-1	-1	-1	6
v	11	{}	-1	-1	-1	0	14
v	12	{}	-1	-1	-1	-2	2
v	13	{}	-1	-1	-1	-1	9
v	14	{}	-1	-1	-1	0	16
v	15	{}	-1	-1	-1	-1	11
v	16	{}	-1	-1	-1	-2	3
v	17	{}	-1	-1	-1	-1	10
e	-1	{7}	6	11	2	-1	-1
e	-2	{7}	6	8	2	-1	-1
e	-3	{7}	11	8	2	-1	-1
e	-4	{7,8}	6	9	3	-1	-1
e	-5	{7,8}	11	9	3	-1	-1
e	-6	{8}	12	9	2	-1	-1
e	-7	{7,11}	6	12	3	-1	-1
e	-8	{7,11}	6	16	3	-1	-1
e	-9	{11}	12	16	2	-1	-1
e	-10	{7,11,12}	6	17	4	-1	-1

(26 rows)

Contraction process steps details

Shortcut building process

A vertex  $v$  is contracted by adding shortcuts replacing former paths of the form  $(u, v, w)$  by an edge  $(u, w)$ . The shortcut  $(u, w)$  is only needed when  $(u, v, w)$  is the only shortest path between  $u$  and  $w$ .

When all shortcuts have been added for a given vertex  $v$ , the incident edges of  $v$  are removed and another vertex is contracted with the remaining graph.

The procedure is destructive for the graph and a copy is made to be able to manipulate it again as a whole. The contraction process adds all discovered shortcuts to the edge set and attributes a metric to each contracted vertex. This metric is giving what is called the *contraction hierarchy*.

Initialize the queue with a first vertices order

For each vertex  $v$  of the graph, a contraction of  $v$  is built:

Node Adjacent nodes

$\backslash(v)$   $\backslash(\{p, r, u\})$

$\backslash(p)$   $\backslash(\{u, v\})$

$\backslash(u)$   $\backslash(\{p, v, w\})$

$\backslash(r)$   $\backslash(\{v, w\})$

$\backslash(w)$   $\backslash(\{r, u\})$

Adjacent edges are removed.

Shortcuts are built from predecessors of  $v$  to successors of  $v$  if and only if the path through  $v$  corresponds to the only shortest path between the predecessor and the successor of  $v$  in the graph. The *edge difference* metric here takes the value of -2.

Then the following vertex is contracted. The process goes on until each node of the graph has been contracted. At the end, there are no more edges in the graph, which has been destroyed by the process.

This first contraction will give a vertices order, given by ordering them in ascending order on the metric (edge difference). A total vertices order is built.  $u < v$ , then  $u$  is less important than  $v$ . The



algorithm keeps the vertices into a queue in this order.

A hierarchy will now be constructed by contracting again the vertices in this order.

Build the final vertex order

Once the first order built, the algorithm uses it to browse the graph once again. For each vertex taken in the queue, the algorithm simulates contraction and calculates its edge difference. If the computed value is greater than the one of the next vertex to be contracted, then the algorithm puts it back in the queue (heuristic approach). Otherwise it contracts it permanently.

Add shortcuts to the initial graph

At the end, the algorithm takes the initial graph (before edges deletions) and adds the shortcut edges to it. It gives you the contracted graph, ready to use with a specialized Dijkstra algorithm, which takes into account the order of the nodes in the hierarchy.

Use the contraction

Build the contraction

```
SELECT * INTO contraction_results
FROM pgr_contractionHierarchies(
'SELECT id, source, target, cost FROM edges',
directed => false);
SELECT 21
```

Add shortcuts and hierarchy in the existing tables

Add new columns in the *vertices* and *edges* tables to store the results:

```
ALTER TABLE edges
ADD is_new BOOLEAN DEFAULT false,
ADD contracted_vertices BIGINT[];
ALTER TABLE

ALTER TABLE vertices
ADD metric INTEGER,
ADD vertex_order INTEGER;
ALTER TABLE
```

Update and insert the results in the two tables.

```
INSERT INTO edges(source, target, cost, reverse_cost, contracted_vertices, is_new)
SELECT source, target, cost, -1, contracted_vertices, true
FROM contraction_results
WHERE type = 'e';
INSERT 0 4

UPDATE vertices
SET metric = c.metric, vertex_order = c.vertex_order
FROM contraction_results c
WHERE c.type = 'v' AND c.id = vertices.id;
UPDATE 17
```

Use a Dijkstra shortest path algorithm on

Then you can use any Dijkstra-like algorithm, waiting for the adapted one which will take into account the built vertices hierarchy. For example:

```
SELECT * FROM pgr_bdDijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
1, 17
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 6 | 1 | 0
2 | 2 | 3 | 7 | 1 | 1
3 | 3 | 7 | 8 | 1 | 2
4 | 4 | 11 | 11 | 1 | 3
5 | 5 | 12 | 13 | 1 | 4
6 | 6 | 17 | -1 | 0 | 5
(6 rows)
```

See Also

- Contraction - Family of functions

Indices and tables

- Index
- Search Page

Introduction

In large graphs, like road graphs or electric networks, graph contraction can be used to speed up some graph algorithms. Contraction can reduce the size of the graph by removing some of the vertices and edges and adding edges that represent a sequence of original edges (the original ones can be kept in some methods). In this way, it decreases the total time and space used by graph algorithms, particularly those searching for an optimal path.

This implementation gives a flexible framework for adding contraction algorithms in the future. Currently, it supports three algorithms.

- Dead end contraction
- Linear contraction
- Contraction hierarchies

The two first ones can be combined through a iterative procedure, via the `pgr_contraction` method. The third one is implemented on its own.

All functions allow the user to forbid contraction on a set of nodes.

See Also

- <https://www.cs.cmu.edu/afs/cs/academic/class/15210-f12/www/lectures/lecture16.pdf>
- [https://ae.iti.kit.edu/download/diploma\\_thesis\\_geisberger.pdf](https://ae.iti.kit.edu/download/diploma_thesis_geisberger.pdf)
- <https://jlazarsfeld.github.io/ch.150.project/contents/>

Indices and tables

- Index
- Search Page

Dijkstra - Family of functions

- `pgr_dijkstra` - Dijkstra's algorithm for the shortest paths.

- [pgr\\_dijkstraCost](#) - Get the aggregate cost of the shortest paths.
- [pgr\\_dijkstraCostMatrix](#) - Use pgr\_dijkstra to create a costs matrix.
- [pgr\\_drivingDistance](#) - Use pgr\_dijkstra to calculate catchment information.
- [pgr\\_KSP](#) - Use Yen algorithm with pgr\_dijkstra to get the K shortest paths.

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.
- [pgr\\_dijkstraVia - Proposed](#) - Get a route of a sequence of vertices.
- [pgr\\_dijkstraNear - Proposed](#) - Get the route to the nearest vertex.
- [pgr\\_dijkstraNearCost - Proposed](#) - Get the cost to the nearest vertex.

[pgr\\_dijkstra](#)

`pgr_dijkstra` — Shortest path using Dijkstra algorithm.

Availability

- Version 3.5.0
  - Standardizing output columns to (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)
    - `pgr_dijkstra`(One to One) added start\_vid and end\_vid columns.
    - `pgr_dijkstra`(One to Many) added end\_vid column.
    - `pgr_dijkstra`(Many to One) added start\_vid column.
- Version 3.1.0
  - New proposed signature:
    - `pgr_dijkstra`(Combinations)
- Version 3.0.0
  - Function promoted to official.
- Version 2.2.0
  - New proposed signatures:
    - `pgr_dijkstra`(One to Many)
    - `pgr_dijkstra`(Many to One)
    - `pgr_dijkstra`(Many to Many)
- Version 2.1.0
  - Signature change on `pgr_dijkstra`(One to One)
- Version 2.0.0
  - Official function.

[Description](#)

Dijkstra's algorithm, conceived by Dutch computer scientist Edsger Dijkstra in 1956. It is a graph search algorithm that solves the shortest path problem for a graph with non-negative edge path costs, producing a shortest path from a starting vertex to an ending vertex. This implementation can be used with a directed graph and an undirected graph.

- Process is done only on edges with positive costs.
  - A negative value on a cost column is interpreted as the edge does not exist.
- Values are returned when there is a path.
- When there is no path:
  - When the starting vertex and ending vertex are the same.
    - The **aggregate cost** of the non included values  $\backslash((v, v)\backslash)$  is  $\backslash(0\backslash)$
  - When the starting vertex and ending vertex are the different and there is no path:
    - The **aggregate cost** the non included values  $\backslash((u, v)\backslash)$  is  $\backslash(\infty\backslash)$
- For optimization purposes, any duplicated value in the starting vertices or on the ending vertices are ignored.
- Running time:  $\backslash(O(\backslash \text{start\ vids} \mid * (V \backslash \log V + E)\backslash))$

Boost Graph Inside

[Signatures](#)

Summary

```
pgr_dijkstra(Edges SQL, start_vid, end_vid, [directed])
pgr_dijkstra(Edges SQL, start_vid, end_vids, [directed])
pgr_dijkstra(Edges SQL, start_vids, end_vid, [directed])
```

pgr\_dijkstra([Edges SQL](#), start vids, end vids, [directed])  
pgr\_dijkstra([Edges SQL](#), [Combinations SQL](#), [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Warning

Breaking change on 3.5.0

Read the [Migration guide](#) about how to migrate from the old result columns to the new result columns.

One to One

pgr\_dijkstra([Edges SQL](#), start vid, end vid, [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertex \10\ on a **directed** graph

```
SELECT * FROM pgr_Dijkstra(
  'select id, source, target, cost, reverse_cost from edges',
  6, 10, true);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(6 rows)
```

One to Many

pgr\_dijkstra([Edges SQL](#), start vid, end vids, [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertices \10, 17\ on a **directed**

```
SELECT * FROM pgr_Dijkstra(
  'select id, source, target, cost, reverse_cost from edges',
  6, ARRAY[10, 17]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
7 | 1 | 6 | 17 | 6 | 4 | 1 | 0
8 | 2 | 6 | 17 | 7 | 8 | 1 | 1
9 | 3 | 6 | 17 | 11 | 9 | 1 | 2
10 | 4 | 6 | 17 | 16 | 15 | 1 | 3
11 | 5 | 6 | 17 | 17 | -1 | 0 | 4
(11 rows)
```

Many to One

pgr\_dijkstra([Edges SQL](#), start vids, end vid, [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \6, 1\ to vertex \17\ on a **directed** graph

```
SELECT * FROM pgr_Dijkstra(
  'select id, source, target, cost, reverse_cost from edges',
  ARRAY[6, 1], 17);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 17 | 1 | 6 | 1 | 0
2 | 2 | 1 | 17 | 3 | 7 | 1 | 1
3 | 3 | 1 | 17 | 7 | 8 | 1 | 2
4 | 4 | 1 | 17 | 11 | 11 | 1 | 3
5 | 5 | 1 | 17 | 12 | 13 | 1 | 4
6 | 6 | 1 | 17 | 17 | -1 | 0 | 5
7 | 1 | 6 | 17 | 6 | 4 | 1 | 0
8 | 2 | 6 | 17 | 7 | 8 | 1 | 1
9 | 3 | 6 | 17 | 11 | 11 | 1 | 2
10 | 4 | 6 | 17 | 12 | 13 | 1 | 3
11 | 5 | 6 | 17 | 17 | -1 | 0 | 4
(11 rows)
```

Many to Many

pgr\_dijkstra([Edges SQL](#), start vids, end vids, [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \6, 1\ to vertices \10, 17\ on an **undirected** graph

```
SELECT * FROM pgr_Dijkstra(
  'select id, source, target, cost, reverse_cost from edges',
  ARRAY[6, 1], ARRAY[10, 17],
  directed => false);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 10 | 1 | 6 | 1 | 0
2 | 2 | 1 | 10 | 3 | 7 | 1 | 1
3 | 3 | 1 | 10 | 7 | 4 | 1 | 2
4 | 4 | 1 | 10 | 6 | 2 | 1 | 3
5 | 5 | 1 | 10 | 10 | -1 | 0 | 4
6 | 1 | 1 | 17 | 1 | 6 | 1 | 0
7 | 2 | 1 | 17 | 3 | 7 | 1 | 1
8 | 3 | 1 | 17 | 7 | 8 | 1 | 2
9 | 4 | 1 | 17 | 11 | 9 | 1 | 3
10 | 5 | 1 | 17 | 16 | 15 | 1 | 4
11 | 6 | 1 | 17 | 17 | -1 | 0 | 5
12 | 1 | 6 | 10 | 6 | 2 | 1 | 0
13 | 2 | 6 | 10 | 10 | -1 | 0 | 1
```

14		1		6		17		6		4		1		0
15		2		6		17		7		8		1		1
16		3		6		17		11		11		1		2
17		4		6		17		12		13		1		3
18		5		6		17		17		-1		0		4
(18 rows)														

Combinations¶

pgr\_dijkstra([Edges SQL](#), [Combinations SQL](#), [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:  
Using a combinations table on an **undirected** graph

The combinations table:  
  
SELECT source, target FROM combinations;  
source | target  
-----+-----  
5 | 6  
5 | 10  
6 | 5  
6 | 15  
6 | 14  
(5 rows)

The query:  
  
SELECT \* FROM pgr\_Dijkstra(  
  'SELECT id, source, target, cost, reverse\_cost FROM edges',  
  'SELECT source, target FROM combinations',  
  false);  
seq | path\_seq | start\_vid | end\_vid | node | edge | cost | agg\_cost  
-----+-----+-----+-----+-----+-----+-----+-----  
1 | 1 | 5 | 6 | 5 | 1 | 1 | 0  
2 | 2 | 5 | 6 | 6 | -1 | 0 | 1  
3 | 1 | 5 | 10 | 5 | 1 | 1 | 0  
4 | 2 | 5 | 10 | 6 | 2 | 1 | 1  
5 | 3 | 5 | 10 | 10 | -1 | 0 | 2  
6 | 1 | 6 | 5 | 6 | 1 | 1 | 0  
7 | 2 | 6 | 5 | 5 | -1 | 0 | 1  
8 | 1 | 6 | 15 | 6 | 2 | 1 | 0  
9 | 2 | 6 | 15 | 10 | 3 | 1 | 1  
10 | 3 | 6 | 15 | 15 | -1 | 0 | 2  
(10 rows)

Parameters¶

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters¶

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries¶

Edges SQL¶

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGERS:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
source	ANY-INTEGERS	Identifier of the departure vertex.
target	ANY-INTEGERS	Identifier of the arrival vertex.

Where:

ANY-INTEGERS:

SMALLINT, INTEGER, BIGINT

Result columns

Returns set of (seq, path\_seq [, start\_vid] [, end\_vid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex. Returned when multiple starting vertices are in the query. <ul style="list-style-type: none"><li>Many to One</li><li>Many to Many</li></ul>
end_vid	BIGINT	Identifier of the ending vertex. Returned when multiple ending vertices are in the query. <ul style="list-style-type: none"><li>One to Many</li><li>Many to Many</li></ul>
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Additional Examples

Example:

Demonstration of repeated values are ignored, and result is sorted.

```
SELECT * FROM pgr_Dijkstra(
'select id, source, target, cost, reverse_cost from edges',
ARRAY[7, 10, 15, 10, 10, 15], ARRAY[10, 7, 10, 15]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1	1	7	10	7	8	1	0
2	2	7	10	11	9	1	1
3	3	7	10	16	16	1	2
4	4	7	10	15	3	1	3
5	5	7	10	10	-1	0	4
6	1	7	15	7	8	1	0
7	2	7	15	11	9	1	1
8	3	7	15	16	16	1	2
9	4	7	15	15	-1	0	3
10	1	10	7	10	5	1	0
11	2	10	7	11	8	1	1
12	3	10	7	7	-1	0	2
13	1	10	15	10	5	1	0
14	2	10	15	11	9	1	1
15	3	10	15	16	16	1	2
16	4	10	15	15	-1	0	3
17	1	15	7	15	16	1	0
18	2	15	7	16	9	1	1
19	3	15	7	11	8	1	2
20	4	15	7	7	-1	0	3
21	1	15	10	15	3	1	0
22	2	15	10	10	-1	0	1
(22 rows)							

Example 2:

Making start\_vids the same as end\_vids

```
SELECT * FROM pgr_Dijkstra(
'select id, source, target, cost, reverse_cost from edges',
ARRAY[7, 10, 15], ARRAY[7, 10, 15]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1	1	7	10	7	8	1	0
2	2	7	10	11	9	1	1
3	3	7	10	16	16	1	2
4	4	7	10	15	3	1	3

5	5	7	10	10	-1	0	4
6	1	7	15	7	8	1	0
7	2	7	15	11	9	1	1
8	3	7	15	16	16	1	2
9	4	7	15	15	-1	0	3
10	1	10	7	10	5	1	0
11	2	10	7	11	8	1	1
12	3	10	7	7	-1	0	2
13	1	10	15	10	5	1	0
14	2	10	15	11	9	1	1
15	3	10	15	16	16	1	2
16	4	10	15	15	-1	0	3
17	1	15	7	15	16	1	0
18	2	15	7	16	9	1	1
19	3	15	7	11	8	1	2
20	4	15	7	7	-1	0	3
21	1	15	10	15	3	1	0
22	2	15	10	10	-1	0	1

(22 rows)

Example:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_Dijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)');
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1	1	6	7	6	4	1	0
2	2	6	7	7	-1	0	1
3	1	6	10	6	4	1	0
4	2	6	10	7	8	1	1
5	3	6	10	11	9	1	2
6	4	6	10	16	16	1	3
7	5	6	10	15	3	1	4
8	6	6	10	10	-1	0	5
9	1	12	10	12	13	1	0
10	2	12	10	17	15	1	1
11	3	12	10	16	16	1	2
12	4	12	10	15	3	1	3
13	5	12	10	10	-1	0	4

(13 rows)

The examples of this section are based on the [Sample Data](#) network.

- For **directed** graphs with `cost` and `reverse_cost` columns
  - 1) Path from `\(6\)` to `\(10\)`
  - 2) Path from `\(6\)` to `\(7\)`
  - 3) Path from `\(12\)` to `\(10\)`
  - 4) Path from `\(12\)` to `\(7\)`
  - 5) Using One to Many to get the solution of examples 1 and 2
  - 6) Using Many to One to get the solution of examples 2 and 4
  - 7) Using Many to Many to get the solution of examples 1 to 4
  - 8) Using Combinations to get the solution of examples 1 to 3
- For **undirected** graphs with `cost` and `reverse_cost` columns
  - 9) Path from `\(6\)` to `\(10\)`
  - 10) Path from `\(6\)` to `\(7\)`
  - 11) Path from `\(12\)` to `\(10\)`
  - 12) Path from `\(12\)` to `\(7\)`
  - 13) Using One to Many to get the solution of examples 9 and 10
  - 14) Using Many to One to get the solution of examples 10 and 12
  - 15) Using Many to Many to get the solution of examples 9 to 12
  - 16) Using Combinations to get the solution of examples 9 to 11
- For **directed** graphs only with `cost` column
  - 17) Path from `\(6\)` to `\(10\)`
  - 18) Path from `\(6\)` to `\(7\)`
  - 19) Path from `\(12\)` to `\(10\)`
  - 20) Path from `\(12\)` to `\(7\)`
  - 21) Using One to Many to get the solution of examples 17 and 18
  - 22) Using Many to One to get the solution of examples 18 and 20
  - 23) Using Many to Many to get the solution of examples 17 to 20
  - 24) Using Combinations to get the solution of examples 17 to 19
- For **undirected** graphs only with `cost` column
  - 25) Path from `\(6\)` to `\(10\)`
  - 26) Path from `\(6\)` to `\(7\)`
  - 27) Path from `\(12\)` to `\(10\)`
  - 28) Path from `\(12\)` to `\(7\)`
  - 29) Using One to Many to get the solution of examples 25 and 26
  - 30) Using Many to One to get the solution of examples 26 and 28
  - 31) Using Many to Many to get the solution of examples 25 to 28
  - 32) Using Combinations to get the solution of examples 25 to 27
- Equivalences between signatures
  - 33) Using One to One

- [34\) Using One to Many](#)
- [35\) Using Many to One](#)
- [36\) Using Many to Many](#)
- [37\) Using Combinations](#)

For directed graphs with cost and reverse cost columns¶

images/Fig1-originalData.png



Directed graph with cost and reverse cost columns¶

1) Path from \6) to \10)¶

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
6, 10
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(6 rows)
```

2) Path from \6) to \7)¶

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
6, 7
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
(2 rows)
```

3) Path from \12) to \10)¶

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
12, 10
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 12 | 10 | 12 | 13 | 1 | 0
2 | 2 | 12 | 10 | 17 | 15 | 1 | 1
3 | 3 | 12 | 10 | 16 | 16 | 1 | 2
4 | 4 | 12 | 10 | 15 | 3 | 1 | 3
5 | 5 | 12 | 10 | 10 | -1 | 0 | 4
(5 rows)
```

4) Path from \12) to \7)¶

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
12, 7
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 12 | 7 | 12 | 13 | 1 | 0
2 | 2 | 12 | 7 | 17 | 15 | 1 | 1
3 | 3 | 12 | 7 | 16 | 9 | 1 | 2
4 | 4 | 12 | 7 | 11 | 8 | 1 | 3
5 | 5 | 12 | 7 | 7 | -1 | 0 | 4
(5 rows)
```

5) Using [One to Many](#) to get the solution of examples 1 and 2¶

Paths  $\{6\} \rightarrow \{10, 7\}$

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
6, ARRAY[10, 7]
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 7 | -1 | 0 | 1
3 | 1 | 6 | 10 | 6 | 4 | 1 | 0
4 | 2 | 6 | 10 | 7 | 8 | 1 | 1
5 | 3 | 6 | 10 | 11 | 9 | 1 | 2
6 | 4 | 6 | 10 | 16 | 16 | 1 | 3
7 | 5 | 6 | 10 | 15 | 3 | 1 | 4
8 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(8 rows)
```

6) Using [Many to One](#) to get the solution of examples 2 and 4¶

Paths  $\{(6, 12) \rightarrow (7)\}$

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  ARRAY[6, 12], 7
);
```

seq	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	6	7	6	4	1	0
2	2	6	7	7	-1	0	1
3	1	12	7	12	13	1	0
4	2	12	7	17	15	1	1
5	3	12	7	16	9	1	2
6	4	12	7	11	8	1	3
7	5	12	7	7	-1	0	4

(7 rows)

7) Using [Many to Many](#) to get the solution of examples 1 to 4

Paths  $\{(6, 12) \rightarrow (10, 7)\}$

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  ARRAY[6, 12], ARRAY[10, 7]
);
```

seq	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	6	7	6	4	1	0
2	2	6	7	7	-1	0	1
3	1	6	10	6	4	1	0
4	2	6	10	7	8	1	1
5	3	6	10	11	9	1	2
6	4	6	10	16	16	1	3
7	5	6	10	15	3	1	4
8	6	6	10	10	-1	0	5
9	1	12	7	12	13	1	0
10	2	12	7	17	15	1	1
11	3	12	7	16	9	1	2
12	4	12	7	11	8	1	3
13	5	12	7	7	-1	0	4
14	1	12	10	12	13	1	0
15	2	12	10	17	15	1	1
16	3	12	10	16	16	1	2
17	4	12	10	15	3	1	3
18	5	12	10	10	-1	0	4

(18 rows)

8) Using [Combinations](#) to get the solution of examples 1 to 3

Paths  $\{(6) \rightarrow (10, 7) \cup (12) \rightarrow (10)\}$

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)'
);
```

seq	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	6	7	6	4	1	0
2	2	6	7	7	-1	0	1
3	1	6	10	6	4	1	0
4	2	6	10	7	8	1	1
5	3	6	10	11	9	1	2
6	4	6	10	16	16	1	3
7	5	6	10	15	3	1	4
8	6	6	10	10	-1	0	5
9	1	12	10	12	13	1	0
10	2	12	10	17	15	1	1
11	3	12	10	16	16	1	2
12	4	12	10	15	3	1	3
13	5	12	10	10	-1	0	4

(13 rows)

[For undirected graphs with cost and reverse\\_cost columns](#)



Undirected graph with cost and reverse cost columns

9) Path from  $(6)$  to  $(10)$

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  6, 10,
  false
);
```

seq	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	6	10	6	2	1	0
2	2	6	10	10	-1	0	1

(2 rows)

10) Path from  $(6)$  to  $(7)$

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
```



```
6, 7,
false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
(2 rows)
```

11) Path from  $\{12\}$  to  $\{10\}$

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
12, 10,
false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 12 | 10 | 12 | 11 | 1 | 0
2 | 2 | 12 | 10 | 11 | 5 | 1 | 1
3 | 3 | 12 | 10 | 10 | -1 | 0 | 2
(3 rows)
```

12) Path from  $\{12\}$  to  $\{7\}$

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
12, 7,
false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 12 | 7 | 12 | 12 | 1 | 0
2 | 2 | 12 | 7 | 8 | 10 | 1 | 1
3 | 3 | 12 | 7 | 7 | -1 | 0 | 2
(3 rows)
```

13) Using [One to Many](#) to get the solution of examples 9 and 10

Paths  $\{6\} \rightarrow \{10, 7\}$

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
6, ARRAY[10,7],
false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
3 | 1 | 6 | 10 | 6 | 2 | 1 | 0
4 | 2 | 6 | 10 | 10 | -1 | 0 | 1
(4 rows)
```

14) Using [Many to One](#) to get the solution of examples 10 and 12

Paths  $\{6, 12\} \rightarrow \{7\}$

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[6,12], 7,
false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
3 | 1 | 12 | 7 | 12 | 12 | 1 | 0
4 | 2 | 12 | 7 | 8 | 10 | 1 | 1
5 | 3 | 12 | 7 | 7 | -1 | 0 | 2
(5 rows)
```

15) Using [Many to Many](#) to get the solution of examples 9 to 12

Paths  $\{6, 12\} \rightarrow \{10, 7\}$

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[6, 12], ARRAY[10,7],
false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
3 | 1 | 6 | 10 | 6 | 2 | 1 | 0
4 | 2 | 6 | 10 | 10 | -1 | 0 | 1
5 | 1 | 12 | 7 | 12 | 12 | 1 | 0
6 | 2 | 12 | 7 | 8 | 10 | 1 | 1
7 | 3 | 12 | 7 | 7 | -1 | 0 | 2
8 | 1 | 12 | 10 | 12 | 11 | 1 | 0
9 | 2 | 12 | 10 | 11 | 5 | 1 | 1
10 | 3 | 12 | 10 | 10 | -1 | 0 | 2
(10 rows)
```

16) Using [Combinations](#) to get the solution of examples 9 to 12

Paths  $\{6\} \rightarrow \{10, 7\} \cup \{12\} \rightarrow \{10\}$

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)',
false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
3 | 1 | 6 | 10 | 6 | 2 | 1 | 0
4 | 2 | 6 | 10 | 10 | -1 | 0 | 1
5 | 1 | 12 | 10 | 12 | 11 | 1 | 0
6 | 2 | 12 | 10 | 11 | 5 | 1 | 1
7 | 3 | 12 | 10 | 10 | -1 | 0 | 2
(7 rows)
```

[For directed graphs only with cost column](#)

images/Fig2-cost.png

Directed graph only with cost column

17) Path from 6 to 10

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost FROM edges',
6, 10
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
(0 rows)
```

18) Path from 6 to 7

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost FROM edges',
6, 7
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
(2 rows)
```

19) Path from 12 to 10

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost FROM edges',
12, 10
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
(0 rows)
```

20) Path from 12 to 7

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost FROM edges',
12, 7
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
(0 rows)
```

21) Using One to Many to get the solution of examples 17 and 18

Paths 6 → 10, 7

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost FROM edges',
6, ARRAY[10,7]
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
(2 rows)
```

22) Using Many to One to get the solution of examples 18 and 20

Paths 6, 12 → 7

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost FROM edges',
ARRAY[6,12], 7
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
(2 rows)
```

23) Using Many to Many to get the solution of examples 17 to 20

Paths 6, 12 → 10, 7

```
SELECT * FROM pgr_dijkstra(
'SELECT id, source, target, cost FROM edges',
ARRAY[6, 12], ARRAY[10,7]
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
(2 rows)
```

24) Using Combinations to get the solution of examples 17 to 19

Paths  $\{(6) \rightarrow (10, 7) \cup (12) \rightarrow (10)\}$

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost FROM edges',
  'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)'
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
(2 rows)
```

For undirected graphs only with cost column¶



Undirected graph only with cost column¶

25) Path from (6) to (10)¶

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost FROM edges',
  6, 10,
  false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 5 | 1 | 2
4 | 4 | 6 | 10 | 10 | -1 | 0 | 3
(4 rows)
```

26) Path from (6) to (7)¶

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost FROM edges',
  6, 7,
  false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
(2 rows)
```

27) Path from (12) to (10)¶

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost FROM edges',
  12, 10,
  false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 12 | 10 | 12 | 11 | 1 | 0
2 | 2 | 12 | 10 | 11 | 5 | 1 | 1
3 | 3 | 12 | 10 | 10 | -1 | 0 | 2
(3 rows)
```

28) Path from (12) to (7)¶

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost FROM edges',
  12, 7,
  false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 12 | 7 | 12 | 12 | 1 | 0
2 | 2 | 12 | 7 | 8 | 10 | 1 | 1
3 | 3 | 12 | 7 | 7 | -1 | 0 | 2
(3 rows)
```

29) Using One to Many to get the solution of examples 25 and 28¶

Paths  $\{(6) \rightarrow (10, 7)\}$

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost FROM edges',
  6, ARRAY[10, 7],
  false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | -1 | 0 | 1
3 | 1 | 6 | 10 | 6 | 4 | 1 | 0
4 | 2 | 6 | 10 | 7 | 8 | 1 | 1
5 | 3 | 6 | 10 | 11 | 5 | 1 | 2
6 | 4 | 6 | 10 | 10 | -1 | 0 | 3
(6 rows)
```

30) Using [Many to One](#) to get the solution of examples 26 and 29

Paths  $\{(6, 12)\} \rightarrow \{7\}$

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost FROM edges',
  ARRAY[6,12], 7,
  false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
3 | 1 | 12 | 7 | 12 | 12 | 1 | 0
4 | 2 | 12 | 7 | 8 | 10 | 1 | 1
5 | 3 | 12 | 7 | 7 | -1 | 0 | 2
(5 rows)
```

31) Using [Many to Many](#) to get the solution of examples 25 to 28

Paths  $\{(6, 12)\} \rightarrow \{10, 7\}$

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost FROM edges',
  ARRAY[6, 12], ARRAY[10,7],
  false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
3 | 1 | 6 | 10 | 6 | 4 | 1 | 0
4 | 2 | 6 | 10 | 7 | 8 | 1 | 1
5 | 3 | 6 | 10 | 11 | 5 | 1 | 2
6 | 4 | 6 | 10 | 10 | -1 | 0 | 3
7 | 1 | 12 | 7 | 12 | 12 | 1 | 0
8 | 2 | 12 | 7 | 8 | 10 | 1 | 1
9 | 3 | 12 | 7 | 7 | -1 | 0 | 2
10 | 1 | 12 | 10 | 12 | 11 | 1 | 0
11 | 2 | 12 | 10 | 11 | 5 | 1 | 1
12 | 3 | 12 | 10 | 10 | -1 | 0 | 2
(12 rows)
```

32) Using [Combinations](#) to get the solution of examples 25 to 27

Paths  $\{(6)\} \rightarrow \{10, 7\} \cup \{12\} \rightarrow \{10\}$

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost FROM edges',
  'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)',
  false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
3 | 1 | 6 | 10 | 6 | 4 | 1 | 0
4 | 2 | 6 | 10 | 7 | 8 | 1 | 1
5 | 3 | 6 | 10 | 11 | 5 | 1 | 2
6 | 4 | 6 | 10 | 10 | -1 | 0 | 3
7 | 1 | 12 | 10 | 12 | 11 | 1 | 0
8 | 2 | 12 | 10 | 11 | 5 | 1 | 1
9 | 3 | 12 | 10 | 10 | -1 | 0 | 2
(9 rows)
```

[Equivalences between signatures](#)

The following examples find the path for  $\{(6)\} \rightarrow \{10\}$

33) Using [One to One](#)

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  6, 10
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(6 rows)
```

34) Using [One to Many](#)

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  6, ARRAY[10]
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(6 rows)
```

35) Using [Many to One](#)

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  ARRAY[6], 10
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(6 rows)
```

36) Using [Many to Many](#)

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  ARRAY[6], ARRAY[10]
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(6 rows)
```

37) Using [Combinations](#)

```
SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  'SELECT * FROM (VALUES(6, 10)) AS combinations (source, target)'
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(6 rows)
```

See Also

- [Sample Data](#)
- [Boost: Dijkstra shortest paths](#)
- [https://en.wikipedia.org/wiki/Dijkstra%27s\\_algorithm](https://en.wikipedia.org/wiki/Dijkstra%27s_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

[pgr\\_dijkstraCost](#)

pgr\_dijkstraCost - Total cost of the shortest path using Dijkstra algorithm.

Availability

- Version 3.1.0
  - New proposed signature:
    - `pgr_dijkstraCost(Combinations)`
- Version 2.2.0
  - Official function.

[Description](#)

The `pgr_dijkstraCost` function summarizes of the cost of the shortest path using Dijkstra Algorithm.

Dijkstra's algorithm, conceived by Dutch computer scientist Edsger Dijkstra in 1956. It is a graph search algorithm that solves the shortest path problem for a graph with non-negative edge path costs, producing a shortest path from a starting vertex to an ending vertex. This implementation can be used with a directed graph and an undirected graph.

- Process is done only on edges with positive costs.
  - A negative value on a cost column is interpreted as the edge does not exist.
- Values are returned when there is a path.
- When there is no path:
  - When the starting vertex and ending vertex are the same.
    - The **aggregate cost** of the non included values  $\setminus((v, v))$  is  $\setminus(0)$
  - When the starting vertex and ending vertex are the different and there is no path:
    - The **aggregate cost** the non included values  $\setminus((u, v))$  is  $\setminus(\infty)$
- For optimization purposes, any duplicated value in the starting vertices or on the ending vertices are ignored.
- Running time:  $\setminus(O(|\text{start} \setminus \text{vids}| * (V \setminus \log V + E)))$
- It does not return a path.
- Returns the sum of the costs of the shortest path of each pair combination of nodes requested.
- Let be the case the values returned are stored in a table, so the unique index would be the `pair(start_vid, end_vid)`.
- Depending on the function and its parameters, the results can be symmetric.
  - The **aggregate cost** of  $\setminus((u, v))$  is the same as for  $\setminus((v, u))$ .
- Any duplicated value in the start or end vertex identifiers are ignored.
- The returned values are ordered:
  - `start_vid` ascending
  - `end_vid` ascending

 Boost Graph Inside

[Signatures](#)

Summary

```
pgr_dijkstraCost(Edges SQL, start_vid, end_vid, [directed])
pgr_dijkstraCost(Edges SQL, start_vid, end_vids, [directed])
pgr_dijkstraCost(Edges SQL, start_vids, end_vid, [directed])
```

pgr\_dijkstraCost([Edges SQL](#), start vids, end vids, [directed])  
pgr\_dijkstraCost([Edges SQL](#), [Combinations SQL](#), [directed])  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

One to One

pgr\_dijkstraCost([Edges SQL](#), start vid, end vid, [directed])

Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertex \10\ on a **directed** graph

```
SELECT * FROM pgr_dijkstraCost(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  6, 10, true);
start_vid | end_vid | agg_cost
-----+-----+-----
        6 |      10 |         5
(1 row)
```

One to Many

pgr\_dijkstraCost([Edges SQL](#), start vid, end vids, [directed])

Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertices \10, 17\ on a **directed** graph

```
SELECT * FROM pgr_dijkstraCost(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  6, ARRAY[10, 17]);
start_vid | end_vid | agg_cost
-----+-----+-----
        6 |      10 |         5
        6 |      17 |         4
(2 rows)
```

Many to One

pgr\_dijkstraCost([Edges SQL](#), start vids, end vid, [directed])

Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertices \6, 1\ to vertex \17\ on a **directed** graph

```
SELECT * FROM pgr_dijkstraCost(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  ARRAY[6, 1], 17);
start_vid | end_vid | agg_cost
-----+-----+-----
        1 |      17 |         5
        6 |      17 |         4
(2 rows)
```

Many to Many

pgr\_dijkstraCost([Edges SQL](#), start vids, end vids, [directed])

Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

From vertices \6, 1\ to vertices \10, 17\ on an **undirected** graph

```
SELECT * FROM pgr_dijkstraCost(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  ARRAY[6, 1], ARRAY[10, 17],
  directed => false);
start_vid | end_vid | agg_cost
-----+-----+-----
        1 |      10 |         4
        1 |      17 |         5
        6 |      10 |         1
        6 |      17 |         4
(4 rows)
```

Combinations

pgr\_dijkstraCost([Edges SQL](#), [Combinations SQL](#), [directed])

Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table on an **undirected** graph

The combinations table:

```
SELECT source, target FROM combinations;
source | target
-----+-----
       5 |       6
       5 |      10
       6 |       5
       6 |      15
       6 |      14
(5 rows)
```

The query:

```
SELECT * FROM pgr_dijkstraCost(
  'SELECT id, source, target, cost, reverse_cost FROM edges',
  'SELECT source, target FROM combinations',
  false);
```

start_vid	end_vid	agg_cost
5	6	1
5	10	2
6	5	1
6	15	2

(4 rows)

Parameters¶

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start_vid	BIGINT	Identifier of the starting vertex of the path.
start_vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end_vid	BIGINT	Identifier of the ending vertex of the path.
end_vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters¶

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries¶

Edges SQL¶

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	<div>Weight of the edge (target, source)<ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul></div>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL¶

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns¶

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.

Column	Type	Description
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

Additional Examples1

Example 1:

Demonstration of repeated values are ignored, and result is sorted.

```
SELECT * FROM pgr_dijkstraCost(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[7, 10, 15, 10, 10, 15], ARRAY[10, 7, 10, 15]);
start_vid | end_vid | agg_cost
```

7	10	4
7	15	3
10	7	2
10	15	3
15	7	3
15	10	1
(6 rows)		

Example 2:

Making **start\_vids** the same as **end\_vids**

```
SELECT * FROM pgr_dijkstraCost(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[7, 10, 15], ARRAY[7, 10, 15]);
start_vid | end_vid | agg_cost
```

7	10	4
7	15	3
10	7	2
10	15	3
15	7	3
15	10	1
(6 rows)		

Example 3:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_dijkstraCost(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)');
start_vid | end_vid | agg_cost
```

6	7	1
6	10	5
12	10	4
(3 rows)		

See Also5

- [Dijkstra - Family of functions](#)
- [Sample Data](#)
- [Boost: Dijkstra shortest paths](#)
- [https://en.wikipedia.org/wiki/Dijkstra%27s\\_algorithm](https://en.wikipedia.org/wiki/Dijkstra%27s_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_dijkstraCostMatrix1

pgr\_dijkstraCostMatrix - Calculates a cost matrix using [pgr\\_dijkstra](#).

Availability

- Version 3.0.0
  - Function promoted to official.
- Version 2.3.0
  - New proposed function.

Description1

Using Dijkstra algorithm, calculate and return a cost matrix.

Dijkstra’s algorithm, conceived by Dutch computer scientist Edsger Dijkstra in 1956. It is a graph search algorithm that solves the shortest path problem for a graph with non-negative edge path costs, producing a shortest path from a starting vertex to an ending vertex. This implementation can be used with a directed graph and an undirected graph.

The main Characteristics are:

- Can be used as input to [pgr\\_TSP](#).
  - Use directly when the resulting matrix is symmetric and there is no $\backslash$ ( $\infty$ ) value.
  - It will be the users responsibility to make the matrix symmetric.
    - By using geometric or harmonic average of the non symmetric values.
    - By using max or min the non symmetric values.
    - By setting the upper triangle to be the mirror image of the lower triangle.
    - By setting the lower triangle to be the mirror image of the upper triangle.
  - It is also the users responsibility to fix an $\backslash$ ( $\infty$ ) value.



- Each function works as part of the family it belongs to.
- It does not return a path.
- Returns the sum of the costs of the shortest path for pair combination of nodes in the graph.
- Process is done only on edges with positive costs.
- Values are returned when there is a path.
  - When the starting vertex and ending vertex are the same, there is no path.
    - The aggregate cost in the non included values  $(v, v)$  is 0.
  - When the starting vertex and ending vertex are the different and there is no path.
    - The aggregate cost in the non included values  $(u, v)$  is  $\backslash(\infty)$ .
- Let be the case the values returned are stored in a table:
  - The unique index would be the pair: (start\_vid, end\_vid).
- Depending on the function and its parameters, the results can be symmetric.
  - The aggregate cost of  $(u, v)$  is the same as for  $(v, u)$ .
- Any duplicated value in the **start vids** are ignored.
- The returned values are ordered:
  - start\_vid ascending
  - end\_vid ascending

Boost Graph Inside

Signatures1

Summary

pgr\_dijkstraCostMatrix([Edges SQL](#), **start vids**, [directed])

Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

Symmetric cost matrix for vertices  $\backslash(\{5, 6, 10, 15\})$  on an **undirected** graph

```
SELECT * FROM pgr_dijkstraCostMatrix(
'SELECT id, source, target, cost, reverse_cost FROM edges',
(SELECT array_agg(id)
 FROM vertices
 WHERE id IN (5, 6, 10, 15)),
false);
start_vid | end_vid | agg_cost
-----+-----+-----
5 | 6 | 1
5 | 10 | 2
5 | 15 | 3
6 | 5 | 1
6 | 10 | 1
6 | 15 | 2
10 | 5 | 2
10 | 6 | 1
10 | 15 | 1
15 | 5 | 3
15 | 6 | 2
15 | 10 | 1
(12 rows)
```

Parameters1

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<b>start vids</b>	ARRAY[BIGINT]	Array of identifiers of starting vertices.

Optional parameters1

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>• When true the graph is considered <i>Directed</i></li><li>• When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries1

Edges SQL1

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.

Column	Type	Default	Description
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

Additional Examples

Example:

Use with [pgr\\_TSP](#).

```
SELECT * FROM pgr_TSP(
$$
SELECT * FROM pgr_dijkstraCostMatrix(
'SELECT id, source, target, cost, reverse_cost FROM edges',
(SELECT array_agg(id)
FROM vertices
WHERE id IN (5, 6, 10, 15)),
false)
$$);
```

NOTICE: pgr\_TSP no longer solving with simulated annealing

HINT: Ignoring annealing parameters

seq | node | cost | agg\_cost

```
-----+-----
1 | 5 | 0 | 0
2 | 6 | 1 | 1
3 | 10 | 1 | 2
4 | 15 | 1 | 3
5 | 5 | 3 | 6
(5 rows)
```

See Also

- [Dijkstra - Family of functions](#)
- [Cost Matrix - Category](#)
- [Traveling Sales Person - Family of functions](#)
- [Sample Data](#)
- [Boost: Dijkstra shortest paths](#)

Indices and tables

- [Index](#)
- [Search Page](#)

[pgr\\_drivingDistance](#)

pgr\_drivingDistance - Returns the driving distance from a start node.

Availability

Version 3.6.0

- Standardizing output columns to (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
  - pgr\_drivingDistance(Single vertex)
    - Added depth and start\_vid result columns.
  - pgr\_drivingDistance(Multiple vertices)
    - Result column name change: from\_v to start\_vid.
    - Added depth and pred result columns.

Version 2.1.0

- Signature change:
  - pgr\_drivingDistance(single vertex)
- New official signature:
  - pgr\_drivingDistance(multiple vertices)

Version 2.0.0

- Official function.

Description

Using the Dijkstra algorithm, extracts all the nodes that have costs less than or equal to the value `distance`. The edges extracted will conform to the corresponding spanning tree.

Boost Graph Inside

Signatures

`pgr_drivingDistance`([Edges SQL](#), **Root vid**, **distance**, [directed])  
`pgr_drivingDistance`([Edges SQL](#), **Root vids**, **distance**, [options])  
**options**: [directed, equicost]  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Single Vertex

`pgr_drivingDistance`([Edges SQL](#), **Root vid**, **distance**, [directed])  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

From vertex `\(11\)` for a distance of `\(3.0\)`

```
SELECT * FROM pgr_drivingDistance(
'SELECT id, source, target, cost, reverse_cost FROM edges',
11, 3.0);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
 1 |    0 |      11 |   11 |   11 |   -1 |    0 |         0
 2 |    1 |      11 |   11 |    7 |    8 |    1 |         1
 3 |    1 |      11 |   11 |   12 |   11 |    1 |         1
 4 |    1 |      11 |   11 |   16 |    9 |    1 |         1
 5 |    2 |      11 |    7 |    3 |    7 |    1 |         2
 6 |    2 |      11 |    7 |    6 |    4 |    1 |         2
 7 |    2 |      11 |    7 |   10 |   10 |    1 |         2
 8 |    2 |      11 |   16 |   15 |   16 |    1 |         2
 9 |    2 |      11 |   16 |   17 |   15 |    1 |         2
10 |    3 |      11 |    3 |    1 |    6 |    1 |         3
11 |    3 |      11 |    6 |    5 |    1 |    1 |         3
12 |    3 |      11 |    8 |    9 |   14 |    1 |         3
13 |    3 |      11 |   15 |   10 |    3 |    1 |         3
(13 rows)
```

Multiple Vertices

`pgr_drivingDistance`([Edges SQL](#), **Root vids**, **distance**, [options])  
**options**: [directed, equicost]  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

From vertices `\(11, 16\)` for a distance of `\(3.0\)` with equi-cost on a directed graph

```
SELECT * FROM pgr_drivingDistance(
'SELECT id, source, target, cost, reverse_cost FROM edges',
array[11, 16], 3.0, equicost => true);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
 1 |    0 |      11 |   11 |   11 |   -1 |    0 |         0
 2 |    1 |      11 |   11 |    7 |    8 |    1 |         1
 3 |    1 |      11 |   11 |   12 |   11 |    1 |         1
 4 |    2 |      11 |    7 |    3 |    7 |    1 |         2
 5 |    2 |      11 |    7 |    6 |    4 |    1 |         2
 6 |    2 |      11 |    7 |   10 |   10 |    1 |         2
 7 |    3 |      11 |    3 |    1 |    6 |    1 |         3
 8 |    3 |      11 |    6 |    5 |    1 |    1 |         3
 9 |    3 |      11 |    8 |    9 |   14 |    1 |         3
10 |    0 |      16 |   16 |   16 |   -1 |    0 |         0
11 |    1 |      16 |   16 |   15 |   16 |    1 |         1
12 |    1 |      16 |   16 |   17 |   15 |    1 |         1
13 |    2 |      16 |   15 |   10 |    3 |    1 |         2
(13 rows)
```

Parameters

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	Edges SQL as described below.
<b>Root vid</b>	BIGINT	Identifier of the root vertex of the tree.
<b>Root vids</b>	ARRAY[ANY-INTEGER]	Array of identifiers of the root vertices. <ul style="list-style-type: none"><li>• <code>\(0\)</code> values are ignored</li><li>• For optimization purposes, any duplicated value is ignored.</li></ul>
<b>distance</b>	FLOAT	Upper limit for the inclusion of a node in the result.

Where:

ANY-NUMERIC:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Optional parameters

Column	Type	Default	Description
--------	------	---------	-------------

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true the graph is considered <i>Directed</i></li> <li>When false the graph is considered as <i>Undirected</i>.</li> </ul>

Driving distance optional parameters

Column	Type	Default	Description
equicost	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true the node will only appear in the closeststart_vid list. Tie brakes are arbitrary.</li> <li>When false which resembles several calls using the single vertex signature.</li> </ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Parameter	Type	Description
seq	BIGINT	Sequential value starting from \(\backslash\).
depth	BIGINT	Depth of the node. <ul style="list-style-type: none"> <li>\(0\backslash\) when node = start_vid.</li> <li>\(depth-1\backslash\) is the depth of pred</li> </ul>
start_vid	BIGINT	Identifier of the root vertex.
pred	BIGINT	Predecessor of node. <ul style="list-style-type: none"> <li>When node = start_vid then has the value node.</li> </ul>
node	BIGINT	Identifier of node reached using edge.
edge	BIGINT	Identifier of the edge used to arrive from pred to node. <ul style="list-style-type: none"> <li>\(-1\backslash\) when node = start_vid.</li> </ul>
cost	FLOAT	Cost to traverse edge.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Additional Examples

Example:

From vertices \(\backslash{11, 16}\backslash\backslash\) for a distance of \(\backslash{3.0}\backslash\) on an undirected graph

```

SELECT * FROM pgr_drivingDistance(
'SELECT id, source, target, cost, reverse_cost FROM edges',
array[11, 16], 3.0, directed => false);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 11 | 11 | 11 | -1 | 0 | 0
2 | 1 | 11 | 11 | 7 | 8 | 1 | 1
3 | 1 | 11 | 11 | 10 | 5 | 1 | 1
4 | 1 | 11 | 11 | 12 | 11 | 1 | 1

```

5	1	11	11	16	9	1	1
6	2	11	7	3	7	1	2
7	2	11	10	6	2	1	2
8	2	11	7	8	10	1	2
9	2	11	10	15	3	1	2
10	2	11	16	17	15	1	2
11	3	11	3	1	6	1	3
12	3	11	6	5	1	1	3
13	3	11	8	9	14	1	3
14	0	16	16	16	-1	0	0
15	1	16	16	11	9	1	1
16	1	16	16	15	16	1	1
17	1	16	16	17	15	1	1
18	2	16	11	7	8	1	2
19	2	16	11	10	5	1	2
20	2	16	17	12	13	1	2
21	3	16	7	3	7	1	3
22	3	16	7	6	4	1	3
23	3	16	7	8	10	1	3

(23 rows)

See Also

- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_KSP

pgr\_ksp — Yen’s algorithm for K shortest paths using Dijkstra.

Availability

Version 3.6.0

- Standardizing output columns to (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)
- pgr\_ksp(One to One)
  - Added start\_vid and end\_vid result columns.
- New proposed signatures:
  - pgr\_ksp(One to Many)
  - pgr\_ksp(Many to One)
  - pgr\_ksp(Many to Many)
  - pgr\_ksp(Combinations)

Version 2.1.0

- Signature change
  - Old signature no longer supported

Version 2.0.0

- Official function.

Description

The K shortest path routing algorithm based on Yen’s algorithm. “K” is the number of shortest paths desired.

Boost Graph Inside

Signatures

Summary

pgr\_KSP([Edges SQL](#), start\_vid, end\_vid, K, [options])  
pgr\_KSP([Edges SQL](#), start\_vid, end\_vids, K, [options])  
pgr\_KSP([Edges SQL](#), start\_vids, end\_vid, K, [options])  
pgr\_KSP([Edges SQL](#), start\_vids, end\_vids, K, [options])  
pgr\_KSP([Edges SQL](#), [Combinations SQL](#), K, [options])  
**options:** [directed, heap\_paths]  
Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

One to One

pgr\_KSP([Edges SQL](#), start\_vid, end\_vid, K, [options])  
**options:** [directed, heap\_paths]  
Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Get 2 paths from \6\ to \17\ on a directed graph.

```
SELECT * FROM pgr_KSP(
'SELECT id, source, target, cost, reverse_cost FROM edges',
6, 17, 2);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 6 | 17 | 6 | 17 | 4 | 0
2 | 1 | 2 | 6 | 17 | 7 | 10 | 1 | 1
3 | 1 | 3 | 6 | 17 | 8 | 12 | 1 | 2
4 | 1 | 4 | 6 | 17 | 12 | 13 | 1 | 3
5 | 1 | 5 | 6 | 17 | 17 | -1 | 0 | 4
6 | 2 | 1 | 6 | 17 | 6 | 4 | 1 | 0
7 | 2 | 2 | 6 | 17 | 7 | 8 | 1 | 1
8 | 2 | 3 | 6 | 17 | 11 | 9 | 1 | 2
9 | 2 | 4 | 6 | 17 | 16 | 15 | 1 | 3
10 | 2 | 5 | 6 | 17 | 17 | -1 | 0 | 4
(10 rows)
```

One to Many

pgr\_KSP([Edges SQL](#), start vid, end vids, K, [options])  
**options:** [directed, heap\_paths]  
Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Get 2 paths from vertex \6) to vertices \10, 17) on a directed graph.

```
SELECT * FROM pgr_KSP(
'select id, source, target, cost, reverse_cost from edges',
6, ARRAY[10, 17], 2);
```

seq	path_id	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	1	6	10	6	4	1	0
2	1	2	6	10	7	8	1	1
3	1	3	6	10	11	9	1	2
4	1	4	6	10	16	16	1	3
5	1	5	6	10	15	3	1	4
6	1	6	6	10	10	-1	0	5
7	2	1	6	10	6	4	1	0
8	2	2	6	10	7	10	1	1
9	2	3	6	10	8	12	1	2
10	2	4	6	10	12	13	1	3
11	2	5	6	10	17	15	1	4
12	2	6	6	10	16	16	1	5
13	2	7	6	10	15	3	1	6
14	2	8	6	10	10	-1	0	7
15	3	1	6	17	6	4	1	0
16	3	2	6	17	7	10	1	1
17	3	3	6	17	8	12	1	2
18	3	4	6	17	12	13	1	3
19	3	5	6	17	17	-1	0	4
20	4	1	6	17	6	4	1	0
21	4	2	6	17	7	8	1	1
22	4	3	6	17	11	9	1	2
23	4	4	6	17	16	15	1	3
24	4	5	6	17	17	-1	0	4

(24 rows)

Many to One

pgr\_KSP([Edges SQL](#), start vids, end vid, K, [options])  
**options:** [directed, heap\_paths]  
Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Get 2 paths from vertices \6, 1) to vertex \17) on a directed graph.

```
SELECT * FROM pgr_KSP(
'select id, source, target, cost, reverse_cost from edges',
ARRAY[6, 1], 17, 2);
```

seq	path_id	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	1	1	17	1	6	1	0
2	1	2	1	17	3	7	1	1
3	1	3	1	17	7	10	1	2
4	1	4	1	17	8	12	1	3
5	1	5	1	17	12	13	1	4
6	1	6	1	17	17	-1	0	5
7	2	1	1	17	1	6	1	0
8	2	2	1	17	3	7	1	1
9	2	3	1	17	7	8	1	2
10	2	4	1	17	11	9	1	3
11	2	5	1	17	16	15	1	4
12	2	6	1	17	17	-1	0	5
13	3	1	6	17	6	4	1	0
14	3	2	6	17	7	10	1	1
15	3	3	6	17	8	12	1	2
16	3	4	6	17	12	13	1	3
17	3	5	6	17	17	-1	0	4
18	4	1	6	17	6	4	1	0
19	4	2	6	17	7	8	1	1
20	4	3	6	17	11	9	1	2
21	4	4	6	17	16	15	1	3
22	4	5	6	17	17	-1	0	4

(22 rows)

Many to Many

pgr\_KSP([Edges SQL](#), start vids, end vids, K, [options])  
**options:** [directed, heap\_paths]  
Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Get 2 paths vertices \6, 1) to vertices \10, 17) on a directed graph.

```
SELECT * FROM pgr_KSP(
'select id, source, target, cost, reverse_cost from edges',
ARRAY[6, 1], ARRAY[10, 17], 2);
```

seq	path_id	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	1	1	10	1	6	1	0
2	1	2	1	10	3	7	1	1
3	1	3	1	10	7	8	1	2
4	1	4	1	10	11	9	1	3
5	1	5	1	10	16	16	1	4
6	1	6	1	10	15	3	1	5
7	1	7	1	10	10	-1	0	6
8	2	1	1	10	1	6	1	0
9	2	2	1	10	3	7	1	1
10	2	3	1	10	7	10	1	2
11	2	4	1	10	8	12	1	3
12	2	5	1	10	12	13	1	4
13	2	6	1	10	17	15	1	5
14	2	7	1	10	16	16	1	6
15	2	8	1	10	15	3	1	7
16	2	9	1	10	10	-1	0	8
17	3	1	1	17	1	6	1	0
18	3	2	1	17	3	7	1	1
19	3	3	1	17	7	10	1	2
20	3	4	1	17	8	12	1	3
21	3	5	1	17	12	13	1	4
22	3	6	1	17	17	-1	0	5
23	4	1	1	17	1	6	1	0
24	4	2	1	17	3	7	1	1
25	4	3	1	17	7	8	1	2
26	4	4	1	17	11	9	1	3
27	4	5	1	17	16	15	1	4

28	4	6	1	17	17	-1	0	5
29	5	1	6	10	6	4	1	0
30	5	2	6	10	7	8	1	1
31	5	3	6	10	11	9	1	2
32	5	4	6	10	16	16	1	3
33	5	5	6	10	15	3	1	4
34	5	6	6	10	10	-1	0	5
35	6	1	6	10	6	4	1	0
36	6	2	6	10	7	10	1	1
37	6	3	6	10	8	12	1	2
38	6	4	6	10	12	13	1	3
39	6	5	6	10	17	15	1	4
40	6	6	6	10	16	16	1	5
41	6	7	6	10	15	3	1	6
42	6	8	6	10	10	-1	0	7
43	7	1	6	17	6	4	1	0
44	7	2	6	17	7	10	1	1
45	7	3	6	17	8	12	1	2
46	7	4	6	17	12	13	1	3
47	7	5	6	17	17	-1	0	4
48	8	1	6	17	6	4	1	0
49	8	2	6	17	7	8	1	1
50	8	3	6	17	11	9	1	2
51	8	4	6	17	16	15	1	3
52	8	5	6	17	17	-1	0	4

(52 rows)

Combinations

pgr\_KSP([Edges SQL](#), [Combinations SQL](#), K, [options])  
**options:** [directed, heap\_paths]  
Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:  
Using a combinations table on an directed graph

The combinations table:  
SELECT source, target FROM combinations;  
source | target  
-----+-----  
5 | 6  
5 | 10  
6 | 5  
6 | 15  
6 | 14  
(5 rows)

The query:  
SELECT \* FROM pgr\_KSP(  
  'SELECT id, source, target, cost, reverse\_cost FROM edges',  
  'SELECT source, target FROM combinations', 2);  
seq | path\_id | path\_seq | start\_vid | end\_vid | node | edge | cost | agg\_cost  
-----+-----+-----+-----+-----+-----+-----+-----+-----  
1 | 1 | 1 | 5 | 6 | 5 | 1 | 1 | 0  
2 | 1 | 2 | 5 | 6 | 6 | -1 | 0 | 1  
3 | 2 | 1 | 5 | 10 | 5 | 1 | 1 | 0  
4 | 2 | 2 | 5 | 10 | 6 | 4 | 1 | 1  
5 | 2 | 3 | 5 | 10 | 7 | 8 | 1 | 2  
6 | 2 | 4 | 5 | 10 | 11 | 9 | 1 | 3  
7 | 2 | 5 | 5 | 10 | 16 | 16 | 1 | 4  
8 | 2 | 6 | 5 | 10 | 15 | 3 | 1 | 5  
9 | 2 | 7 | 5 | 10 | 10 | -1 | 0 | 6  
10 | 3 | 1 | 5 | 10 | 5 | 1 | 1 | 0  
11 | 3 | 2 | 5 | 10 | 6 | 4 | 1 | 1  
12 | 3 | 3 | 5 | 10 | 7 | 10 | 1 | 2  
13 | 3 | 4 | 5 | 10 | 8 | 12 | 1 | 3  
14 | 3 | 5 | 5 | 10 | 12 | 13 | 1 | 4  
15 | 3 | 6 | 5 | 10 | 17 | 15 | 1 | 5  
16 | 3 | 7 | 5 | 10 | 16 | 16 | 1 | 6  
17 | 3 | 8 | 5 | 10 | 15 | 3 | 1 | 7  
18 | 3 | 9 | 5 | 10 | 10 | -1 | 0 | 8  
19 | 4 | 1 | 6 | 5 | 6 | 1 | 1 | 0  
20 | 4 | 2 | 6 | 5 | 5 | -1 | 0 | 1  
21 | 5 | 1 | 6 | 15 | 6 | 4 | 1 | 0  
22 | 5 | 2 | 6 | 15 | 7 | 8 | 1 | 1  
23 | 5 | 3 | 6 | 15 | 11 | 9 | 1 | 2  
24 | 5 | 4 | 6 | 15 | 16 | 16 | 1 | 3  
25 | 5 | 5 | 6 | 15 | 15 | -1 | 0 | 4  
26 | 6 | 1 | 6 | 15 | 6 | 4 | 1 | 0  
27 | 6 | 2 | 6 | 15 | 7 | 10 | 1 | 1  
28 | 6 | 3 | 6 | 15 | 8 | 12 | 1 | 2  
29 | 6 | 4 | 6 | 15 | 12 | 13 | 1 | 3  
30 | 6 | 5 | 6 | 15 | 17 | 15 | 1 | 4  
31 | 6 | 6 | 6 | 15 | 16 | 16 | 1 | 5  
32 | 6 | 7 | 6 | 15 | 15 | -1 | 0 | 6  
(32 rows)

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	SQL query as described.
start vid	ANY-INTEGER	Identifier of the departure vertex.
end vid	ANY-INTEGER	Identifier of the destination vertex.
K	ANY-INTEGER	Number of required paths.

Where:  
ANY-INTEGER:  
SMALLINT, INTEGER, BIGINT

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

KSP Optional parameters

Column	Type	Default	Description
heap_paths	BOOLEAN	false	<ul style="list-style-type: none"><li>When false Returns at most K paths.</li><li>When true all the calculated paths while processing are returned.</li><li>Roughly, when the shortest path hasN edges, the heap will contain about thanN * K paths for small value ofK and K &gt; 5.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Path identifier. <ul style="list-style-type: none"><li>Has value 1 for the first of a path fromstart_vid to end_vid</li></ul>
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence. <ul style="list-style-type: none"><li>\(0\) for the last node of the path.</li></ul>



Column Type Description

agg\_cost FLOAT Aggregate cost from **start vid** to node.

Additional Examples¶

Example:

Get 2 paths from  $\backslash(6\backslash)$  to  $\backslash(17\backslash)$  on an undirected graph

Also get the paths in the heap.

```
SELECT * FROM pgr_KSP(
'SELECT id, source, target, cost, reverse_cost FROM edges',
6, 17, 2,
directed => false, heap_paths => true
);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 6 | 17 | 6 | 4 | 1 | 0
2 | 1 | 2 | 6 | 17 | 7 | 10 | 1 | 1
3 | 1 | 3 | 6 | 17 | 8 | 12 | 1 | 2
4 | 1 | 4 | 6 | 17 | 12 | 13 | 1 | 3
5 | 1 | 5 | 6 | 17 | 17 | -1 | 0 | 4
6 | 2 | 1 | 6 | 17 | 6 | 4 | 1 | 0
7 | 2 | 2 | 6 | 17 | 7 | 8 | 1 | 1
8 | 2 | 3 | 6 | 17 | 11 | 11 | 1 | 2
9 | 2 | 4 | 6 | 17 | 12 | 13 | 1 | 3
10 | 2 | 5 | 6 | 17 | 17 | -1 | 0 | 4
11 | 3 | 1 | 6 | 17 | 6 | 4 | 1 | 0
12 | 3 | 2 | 6 | 17 | 7 | 8 | 1 | 1
13 | 3 | 3 | 6 | 17 | 11 | 9 | 1 | 2
14 | 3 | 4 | 6 | 17 | 16 | 15 | 1 | 3
15 | 3 | 5 | 6 | 17 | 17 | -1 | 0 | 4
16 | 4 | 1 | 6 | 17 | 6 | 2 | 1 | 0
17 | 4 | 2 | 6 | 17 | 10 | 5 | 1 | 1
18 | 4 | 3 | 6 | 17 | 11 | 9 | 1 | 2
19 | 4 | 4 | 6 | 17 | 16 | 15 | 1 | 3
20 | 4 | 5 | 6 | 17 | 17 | -1 | 0 | 4
(20 rows)
```

Example:

Get 2 paths using combinations table on an undirected graph

Also get the paths in the heap.

```
SELECT * FROM pgr_KSP(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT source, target FROM combinations', 2, directed => false, heap_paths => true);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 5 | 10 | 5 | 1 | 1 | 0
2 | 1 | 2 | 5 | 10 | 6 | -1 | 0 | 1
3 | 2 | 1 | 5 | 10 | 5 | 1 | 1 | 0
4 | 2 | 2 | 5 | 10 | 6 | 2 | 1 | 1
5 | 2 | 3 | 5 | 10 | 10 | -1 | 0 | 2
6 | 3 | 1 | 5 | 10 | 5 | 1 | 1 | 0
7 | 3 | 2 | 5 | 10 | 6 | 4 | 1 | 1
8 | 3 | 3 | 5 | 10 | 7 | 8 | 1 | 2
9 | 3 | 4 | 5 | 10 | 11 | 5 | 1 | 3
10 | 3 | 5 | 5 | 10 | 10 | -1 | 0 | 4
11 | 4 | 1 | 6 | 5 | 6 | 1 | 1 | 0
12 | 4 | 2 | 6 | 5 | 5 | -1 | 0 | 1
13 | 5 | 1 | 6 | 15 | 6 | 2 | 1 | 0
14 | 5 | 2 | 6 | 15 | 10 | 3 | 1 | 1
15 | 5 | 3 | 6 | 15 | 15 | -1 | 0 | 2
16 | 6 | 1 | 6 | 15 | 6 | 4 | 1 | 0
17 | 6 | 2 | 6 | 15 | 7 | 8 | 1 | 1
18 | 6 | 3 | 6 | 15 | 11 | 9 | 1 | 2
19 | 6 | 4 | 6 | 15 | 16 | 16 | 1 | 3
20 | 6 | 5 | 6 | 15 | 15 | -1 | 0 | 4
21 | 7 | 1 | 6 | 15 | 6 | 2 | 1 | 0
22 | 7 | 2 | 6 | 15 | 10 | 5 | 1 | 1
23 | 7 | 3 | 6 | 15 | 11 | 9 | 1 | 2
24 | 7 | 4 | 6 | 15 | 16 | 16 | 1 | 3
25 | 7 | 5 | 6 | 15 | 15 | -1 | 0 | 4
(25 rows)
```

Example:

Get 2 paths from vertices  $\backslash(\{6, 1\}\backslash)$  to vertex  $\backslash(17\backslash)$  on a undirected graph.

```
SELECT * FROM pgr_KSP(
'select id, source, target, cost, reverse_cost from edges',
ARRAY[6, 1], 17, 2, directed => false);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | 17 | 1 | 6 | 1 | 0
2 | 1 | 2 | 1 | 17 | 3 | 7 | 1 | 1
3 | 1 | 3 | 1 | 17 | 7 | 10 | 1 | 2
4 | 1 | 4 | 1 | 17 | 8 | 12 | 1 | 3
5 | 1 | 5 | 1 | 17 | 12 | 13 | 1 | 4
6 | 1 | 6 | 1 | 17 | 17 | -1 | 0 | 5
7 | 2 | 1 | 1 | 17 | 1 | 6 | 1 | 0
8 | 2 | 2 | 1 | 17 | 3 | 7 | 1 | 1
9 | 2 | 3 | 1 | 17 | 7 | 8 | 1 | 2
10 | 2 | 4 | 1 | 17 | 11 | 9 | 1 | 3
11 | 2 | 5 | 1 | 17 | 16 | 15 | 1 | 4
12 | 2 | 6 | 1 | 17 | 17 | -1 | 0 | 5
13 | 3 | 1 | 6 | 17 | 6 | 4 | 1 | 0
14 | 3 | 2 | 6 | 17 | 7 | 10 | 1 | 1
15 | 3 | 3 | 6 | 17 | 8 | 12 | 1 | 2
16 | 3 | 4 | 6 | 17 | 12 | 13 | 1 | 3
17 | 3 | 5 | 6 | 17 | 17 | -1 | 0 | 4
18 | 4 | 1 | 6 | 17 | 6 | 4 | 1 | 0
19 | 4 | 2 | 6 | 17 | 7 | 8 | 1 | 1
20 | 4 | 3 | 6 | 17 | 11 | 11 | 1 | 2
21 | 4 | 4 | 6 | 17 | 12 | 13 | 1 | 3
22 | 4 | 5 | 6 | 17 | 17 | -1 | 0 | 4
(22 rows)
```

Example:

Get 2 paths vertices  $\backslash(\{6, 1\}\backslash)$  to vertices  $\backslash(\{10, 17\}\backslash)$  on a directed graph.

Also get the paths in the heap.

```
SELECT * FROM pgr_KSP(
'select id, source, target, cost, reverse_cost from edges',
ARRAY[6, 1], ARRAY[10, 17], 2, heap_paths => true);
```

seq	path_id	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	1	1	10	1	6	1	0
2	1	2	1	10	3	7	1	1
3	1	3	1	10	7	8	1	2
4	1	4	1	10	11	9	1	3
5	1	5	1	10	16	16	1	4
6	1	6	1	10	15	3	1	5
7	1	7	1	10	10	-1	0	6
8	2	1	1	10	1	6	1	0
9	2	2	1	10	3	7	1	1
10	2	3	1	10	7	10	1	2
11	2	4	1	10	8	12	1	3
12	2	5	1	10	12	13	1	4
13	2	6	1	10	17	15	1	5
14	2	7	1	10	16	16	1	6
15	2	8	1	10	15	3	1	7
16	2	9	1	10	10	-1	0	8
17	3	1	1	10	1	6	1	0
18	3	2	1	10	3	7	1	1
19	3	3	1	10	7	8	1	2
20	3	4	1	10	11	11	1	3
21	3	5	1	10	12	13	1	4
22	3	6	1	10	17	15	1	5
23	3	7	1	10	16	16	1	6
24	3	8	1	10	15	3	1	7
25	3	9	1	10	10	-1	0	8
26	4	1	1	17	1	6	1	0
27	4	2	1	17	3	7	1	1
28	4	3	1	17	7	10	1	2
29	4	4	1	17	8	12	1	3
30	4	5	1	17	12	13	1	4
31	4	6	1	17	17	-1	0	5
32	5	1	1	17	1	6	1	0
33	5	2	1	17	3	7	1	1
34	5	3	1	17	7	8	1	2
35	5	4	1	17	11	11	1	3
36	5	5	1	17	12	13	1	4
37	5	6	1	17	17	-1	0	5
38	6	1	1	17	1	6	1	0
39	6	2	1	17	3	7	1	1
40	6	3	1	17	7	8	1	2
41	6	4	1	17	11	9	1	3
42	6	5	1	17	16	15	1	4
43	6	6	1	17	17	-1	0	5
44	7	1	6	10	6	4	1	0
45	7	2	6	10	7	8	1	1
46	7	3	6	10	11	9	1	2
47	7	4	6	10	16	16	1	3
48	7	5	6	10	15	3	1	4
49	7	6	6	10	10	-1	0	5
50	8	1	6	10	6	4	1	0
51	8	2	6	10	7	10	1	1
52	8	3	6	10	8	12	1	2
53	8	4	6	10	12	13	1	3
54	8	5	6	10	17	15	1	4
55	8	6	6	10	16	16	1	5
56	8	7	6	10	15	3	1	6
57	8	8	6	10	10	-1	0	7
58	9	1	6	10	6	4	1	0
59	9	2	6	10	7	8	1	1
60	9	3	6	10	11	11	1	2
61	9	4	6	10	12	13	1	3
62	9	5	6	10	17	15	1	4
63	9	6	6	10	16	16	1	5
64	9	7	6	10	15	3	1	6
65	9	8	6	10	10	-1	0	7
66	10	1	6	17	6	4	1	0
67	10	2	6	17	7	10	1	1
68	10	3	6	17	8	12	1	2
69	10	4	6	17	12	13	1	3
70	10	5	6	17	17	-1	0	4
71	11	1	6	17	6	4	1	0
72	11	2	6	17	7	8	1	1
73	11	3	6	17	11	11	1	2
74	11	4	6	17	12	13	1	3
75	11	5	6	17	17	-1	0	4
76	12	1	6	17	6	4	1	0
77	12	2	6	17	7	8	1	1
78	12	3	6	17	11	9	1	2
79	12	4	6	17	16	15	1	3
80	12	5	6	17	17	-1	0	4

(80 rows)

See Also

- [K shortest paths - Category](#)
- [Sample Data](#)
- [https://en.wikipedia.org/wiki/K\\_shortest\\_path\\_routing](https://en.wikipedia.org/wiki/K_shortest_path_routing)

Indices and tables

- [Index](#)
- [Search Page](#)

pg\_r\_dijkstraVia - Proposed

pg\_r\_dijkstraVia — Route that goes through a list of vertices.

Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.

- Documentation might need refinement.

Availability

- Version 2.2.0
- New proposed function.

Description

Given a list of vertices and a graph, this function is equivalent to finding the shortest path between  $(vertex\_i)$  and  $(vertex_{i+1})$  for all  $(i < size\_of(via,vertices))$ .

Route:

is a sequence of paths.

Path:

is a section of the route.

 Boost Graph Inside

Signatures

One Via

`pgr_dijkstraVia(Edges SQL, via vertices, [options])`  
**options:** [directed, strict, U\_turn\_on\_edge]  
Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost, route\_agg\_cost)  
OR EMPTY SET

Example:

Find the route that visits the vertices  $(5, 1, 8)$  in that order on an directed graph.

```
SELECT * FROM pgr_dijkstraVia(
'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
ARRAY[5, 1, 8]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 5 | 1 | 5 | 1 | 0 | 0 | 0
2 | 1 | 2 | 5 | 1 | 6 | 4 | 1 | 1 | 1
3 | 1 | 3 | 5 | 1 | 7 | 7 | 1 | 2 | 2
4 | 1 | 4 | 5 | 1 | 3 | 6 | 1 | 3 | 3
5 | 1 | 5 | 5 | 1 | 1 | -1 | 0 | 4 | 4
6 | 2 | 1 | 1 | 8 | 1 | 6 | 1 | 0 | 4
7 | 2 | 2 | 1 | 8 | 3 | 7 | 1 | 1 | 5
8 | 2 | 3 | 1 | 8 | 7 | 10 | 1 | 2 | 6
9 | 2 | 4 | 1 | 8 | 8 | -2 | 0 | 3 | 7
(9 rows)
```

Parameters

Parameter	Type	Default	Description
Edges SQL	TEXT		SQL query as described.
via vertices	ARRAY [ ANY-INTEGER ]		Array of ordered vertices identifiers that are going to be visited.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>• When true the graph is considered <i>Directed</i></li><li>• When false the graph is considered as <i>Undirected</i>.</li></ul>

Via optional parameters

Parameter	Type	Default	Description
strict	BOOLEAN	false	<ul style="list-style-type: none"><li>• When true if a path is missing stops and returns <b>EMPTY SET</b></li><li>• When false ignores missing paths returning all paths found</li></ul>
U_turn_on_edge	BOOLEAN	true	<ul style="list-style-type: none"><li>• When true departing from a visited vertex will not try to avoid</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.

Column	Type	Default	Description
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns¶

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Identifier of a path. Has value 1 for the first path.
path_seq	INTEGER	Relative position in the path. Has value 1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex of the path.
end_vid	BIGINT	Identifier of the ending vertex of the path.
node	BIGINT	Identifier of the node in the path from start_vid to end_vid.  Identifier of the edge used to go from node to the next node in the path sequence.
edge	BIGINT	<ul style="list-style-type: none"><li>-1 for the last node of the path.</li><li>-2 for the last node of the route.</li></ul>
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.
route_agg_cost	FLOAT	Total cost from start_vid of seq = 1 to end_vid of the current seq.

Additional Examples¶

- [The main query](#)
  - [Aggregate cost of the third path.](#)
  - [Route's aggregate cost of the route at the end of the third path.](#)
  - [Nodes visited in the route.](#)
  - [The aggregate costs of the route when the visited vertices are reached.](#)
  - [Status of "passes in front" or "visits" of the nodes.](#)

All this examples are about the route that visits the vertices (5, 7, 1, 8, 15) in that order on a directed graph.

[The main query¶](#)

```
SELECT * FROM pgr_dijkstraVia(
  'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
  ARRAY[5, 7, 1, 8, 15]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 5 | 7 | 5 | 1 | 0 | 0 | 0
2 | 1 | 2 | 5 | 7 | 6 | 4 | 1 | 1 | 1
3 | 1 | 3 | 5 | 7 | 7 | -1 | 0 | 2 | 2
4 | 2 | 1 | 7 | 1 | 7 | 7 | 0 | 0 | 2
5 | 2 | 2 | 7 | 1 | 3 | 6 | 1 | 1 | 3
6 | 2 | 3 | 7 | 1 | 1 | -1 | 0 | 2 | 4
7 | 3 | 1 | 1 | 8 | 1 | 6 | 1 | 0 | 4
8 | 3 | 2 | 1 | 8 | 3 | 7 | 1 | 1 | 5
9 | 3 | 3 | 1 | 8 | 7 | 10 | 1 | 2 | 6
10 | 3 | 4 | 1 | 8 | 8 | -1 | 0 | 3 | 7
11 | 4 | 1 | 8 | 15 | 8 | 12 | 1 | 0 | 7
12 | 4 | 2 | 8 | 15 | 12 | 13 | 1 | 1 | 8
13 | 4 | 3 | 8 | 15 | 17 | 15 | 1 | 2 | 9
14 | 4 | 4 | 8 | 15 | 16 | 16 | 1 | 3 | 10
15 | 4 | 5 | 8 | 15 | 15 | -2 | 0 | 4 | 11
(15 rows)
```

[Aggregate cost of the third path¶](#)

```
SELECT agg_cost FROM pgr_dijkstraVia(
  'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
  ARRAY[5, 7, 1, 8, 15])
WHERE path_id = 3 AND edge < 0;
agg_cost
-----
3
(1 row)
```

Route's aggregate cost of the route at the end of the third path¶

```
SELECT route_agg_cost FROM pgr_dijkstraVia(
  'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
  ARRAY[5, 7, 1, 8, 15])
WHERE path_id = 3 AND edge < 0;
route_agg_cost
-----
7
(1 row)
```

Nodes visited in the route.¶

```
SELECT row_number() over () as node_seq, node
FROM pgr_dijkstraVia(
  'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
  ARRAY[5, 7, 1, 8, 15])
WHERE edge <> -1 ORDER BY seq;
node_seq | node
-----+-----
1 | 5
2 | 6
3 | 7
4 | 3
5 | 1
6 | 3
7 | 7
8 | 8
9 | 12
10 | 17
11 | 16
12 | 15
(12 rows)
```

The aggregate costs of the route when the visited vertices are reached.¶

```
SELECT path_id, route_agg_cost FROM pgr_dijkstraVia(
  'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
  ARRAY[5, 7, 1, 8, 15])
WHERE edge < 0;
path_id | route_agg_cost
-----+-----
1 | 2
2 | 4
3 | 7
4 | 11
(4 rows)
```

Status of "passes in front" or "visits" of the nodes.¶

```
SELECT seq, route_agg_cost, node, agg_cost ,
CASE WHEN edge = -1 THEN 'visits'
ELSE 'passes in front'
END as status
FROM pgr_dijkstraVia(
  'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
  ARRAY[5, 7, 1, 8, 15])
WHERE agg_cost <> 0 or seq = 1;
seq | route_agg_cost | node | agg_cost | status
-----+-----+-----+-----+-----
1 | 0 | 5 | 0 | passes in front
2 | 1 | 6 | 1 | passes in front
3 | 2 | 7 | 2 | visits
5 | 3 | 3 | 1 | passes in front
6 | 4 | 1 | 2 | visits
8 | 5 | 3 | 1 | passes in front
9 | 6 | 7 | 2 | passes in front
10 | 7 | 8 | 3 | visits
12 | 8 | 12 | 1 | passes in front
13 | 9 | 17 | 2 | passes in front
14 | 10 | 16 | 3 | passes in front
15 | 11 | 15 | 4 | passes in front
(12 rows)
```

See Also¶

- [Via - Category.](#)
- [Dijkstra - Family of functions.](#)
- [Sample Data](#)
- [Boost: Dijkstra shortest paths](#)
- [https://en.wikipedia.org/wiki/Dijkstra%27s\\_algorithm](https://en.wikipedia.org/wiki/Dijkstra%27s_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_dijkstraNear - Proposed¶

pgr\_dijkstraNear — Using Dijkstra's algorithm, finds the route that leads to the nearest vertex.

□ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.

- Documentation might need refinement.

Availability

- Version 3.3.0
  - Function promoted to proposed.
- Version 3.2.0
  - New experimental function.

Description

Given a graph, a starting vertex and a set of ending vertices, this function finds the shortest path from the starting vertex to the nearest ending vertex.

Characteristics

- Uses Dijkstra algorithm.
- Works for **directed** and **undirected** graphs.
- When there are more than one path to the same vertex with same cost:
  - The algorithm will return just one path
- Optionally allows to find more than one path.
  - When more than one path is to be returned:
    - Results are sorted in increasing order of:
      - aggregate cost
      - Within the same value of aggregate costs:
        - results are sorted by (source, target)
- Running time: Dijkstra running time:  $\backslash(drt = O((|E| + |V|)\log|V|))\backslash$ 
  - One to Many:  $\backslash(drt)\backslash$
  - Many to One:  $\backslash(drt)\backslash$
  - Many to Many:  $\backslash(drt * |Starting\ vids|)\backslash$
  - Combinations:  $\backslash(drt * |Starting\ vids|)\backslash$

 Boost Graph Inside

Signatures

Summary

`pgr_dijkstraNear(Edges SQL, start vid, end vids, [options A])`  
`pgr_dijkstraNear(Edges SQL, start vids, end vid, [options A])`  
`pgr_dijkstraNear(Edges SQL, start vids, end vids, [options B])`  
`pgr_dijkstraNear(Edges SQL, Combinations SQL, [options B])`  
**options A:** [directed, cap]  
**options B:** [directed, cap, global]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

One to Many

`pgr_dijkstraNear(Edges SQL, start vid, end vids, [options])`  
**options:** [directed, cap]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Departing on car from vertex  $\backslash(6)\backslash$  find the nearest subway station.

- Using a **directed** graph for car routing.
- The subway stations are on the following vertices  $\backslash(\{1, 10, 11\})\backslash$
- The defaults used:
  - `directed => true`
  - `cap => 1`

```
1 SELECT * FROM pgr_dijkstraNear(
2   'SELECT id, source, target, cost, reverse_cost FROM edges',
3   6, ARRAY(10, 11));
4 seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
5 -----+-----+-----+-----+-----+-----+-----+-----
6 1 |      1 |         6 |        11 |    6 |    4 |    1 |         0
7 2 |      2 |         6 |        11 |    7 |    8 |    1 |         1
8 3 |      3 |         6 |        11 |   11 |   -1 |    0 |         2
9(3 rows)
10
```

The result shows that station at vertex  $\backslash(11)\backslash$  is the nearest.

Many to One

`pgr_dijkstraNear(Edges SQL, start vids, end vid, [options])`  
**options:** [directed, cap]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Departing on a car from a subway station find the nearest **two** stations to vertex  $\backslash(2)\backslash$

- Using a **directed** graph for car routing.
- The subway stations are on the following vertices  $\backslash(\{1, 10, 11\})\backslash$
- On line 4: using the positional parameter: `directed` set to `true`
- In line 5: using named parameter `cap => 2`

```
1 SELECT * FROM pgr_dijkstraNear(
2   'SELECT id, source, target, cost, reverse_cost FROM edges',
3   ARRAY[10, 11, 1], 6,
4   true,
5   cap => 2);
6 seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
7 -----+-----+-----+-----+-----+-----+-----+-----
8 1 | 1 | 10 | 6 | 10 | 2 | 1 | 0
9 2 | 2 | 10 | 6 | 6 | -1 | 0 | 1
10 3 | 1 | 11 | 6 | 11 | 8 | 1 | 0
11 4 | 2 | 11 | 6 | 7 | 4 | 1 | 1
12 5 | 3 | 11 | 6 | 6 | -1 | 0 | 2
13 (5 rows)
14
```

The result shows that station at vertex\10\ is the nearest and the next best is\11\).

Many to Many

```
pgr_dijkstraNear(Edges SQL, start_vids, end_vids, [options])
options: [directed, cap, global]
Returns set of (seq, path_seq, start_vid, end_vid, node, edge, cost, agg_cost)
OR EMPTY SET
```

Example:

- Find the best pedestrian connection between two lines of buses
- Using an **undirected** graph for pedestrian routing
  - The first subway line stations are at\15, 16\
  - The second subway line stations stops are at\1, 10, 11\
  - On line 4: using the named parameter: *directed* => *false*
  - The defaults used:
    - cap => 1
    - global => true

```
1 SELECT * FROM pgr_dijkstraNear(
2   'SELECT id, source, target, cost, reverse_cost FROM edges',
3   ARRAY[15, 16], ARRAY[10, 11, 1],
4   directed => false);
5 seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
6 -----+-----+-----+-----+-----+-----+-----+-----
7 1 | 1 | 15 | 10 | 15 | 3 | 1 | 0
8 2 | 2 | 15 | 10 | 10 | -1 | 0 | 1
9 (2 rows)
10
```

For a pedestrian the best connection is to get on/off is at vertex\15\ of the first subway line and at vertex\10\ of the second subway line.

Only *one* route is returned because *global* is true and *cap* is 1

Combinations

```
pgr_dijkstraNear(Edges SQL, Combinations SQL, [options])
options: [directed, cap, global]
Returns set of (seq, path_seq, start_vid, end_vid, node, edge, cost, agg_cost)
OR EMPTY SET
```

Example:

- Find the best car connection between all the stations of two subway lines
- Using a **directed** graph for car routing.
  - The first subway line stations stops are at\1, 10, 11\
  - The second subway line stations are at\15, 16\

The combinations contents:

```
SELECT unnest(ARRAY[10, 11]) as source, target
FROM (SELECT unnest(ARRAY[15, 16]) AS target) a
UNION
SELECT unnest(ARRAY[15, 16]), target
FROM (SELECT unnest(ARRAY[10, 11]) AS target) b ORDER BY source, target;
source | target
-----+-----
1 | 15
1 | 16
10 | 15
10 | 16
11 | 15
11 | 16
15 | 1
15 | 10
15 | 11
16 | 1
16 | 10
16 | 11
(12 rows)
```

The query:

- lines 3~4 sets the start vertices to be from the first subway line and the ending vertices to be from the second subway line
- lines 6~7 sets the start vertices to be from the first subway line and the ending vertices to be from the first subway line
- On line 8: using the named parameter is *global* => *false*
- The defaults used:
  - directed => true
  - cap => 1

```
1 SELECT * FROM pgr_dijkstraNear(
2   'SELECT id, source, target, cost, reverse_cost FROM edges',
3   'SELECT unnest(ARRAY[10, 11]) as source, target
4   FROM (SELECT unnest(ARRAY[15, 16]) AS target) a
5   UNION
6   SELECT unnest(ARRAY[15, 16]), target
7   FROM (SELECT unnest(ARRAY[10, 11]) AS target) b',
8   global => false);
9 seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```





ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Returns (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex of the current path.
end_vid	BIGINT	Identifier of the ending vertex of the current path.
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

See Also

- [Dijkstra - Family of functions](#)
- [pgr\\_dijkstraNearCost - Proposed](#)
- [Sample Data](#)
- [Boost: Dijkstra shortest paths](#)
- Wikipedia: [https://en.wikipedia.org/wiki/Dijkstra%27s\\_algorithm](https://en.wikipedia.org/wiki/Dijkstra%27s_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_dijkstraNearCost - Proposed

pgr\_dijkstraNearCost — Using dijkstra algorithm, finds the route that leads to the nearest vertex.

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Availability

- Version 3.3.0
  - Function promoted to proposed.
- Version 3.2.0
  - New experimental function.

Description

Given a graph, a starting vertex and a set of ending vertices, this function finds the shortest path from the starting vertex to the nearest ending vertex.

Characteristics

- Uses Dijkstra algorithm.
- Works for **directed** and **undirected** graphs.
- When there are more than one path to the same vertex with same cost:
  - The algorithm will return just one path
- Optionally allows to find more than one path.
  - When more than one path is to be returned:
    - Results are sorted in increasing order of:
      - aggregate cost
      - Within the same value of aggregate costs:
        - results are sorted by (source, target)
- Running time: Dijkstra running time:  $O((|E| + |V|)\log|V|)$ 
  - One to Many:  $O(|E|)$
  - Many to One:  $O(|E|)$
  - Many to Many:  $O(|E| * |Starting\ vids|)$
  - Combinations:  $O(|E| * |Starting\ vids|)$

Boost Graph Inside

Signatures

Summary

`pgr_dijkstraNearCost(Edges SQL, start vid, end vids, [options A])`  
`pgr_dijkstraNearCost(Edges SQL, start vids, end vid, [options A])`  
`pgr_dijkstraNearCost(Edges SQL, start vids, end vids, [options B])`  
`pgr_dijkstraNearCost(Edges SQL, Combinations SQL, [options B])`  
**options A:** [directed, cap]  
**options B:** [directed, cap, global]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

One to Many

`pgr_dijkstraNearCost(Edges SQL, start vid, end vids, [options])`  
**options:** [directed, cap]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

- Departing on car from vertex \6\ find the nearest subway station.
- Using a **directed** graph for car routing.
  - The subway stations are on the following vertices \1, 10, 11\
  - The defaults used:
    - `directed => true`
    - `cap => 1`

```
1SELECT * FROM pgr_dijkstraNearCost(
2 'SELECT id, source, target, cost, reverse_cost FROM edges',
3 6, ARRAY[10, 11, 1]);
4 start_vid | end_vid | agg_cost
5-----+-----+-----
6      6 |    11 |         2
7(1 row)
8
```

The result shows that station at vertex \11\ is the nearest.

Many to One

`pgr_dijkstraNearCost(Edges SQL, start vids, end vid, [options])`  
**options:** [directed, cap]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

- Departing on a car from a subway station find the nearest **two** stations to vertex \6\
- Using a **directed** graph for car routing.
  - The subway stations are on the following vertices \1, 10, 11\
  - On line 4: using the positional parameter: `directed` set to `true`
  - In line 5: using named parameter `cap => 2`

```
1SELECT * FROM pgr_dijkstraNearCost(
2 'SELECT id, source, target, cost, reverse_cost FROM edges',
3 ARRAY[10, 11, 1], 6,
4 true,
5 cap => 2) ORDER BY agg_cost;
6 start_vid | end_vid | agg_cost
7-----+-----+-----
8      10 |      6 |         1
9      11 |      6 |         2
10(2 rows)
11
```

The result shows that station at vertex \10\ is the nearest and the next best is \11\.

Many to Many

pgr\_dijkstraNearCost([Edges SQL](#), start vids, end vids, [options])  
**options:** [directed, cap, global]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

- Find the best pedestrian connection between two lines of buses
- Using an **undirected** graph for pedestrian routing
  - The first subway line stations are at \(\{15, 16\}\)
  - The second subway line stations stops are at \(\{1, 10, 11\}\)
  - On line 4: using the named parameter: *directed* => *false*
  - The defaults used:
    - cap* => 1
    - global* => true

```
1 SELECT * FROM pgr_dijkstraNearCost(
2   'SELECT id, source, target, cost, reverse_cost FROM edges',
3   ARRAY[15, 16], ARRAY[10, 11, 1],
4   directed => false);
5 start_vid | end_vid | agg_cost
6 -----+-----+-----
7      15 |      10 |          1
8 (1 row)
```

For a pedestrian the best connection is to get on/off is at vertex \(\{15\}\) of the first subway line and at vertex \(\{10\}\) of the second subway line.

Only *one* route is returned because *global* is true and *cap* is 1

Combinations

pgr\_dijkstraNearCost([Edges SQL](#), [Combinations SQL](#), [options])  
**options:** [directed, cap, global]  
Returns set of (start\_vid, end\_vid, agg\_cost)  
OR EMPTY SET

Example:

- Find the best car connection between all the stations of two subway lines
- Using a **directed** graph for car routing.
  - The first subway line stations stops are at \(\{1, 10, 11\}\)
  - The second subway line stations are at \(\{15, 16\}\)

The combinations contents:

```
SELECT unnest(ARRAY[10, 11, 1]) as source, target
FROM (SELECT unnest(ARRAY[15, 16]) AS target) a
UNION
SELECT unnest(ARRAY[15, 16]), target
FROM (SELECT unnest(ARRAY[10, 11, 1]) AS target) b ORDER BY source, target;
source | target
-----+-----
1 | 15
1 | 16
10 | 15
10 | 16
11 | 15
11 | 16
15 | 1
15 | 10
15 | 11
16 | 1
16 | 10
16 | 11
(12 rows)
```

The query:

- lines 3~4 sets the start vertices to be from the first subway line and the ending vertices to be from the second subway line
- lines 6~7 sets the start vertices to be from the first subway line and the ending vertices to be from the first subway line
- On line 8: using the named parameter is *global* => *false*
- The defaults used:
  - directed* => true
  - cap* => 1

```
1 SELECT * FROM pgr_dijkstraNearCost(
2   'SELECT id, source, target, cost, reverse_cost FROM edges',
3   'SELECT unnest(ARRAY[10, 11, 1]) as source, target
4   FROM (SELECT unnest(ARRAY[15, 16]) AS target) a
5   UNION
6   SELECT unnest(ARRAY[15, 16]), target
7   FROM (SELECT unnest(ARRAY[10, 11, 1]) AS target) b',
8   global => false);
9 start_vid | end_vid | agg_cost
10 -----+-----+-----
11      11 |      16 |          1
12      15 |      10 |          1
13      16 |      11 |          1
14      10 |      16 |          2
15       1 |      16 |          4
16 (5 rows)
```

From the results:

- making a connection from the first subway line \(\{1, 10, 11\}\) to the second \(\{15, 16\}\):
  - The best connections from all the stations from the first line are: \(\{(1 \rightarrow 16) (10 \rightarrow 16) (11 \rightarrow 16)\}\)
  - The best one is \(\{(11 \rightarrow 16)\}\) with a cost of \(\{1\}\) (lines: 1)
- making a connection from the second subway line \(\{15, 16\}\) to the first \(\{1, 10, 11\}\):
  - The best connections from all the stations from the second line are: \(\{(15 \rightarrow 10) (16 \rightarrow 11)\}\)

- Both are equally good as they have the same cost. (lines: 12 and 13)

Parameters¶

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Dijkstra optional parameters¶

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>• When true the graph is considered <i>Directed</i></li><li>• When false the graph is considered as <i>Undirected</i>.</li></ul>

Near optional parameters¶

Parameter	Type	Default	Description
cap	BIGINT	1	Find at most cap number of nearest shortest paths
global	BOOLEAN	true	<ul style="list-style-type: none"><li>• When true: only cap limit results will be returned</li><li>• When false: cap limit per Start vid will be returned</li></ul>

Inner Queries¶

Edges SQL¶

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL¶

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns¶

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

See Also

- Dijkstra - Family of functions
- pgr\_dijkstraNear - Proposed
- Sample Data
- Boost: Dijkstra shortest paths
- Wikipedia: [https://en.wikipedia.org/wiki/Dijkstra%27s\\_algorithm](https://en.wikipedia.org/wiki/Dijkstra%27s_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

Introduction

Dijkstra's algorithm, conceived by Dutch computer scientist Edsger Dijkstra in 1956. It is a graph search algorithm that solves the shortest path problem for a graph with non-negative edge path costs, producing a shortest path from a starting vertex to an ending vertex. This implementation can be used with a directed graph and an undirected graph.

The main characteristics are:

- Process is done only on edges with positive costs.
  - A negative value on a cost column is interpreted as the edge does not exist.
- Values are returned when there is a path.
- When there is no path:
  - When the starting vertex and ending vertex are the same.
    - The **aggregate cost** of the non included values  $\backslash((v, v)\backslash)$  is  $\backslash(0\backslash)$
  - When the starting vertex and ending vertex are the different and there is no path:
    - The **aggregate cost** the non included values  $\backslash((u, v)\backslash)$  is  $\backslash(\infty\backslash)$
- For optimization purposes, any duplicated value in the starting vertices or on the ending vertices are ignored.
- Running time:  $\backslash(O(\backslash \text{start}\backslash \text{vids} \mid * (V \backslash \log V + E))\backslash)$

The Dijkstra family functions are based on the Dijkstra algorithm.

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start_vid	BIGINT	Identifier of the starting vertex of the path.
start_vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end_vid	BIGINT	Identifier of the ending vertex of the path.
end_vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.



- $\backslash(\text{node\_}\{|\ \pi|\} = \text{end\_}\{\text{vid}\})\backslash$
- $\backslash(\text{forall } i \neq |\ \pi|, \quad (\text{node\_}i, \text{node\_}\{i+1\}, \text{cost\_}i) \in E)$
- $\backslash(\text{edge\_}i = \begin{cases} \text{id\_}\{(\text{node\_}i, \text{node\_}\{i+1\}, \text{cost\_}i)\} & \text{when } i \neq |\ \pi| - 1 \\ \quad & \text{when } i = |\ \pi| \end{cases})\backslash$
- $\backslash(\text{cost\_}i = \text{cost\_}\{(\text{node\_}i, \text{node\_}\{i+1\})\})\backslash$
- $\backslash(\text{agg\_cost\_}i = \begin{cases} 0 & \text{when } i = 1 \\ \sum_{k=1}^i \text{cost\_}\{(\text{node\_}\{k-1\}, \text{node\_}k)\} & \text{when } i \neq 1 \end{cases})\backslash$

In other words: The algorithm returns a the shortest path between  $(\text{start\_}\{\text{vid}\})$  and  $(\text{end\_}\{\text{vid}\})$ , if it exists, in terms of a sequence of nodes and of edges,

- $\backslash(\text{path\_seq})\backslash$  indicates the relative position in the path of the  $(\text{node})$  or  $(\text{edge})$ .
- $\backslash(\text{cost})\backslash$  is the cost of the edge to be used to go to the next node.
- $\backslash(\text{agg\_cost})\backslash$  is the cost from the  $(\text{start\_}\{\text{vid}\})$  up to the node.

If there is no path, the resulting set is empty.

See Also

Indices and tables

- [Index](#)
- [Search Page](#)

## Flow - Family of functions

- [pgr\\_maxFlow](#) - Only the Max flow calculation using Push and Relabel algorithm.
- [pgr\\_boykovKolmogorov](#) - Boykov and Kolmogorov with details of flow on edges.
- [pgr\\_edmondsKarp](#) - Edmonds and Karp algorithm with details of flow on edges.
- [pgr\\_pushRelabel](#) - Push and relabel algorithm with details of flow on edges.
- Applications
  - [pgr\\_edgeDisjointPaths](#) - Calculates edge disjoint paths between two groups of vertices.
  - [pgr\\_maxCardinalityMatch](#) - Calculates a maximum cardinality matching in a graph.

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting
- [pgr\\_maxFlowMinCost - Experimental](#) - Details of flow and cost on edges.
- [pgr\\_maxFlowMinCost\\_Cost - Experimental](#) - Only the Min Cost calculation.

## pgr\_maxFlow

`pgr_maxFlow` — Calculates the maximum flow in a directed graph from the source(s) to the targets(s) using the Push Relabel algorithm.

Availability

- Version 3.2.0
  - New proposed signature:
    - `pgr_maxFlow(Combinations)`
- Version 3.0.0
  - Function promoted to official.
- Version 2.4.0
  - New proposed function.

Description

The main characteristics are:

- The graph is **directed**.

- Calculates the maximum flow from the sources to the targets.
  - When the maximum flow is **0** then there is no flow and **0** is returned.
  - There is no flow when source has the same value as target.
- Any duplicated values in source or target are ignored.
- Uses the [pgr\\_pushRelabel](#) algorithm.
- Running time:  $\backslash(O( V ^ 3)\backslash)$

Boost Graph Inside

Signatures

Summary

`pgr_maxFlow(Edges SQL, start vid, end vid)`  
`pgr_maxFlow(Edges SQL, start vid, end vids)`  
`pgr_maxFlow(Edges SQL, start vids, end vid)`  
`pgr_maxFlow(Edges SQL, start vids, end vids)`  
`pgr_maxFlow(Edges SQL, Combinations SQL)`  
RETURNS BIGINT

One to One

`pgr_maxFlow(Edges SQL, start vid, end vid)`  
RETURNS BIGINT

Example:

From vertex  $\backslash(11\backslash)$  to vertex  $\backslash(12\backslash)$

```
SELECT * FROM pgr_maxFlow(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
11, 12);
pgr_maxflow
-----
230
(1 row)
```

One to Many

`pgr_maxFlow(Edges SQL, start vid, end vids)`  
RETURNS BIGINT

Example:

From vertex  $\backslash(11\backslash)$  to vertices  $\backslash(\{5, 10, 12\}\backslash)$

```
SELECT * FROM pgr_maxFlow(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
11, ARRAY[5, 10, 12]);
pgr_maxflow
-----
340
(1 row)
```

Many to One

`pgr_maxFlow(Edges SQL, start vids, end vid)`  
RETURNS BIGINT

Example:

From vertices  $\backslash(\{11, 3, 17\}\backslash)$  to vertex  $\backslash(12\backslash)$

```
SELECT * FROM pgr_maxFlow(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
ARRAY[11, 3, 17], 12);
pgr_maxflow
-----
230
(1 row)
```

Many to Many

`pgr_maxFlow(Edges SQL, start vids, end vids)`  
RETURNS BIGINT

Example:

From vertices  $\backslash(\{11, 3, 17\}\backslash)$  to vertices  $\backslash(\{5, 10, 12\}\backslash)$

```
SELECT * FROM pgr_maxFlow(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
ARRAY[11, 3, 17], ARRAY[5, 10, 12]);
pgr_maxflow
-----
360
(1 row)
```

Combinations

`pgr_maxFlow(Edges SQL, Combinations SQL)`  
RETURNS BIGINT

Example:

Using a combinations table, equivalent to calculating result from vertices  $\backslash(\{5, 6\}\backslash)$  to vertices  $\backslash(\{10, 15, 14\}\backslash)$ .

The combinations table:

```
SELECT source, target FROM combinations
WHERE target NOT IN (5, 6);
source | target
-----+-----
5 | 10
6 | 15
6 | 14
```



(3 rows)

The query:

```
SELECT * FROM pgr_maxFlow(  
  'SELECT id, source, target, capacity, reverse_capacity  
  FROM edges;  
  'SELECT * FROM combinations WHERE target NOT IN (5, 6)');  
pgr_maxflow  
-----  
      80  
(1 row)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
capacity	ANY-INTEGER		Weight of the edge (source, target)
reverse_capacity	ANY-INTEGER	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Type	Description
BIGINT	Maximum flow possible from the source(s) to the target(s)

Additional Examples

Example:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_maxFlow(  
  'SELECT id, source, target, capacity, reverse_capacity
```

```
FROM edges',
'SELECT * FROM (VALUES (5, 10), (6, 15), (6, 14)) AS t(source, target));
pgr_maxflow
-----
      80
(1 row)
```

See Also

- [Flow - Family of functions](#)
  - [pgr\\_pushRelabel](#)
- [Boost: push relabel max flow](#)
- [https://en.wikipedia.org/wiki/Push%E2%80%93relabel\\_maximum\\_flow\\_algorithm](https://en.wikipedia.org/wiki/Push%E2%80%93relabel_maximum_flow_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_boykovKolmogorov

pgr\_boykovKolmogorov — Calculates the flow on the graph edges that maximizes the flow from the sources to the targets using Boykov Kolmogorov algorithm.

Availability

- Version 3.2.0
  - New proposed signature:
    - pgr\_boykovKolmogorov(Combinations)
- Version 3.0.0
  - Function promoted to official.
- Version 2.5.0
  - Renamed from pgr\_maxFlowBoykovKolmogorov
  - Function promoted to proposed.
- Version 2.3.0
  - New experimental function.

Description

The main characteristics are:

- The graph is **directed**.
- Process is done only on edges with positive capacities.
- When the maximum flow is 0 then there is no flow and**EMPTY SET** is returned.
  - There is no flow when source has the same value as target.
- Any duplicated values in source or target are ignored.
- Calculates the flow/residual capacity for each edge. In the output
  - Edges with zero flow are omitted.
- Creates
  - a **super source** and edges from it to all the sources,
  - a **super target** and edges from it to all the targetss.
- The maximum flow through the graph is guaranteed to be the value returned by[pgr\\_maxFlow](#) when executed with the same parameters and can be calculated:
  - By aggregation of the outgoing flow from the sources
  - By aggregation of the incoming flow to the targets
- Running time: Polynomial

Boost Graph Inside

Signatures

Summary

```
pgr_boykovKolmogorov(Edges SQL, start vid, end vid)
pgr_boykovKolmogorov(Edges SQL, start vid, end vids)
pgr_boykovKolmogorov(Edges SQL, start vids, end vid)
pgr_boykovKolmogorov(Edges SQL, start vids, end vids)
pgr_boykovKolmogorov(Edges SQL, Combinations SQL)
Returns set of (seq, edge, start_vid, end_vid, flow, residual_capacity)
OR EMPTY SET
```

One to One

```
pgr_boykovKolmogorov(Edges SQL, start vid, end vid)
Returns set of (seq, edge, start_vid, end_vid, flow, residual_capacity)
OR EMPTY SET
```

Example:

```
From vertex \{11\} to vertex \{12\}

SELECT * FROM pgr_boykovKolmogorov(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
11, 12);
seq | edge | start_vid | end_vid | flow | residual_capacity
-----+-----+-----+-----+-----+-----
1 | 10 | 7 | 8 | 100 | 30
2 | 12 | 8 | 12 | 100 | 0
3 | 8 | 11 | 7 | 100 | 30
```

4	11	11	12	130	0
---	----	----	----	-----	---

(4 rows)

One to Many

pgr\_boykovKolmogorov([Edges SQL](#), start\_vid, end\_vids)  
Returns set of (seq, edge, start\_vid, end\_vid, flow, residual\_capacity)  
OR EMPTY SET

Example:

From vertex \{11\} to vertices \{5, 10, 12\}

```
SELECT * FROM pgr_boykovKolmogorov(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
11, ARRAY[5, 10, 12]);
```

seq	edge	start_vid	end_vid	flow	residual_capacity
1	1	6	5	50	80
2	4	7	6	50	0
3	10	7	8	80	50
4	12	8	12	80	20
5	8	11	7	130	0
6	11	11	12	130	0
7	9	11	16	80	50
8	3	15	10	80	50
9	16	16	15	80	0

(9 rows)

Many to One

pgr\_boykovKolmogorov([Edges SQL](#), start\_vids, end\_vid)  
Returns set of (seq, edge, start\_vid, end\_vid, flow, residual\_capacity)  
OR EMPTY SET

Example:

From vertices \{11, 3, 17\} to vertex \{12\}

```
SELECT * FROM pgr_boykovKolmogorov(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
ARRAY[11, 3, 17], 12);
```

seq	edge	start_vid	end_vid	flow	residual_capacity
1	7	3	7	50	0
2	10	7	8	100	30
3	12	8	12	100	0
4	8	11	7	50	80
5	11	11	12	130	0

(5 rows)

Many to Many

pgr\_boykovKolmogorov([Edges SQL](#), start\_vids, end\_vids)  
Returns set of (seq, edge, start\_vid, end\_vid, flow, residual\_capacity)  
OR EMPTY SET

Example:

From vertices \{11, 3, 17\} to vertices \{5, 10, 12\}

```
SELECT * FROM pgr_boykovKolmogorov(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
ARRAY[11, 3, 17], ARRAY[5, 10, 12]);
```

seq	edge	start_vid	end_vid	flow	residual_capacity
1	7	3	7	50	0
2	1	6	5	50	80
3	4	7	6	50	0
4	10	7	8	100	30
5	12	8	12	100	0
6	8	11	7	100	30
7	11	11	12	130	0
8	9	11	16	80	50
9	3	15	10	80	50
10	16	16	15	80	0

(10 rows)

Combinations

pgr\_boykovKolmogorov([Edges SQL](#), [Combinations SQL](#))  
Returns set of (seq, edge, start\_vid, end\_vid, flow, residual\_capacity)  
OR EMPTY SET

Example:

Using a combinations table, equivalent to calculating result from vertices \{5, 6\} to vertices \{10, 15, 14\}.

The combinations table:

```
SELECT source, target FROM combinations
WHERE target NOT IN (5, 6);
```

source	target
5	10
6	15
6	14

(3 rows)

The query:

```
SELECT * FROM pgr_boykovKolmogorov(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
'SELECT * FROM combinations WHERE target NOT IN (5, 6)');
```

seq	edge	start_vid	end_vid	flow	residual_capacity
1	4	6	7	80	20
2	8	7	11	80	20
3	9	11	16	80	50
4	16	16	15	80	0

(4 rows)

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
<b>start vid</b>	BIGINT	Identifier of the starting vertex of the path.
<b>start vids</b>	ARRAY[BIGINT]	Array of identifiers of starting vertices.
<b>end vid</b>	BIGINT	Identifier of the ending vertex of the path.
<b>end vids</b>	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Inner Queries<sup>1</sup>

Edges SQL<sup>1</sup>

Column	Type	Default	Description
id	<b>ANY-INTEGER</b>		Identifier of the edge.
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
target	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
capacity	<b>ANY-INTEGER</b>		Weight of the edge (source, target)
reverse_capacity	<b>ANY-INTEGER</b>	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL<sup>1</sup>

Parameter	Type	Description
source	<b>ANY-INTEGER</b>	Identifier of the departure vertex.
target	<b>ANY-INTEGER</b>	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns<sup>1</sup>

Column	Type	Description
<b>seq</b>	INT	Sequential value starting from 1.
<b>edge</b>	BIGINT	Identifier of the edge in the original query (edges_sql).
<b>start_vid</b>	BIGINT	Identifier of the first end point vertex of the edge.
<b>end_vid</b>	BIGINT	Identifier of the second end point vertex of the edge.
<b>flow</b>	BIGINT	Flow through the edge in the direction (start_vid, end_vid).
<b>residual_capacity</b>	BIGINT	Residual capacity of the edge in the direction (start_vid, end_vid).

Additional Examples<sup>1</sup>

Example:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_boykovKolmogorov(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
```

```
'SELECT * FROM (VALUES (5, 10), (6, 15), (6, 14)) AS t(source, target));
seq | edge | start_vid | end_vid | flow | residual_capacity
-----+-----+-----+-----+-----+-----
1 | 4 | 6 | 7 | 80 | 20
2 | 8 | 7 | 11 | 80 | 20
3 | 9 | 11 | 16 | 80 | 50
4 | 16 | 16 | 15 | 80 | 0
(4 rows)
```

See Also

- [Flow - Family of functions](#)
  - [pgr\\_edmondsKarp](#)
  - [pgr\\_pushRelabel](#)
- [Boost: Boykov Kolmogorov max flow](#)

Indices and tables

- [Index](#)
- [Search Page](#)

**pgr\_edmondsKarp**

**pgr\_edmondsKarp** — Calculates the flow on the graph edges that maximizes the flow from the sources to the targets using Edmonds Karp Algorithm.

Availability

- Version 3.2.0
  - New proposed signature:
    - `pgr_edmondsKarp(Combinations)`
- Version 3.0.0
  - Function promoted to official.
- Version 2.5.0
  - Renamed from `pgr_maxFlowEdmondsKarp`
  - Function promoted to proposed.
- Version 2.3.0
  - New experimental function.

**Description**

The main characteristics are:

- The graph is **directed**.
- Process is done only on edges with positive capacities.
- When the maximum flow is 0 then there is no flow and **EMPTY SET** is returned.
  - There is no flow when source has the same value as target.
- Any duplicated values in source or target are ignored.
- Calculates the flow/residual capacity for each edge. In the output
  - Edges with zero flow are omitted.
- Creates
  - a **super source** and edges from it to all the sources,
  - a **super target** and edges from it to all the targetss.
- The maximum flow through the graph is guaranteed to be the value returned by [pgr\\_maxFlow](#) when executed with the same parameters and can be calculated:
  - By aggregation of the outgoing flow from the sources
  - By aggregation of the incoming flow to the targets
- Running time:  $\mathcal{O}(V * E^2)$

Boost Graph Inside

**Signatures**

Summary

```
pgr_edmondsKarp(Edges SQL, start vid, end vid)
pgr_edmondsKarp(Edges SQL, start vid, end vids)
pgr_edmondsKarp(Edges SQL, start vids, end vid)
pgr_edmondsKarp(Edges SQL, start vids, end vids)
pgr_edmondsKarp(Edges SQL, Combinations SQL)
Returns set of (seq, edge, start_vid, end_vid, flow, residual_capacity)
OR EMPTY SET
```

**One to One**

```
pgr_edmondsKarp(Edges SQL, start vid, end vid)
Returns set of (seq, edge, start_vid, end_vid, flow, residual_capacity)
OR EMPTY SET
```

Example:

From vertex  $\backslash(11\backslash)$  to vertex  $\backslash(12\backslash)$

```
SELECT * FROM pgr_edmondsKarp(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
11, 12);
seq | edge | start_vid | end_vid | flow | residual_capacity
-----+-----+-----+-----+-----+-----
1 | 10 | 7 | 8 | 100 | 30
```

2	12	8	12	100	0
3	8	11	7	100	30
4	11	11	12	130	0

(4 rows)

One to Many

`pgr_edmondsKarp(Edges SQL, start vid, end vids)`  
Returns set of (seq, edge, start\_vid, end\_vid, flow, residual\_capacity)  
OR EMPTY SET

Example:

From vertex \11\ to vertices \5, 10, 12\

```
SELECT * FROM pgr_edmondsKarp(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
11, ARRAY[5, 10, 12]);
seq | edge | start_vid | end_vid | flow | residual_capacity
```

1	1	6	5	50	80
2	4	7	6	50	0
3	10	7	8	80	50
4	12	8	12	80	20
5	8	11	7	130	0
6	11	11	12	130	0
7	9	11	16	80	50
8	3	15	10	80	50
9	16	16	15	80	0

(9 rows)

Many to One

`pgr_edmondsKarp(Edges SQL, start vids, end vid)`  
Returns set of (seq, edge, start\_vid, end\_vid, flow, residual\_capacity)  
OR EMPTY SET

Example:

From vertices \11, 3, 17\ to vertex \12\

```
SELECT * FROM pgr_edmondsKarp(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
ARRAY[11, 3, 17], 12);
seq | edge | start_vid | end_vid | flow | residual_capacity
```

1	7	3	7	50	0
2	10	7	8	100	30
3	12	8	12	100	0
4	8	11	7	50	80
5	11	11	12	130	0

(5 rows)

Many to Many

`pgr_edmondsKarp(Edges SQL, start vids, end vids)`  
Returns set of (seq, edge, start\_vid, end\_vid, flow, residual\_capacity)  
OR EMPTY SET

Example:

From vertices \11, 3, 17\ to vertices \5, 10, 12\

```
SELECT * FROM pgr_edmondsKarp(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
ARRAY[11, 3, 17], ARRAY[5, 10, 12]);
seq | edge | start_vid | end_vid | flow | residual_capacity
```

1	7	3	7	50	0
2	1	6	5	50	80
3	4	7	6	50	0
4	10	7	8	100	30
5	12	8	12	100	0
6	8	11	7	100	30
7	11	11	12	130	0
8	9	11	16	80	50
9	3	15	10	80	50
10	16	16	15	80	0

(10 rows)

Combinations

`pgr_edmondsKarp(Edges SQL, Combinations SQL)`  
Returns set of (seq, edge, start\_vid, end\_vid, flow, residual\_capacity)  
OR EMPTY SET

Example:

Using a combinations table, equivalent to calculating result from vertices \5, 6\ to vertices \10, 15, 14\.

The combinations table:

```
SELECT source, target FROM combinations
WHERE target NOT IN (5, 6);
source | target
```

5	10
6	15
6	14

(3 rows)

The query:

```
SELECT * FROM pgr_edmondsKarp(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
'SELECT * FROM combinations WHERE target NOT IN (5, 6)');
seq | edge | start_vid | end_vid | flow | residual_capacity
```

1	4	6	7	80	20
2	8	7	11	80	20
3	9	11	16	80	50
4	16	16	15	80	0

(4 rows)

Parameters¶

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Inner Queries¶

Edges SQL¶

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
capacity	ANY-INTEGER		Weight of the edge (source, target)
reverse_capacity	ANY-INTEGER	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL¶

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns¶

Column	Type	Description
seq	INT	Sequential value starting from 1.
edge	BIGINT	Identifier of the edge in the original query (edges_sql).
start_vid	BIGINT	Identifier of the first end point vertex of the edge.
end_vid	BIGINT	Identifier of the second end point vertex of the edge.
flow	BIGINT	Flow through the edge in the direction (start_vid, end_vid).
residual_capacity	BIGINT	Residual capacity of the edge in the direction (start_vid, end_vid).

Additional Examples¶

Example:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_edmondsKarp(
  'SELECT id, source, target, capacity, reverse_capacity
  FROM edges',
  'SELECT * FROM (VALUES (5, 10), (6, 15), (6, 14)) AS t(source, target)');
seq | edge | start_vid | end_vid | flow | residual_capacity
-----+-----+-----+-----+-----+-----
1 | 4 | 6 | 7 | 80 | 20
2 | 8 | 7 | 11 | 80 | 20
3 | 9 | 11 | 16 | 80 | 50
4 | 16 | 16 | 15 | 80 | 0
(4 rows)
```

See Also

- [Flow - Family of functions](#)
  - [pgr\\_boykovKolmogorov](#)
  - [pgr\\_pushRelabel](#)
- [Boost: Edmonds Karp max flow](#)
- [https://en.wikipedia.org/wiki/Edmonds%E2%80%93Karp\\_algorithm](https://en.wikipedia.org/wiki/Edmonds%E2%80%93Karp_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_pushRelabel

pgr\_pushRelabel — Calculates the flow on the graph edges that maximizes the flow from the sources to the targets using Push Relabel Algorithm.

Availability

- Version 3.2.0
  - New proposed signature:
    - `pgr_pushRelabel(Combinations)`
- Version 3.0.0
  - Function promoted to official.
- Version 2.5.0
  - Renamed from `pgr_maxFlowPushRelabel`
  - Function promoted to proposed.
- Version 2.3.0
  - New experimental function.

Description

The main characteristics are:

- The graph is **directed**.
- Process is done only on edges with positive capacities.
- When the maximum flow is 0 then there is no flow and **EMPTY SET** is returned.
  - There is no flow when source has the same value as target.
- Any duplicated values in source or target are ignored.
- Calculates the flow/residual capacity for each edge. In the output
  - Edges with zero flow are omitted.
- Creates
  - a **super source** and edges from it to all the sources,
  - a **super target** and edges from it to all the targetss.
- The maximum flow through the graph is guaranteed to be the value returned by [pgr\\_maxFlow](#) when executed with the same parameters and can be calculated:
  - By aggregation of the outgoing flow from the sources
  - By aggregation of the incoming flow to the targets
- Running time:  $\mathcal{O}(V^3)$

 Boost Graph Inside

Signatures

Summary

```
pgr_pushRelabel(Edges SQL, start_vid, end_vid)
pgr_pushRelabel(Edges SQL, start_vid, end_vids)
pgr_pushRelabel(Edges SQL, start_vids, end_vid)
pgr_pushRelabel(Edges SQL, start_vids, end_vids)
pgr_pushRelabel(Edges SQL, Combinations SQL)
Returns set of (seq, edge, start_vid, end_vid, flow, residual_capacity)
OR EMPTY SET
```

One to One

```
pgr_pushRelabel(Edges SQL, start_vid, end_vid)
Returns set of (seq, edge, start_vid, end_vid, flow, residual_capacity)
OR EMPTY SET
```

Example:

From vertex  $\backslash(11\backslash)$  to vertex  $\backslash(12\backslash)$

```
SELECT * FROM pgr_pushRelabel(
  'SELECT id, source, target, capacity, reverse_capacity
```



```
FROM edges',
11, 12);
seq | edge | start_vid | end_vid | flow | residual_capacity
-----+-----+-----+-----+-----+-----
1 | 10 | 7 | 8 | 100 | 30
2 | 12 | 8 | 12 | 100 | 0
3 | 8 | 11 | 7 | 100 | 30
4 | 11 | 11 | 12 | 130 | 0
(4 rows)
```

One to Many

pgr\_pushRelabel([Edges SQL](#), start\_vid, end\_vids)  
Returns set of (seq, edge, start\_vid, end\_vid, flow, residual\_capacity)  
OR EMPTY SET

Example:  
From vertex (11) to vertices (5, 10, 12)

```
SELECT * FROM pgr_pushRelabel(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
11, ARRAY[5, 10, 12]);
seq | edge | start_vid | end_vid | flow | residual_capacity
-----+-----+-----+-----+-----+-----
1 | 6 | 1 | 3 | 50 | 0
2 | 6 | 3 | 1 | 50 | 50
3 | 7 | 3 | 7 | 50 | 0
4 | 1 | 6 | 5 | 30 | 100
5 | 7 | 7 | 3 | 50 | 80
6 | 4 | 7 | 6 | 30 | 20
7 | 10 | 7 | 8 | 100 | 30
8 | 12 | 8 | 12 | 100 | 0
9 | 8 | 11 | 7 | 130 | 0
10 | 11 | 11 | 12 | 130 | 0
11 | 9 | 11 | 16 | 80 | 50
12 | 3 | 15 | 10 | 80 | 50
13 | 16 | 16 | 15 | 80 | 0
(13 rows)
```

Many to One

pgr\_pushRelabel([Edges SQL](#), start\_vids, end\_vid)  
Returns set of (seq, edge, start\_vid, end\_vid, flow, residual\_capacity)  
OR EMPTY SET

Example:  
From vertices (11, 3, 17) to vertex (12)

```
SELECT * FROM pgr_pushRelabel(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
ARRAY[11, 3, 17], 12);
seq | edge | start_vid | end_vid | flow | residual_capacity
-----+-----+-----+-----+-----+-----
1 | 10 | 7 | 8 | 100 | 30
2 | 12 | 8 | 12 | 100 | 0
3 | 8 | 11 | 7 | 100 | 30
4 | 11 | 11 | 12 | 130 | 0
(4 rows)
```

Many to Many

pgr\_pushRelabel([Edges SQL](#), start\_vids, end\_vids)  
Returns set of (seq, edge, start\_vid, end\_vid, flow, residual\_capacity)  
OR EMPTY SET

Example:  
From vertices (11, 3, 17) to vertices (5, 10, 12)

```
SELECT * FROM pgr_pushRelabel(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
ARRAY[11, 3, 17], ARRAY[5, 10, 12]);
seq | edge | start_vid | end_vid | flow | residual_capacity
-----+-----+-----+-----+-----+-----
1 | 7 | 3 | 7 | 20 | 30
2 | 1 | 6 | 5 | 50 | 80
3 | 4 | 7 | 6 | 50 | 0
4 | 10 | 7 | 8 | 100 | 30
5 | 12 | 8 | 12 | 100 | 0
6 | 8 | 11 | 7 | 130 | 0
7 | 11 | 11 | 12 | 130 | 0
8 | 9 | 11 | 16 | 80 | 50
9 | 3 | 15 | 10 | 80 | 50
10 | 16 | 16 | 15 | 80 | 0
(10 rows)
```

Combinations

pgr\_pushRelabel([Edges SQL](#), [Combinations SQL](#))  
Returns set of (seq, edge, start\_vid, end\_vid, flow, residual\_capacity)  
OR EMPTY SET

Example:  
Using a combinations table, equivalent to calculating result from vertices (5, 6) to vertices (10, 15, 14).

The combinations table:

```
SELECT source, target FROM combinations
WHERE target NOT IN (5, 6);
source | target
-----+-----
5 | 10
6 | 15
6 | 14
(3 rows)
```

The query:

```
SELECT * FROM pgr_pushRelabel(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
'SELECT * FROM combinations WHERE target NOT IN (5, 6)');
```

seq	edge	start_vid	end_vid	flow	residual_capacity
1	4	6	7	80	20
2	8	7	11	80	20
3	11	11	12	50	80
4	9	11	16	30	100
5	13	12	17	50	50
6	16	16	15	80	0
7	15	17	16	50	0
(7 rows)					

Parameters1

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start_vid	BIGINT	Identifier of the starting vertex of the path.
start_vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end_vid	BIGINT	Identifier of the ending vertex of the path.
end_vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Inner Queries1

Edges SQL1

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
capacity	ANY-INTEGER		Weight of the edge (source, target)
reverse_capacity	ANY-INTEGER	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL1

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns1

Column	Type	Description
seq	INT	Sequential value starting from 1.
edge	BIGINT	Identifier of the edge in the original query (edges_sql).
start_vid	BIGINT	Identifier of the first end point vertex of the edge.
end_vid	BIGINT	Identifier of the second end point vertex of the edge.
flow	BIGINT	Flow through the edge in the direction (start_vid, end_vid).

Column	Type	Description
<b>residual_capacity</b>	BIGINT	Residual capacity of the edge in the direction (start_vid, end_vid).

Additional Examples

Example:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_pushRelabel(
'SELECT id, source, target, capacity, reverse_capacity
FROM edges',
'SELECT * FROM (VALUES (5, 10), (6, 15), (6, 14)) AS t(source, target)');
seq | edge | start_vid | end_vid | flow | residual_capacity
-----+-----+-----+-----+-----+-----
1 | 4 | 6 | 7 | 80 | 20
2 | 8 | 7 | 11 | 80 | 20
3 | 11 | 11 | 12 | 50 | 80
4 | 9 | 11 | 16 | 30 | 100
5 | 13 | 12 | 17 | 50 | 50
6 | 16 | 16 | 15 | 80 | 0
7 | 15 | 17 | 16 | 50 | 0
(7 rows)
```

See Also

- Flow - Family of functions
  - [pgr\\_boykovKolmogorov](#)
  - [pgr\\_edmondsKarp](#)
- Boost: push relabel max flow
- [https://en.wikipedia.org/wiki/Push%E2%80%93relabel\\_maximum\\_flow\\_algorithm](https://en.wikipedia.org/wiki/Push%E2%80%93relabel_maximum_flow_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_edgeDisjointPaths

pgr\_edgeDisjointPaths — Calculates edge disjoint paths between two groups of vertices.

Availability

- Version 3.2.0
  - New proposed signature:
    - pgr\_edgeDisjointPaths(Combinations)
- Version 3.0.0
  - Function promoted to official.
- Version 2.5.0
  - Function promoted to proposed.
- Version 2.3.0
  - New experimental function.

Description

Calculates the edge disjoint paths between two groups of vertices. Utilizes underlying maximum flow algorithms to calculate the paths.

The main characteristics are:

- Calculates the edge disjoint paths between any two groups of vertices.
- Returns EMPTY SET when source and destination are the same, or cannot be reached.
- The graph can be directed or undirected.
- Uses [pgr\\_boykovKolmogorov](#) to calculate the paths.

Boost Graph Inside

Signatures

Summary

```
pgr_edgeDisjointPaths(Edges SQL, start_vid, end_vid, [directed])
pgr_edgeDisjointPaths(Edges SQL, start_vid, end_vids, [directed])
pgr_edgeDisjointPaths(Edges SQL, start_vids, end_vid, [directed])
pgr_edgeDisjointPaths(Edges SQL, start_vids, end_vids, [directed])
pgr_edgeDisjointPaths(Edges SQL, Combinations SQL, [directed])
Returns set of (seq, path_id, path_seq, [start_vid,] [end_vid,] node, edge, cost, agg_cost)
OR EMPTY SET
```

One to One

```
pgr_edgeDisjointPaths(Edges SQL, start_vid, end_vid, [directed])
Returns set of (seq, path_id, path_seq, node, edge, cost, agg_cost)
OR EMPTY SET
```

Example:

From vertex \11\ to vertex \12\

```
SELECT * FROM pgr_edgeDisjointPaths(
'SELECT id, source, target, cost, reverse_cost
FROM edges',
11, 12);
```

seq	path_id	path_seq	node	edge	cost	agg_cost
1	1	1	11	8	1	0
2	1	2	7	10	1	1
3	1	3	8	12	1	2
4	1	4	12	-1	0	3
5	2	1	11	11	1	0
6	2	2	12	-1	0	1

(6 rows)

#### One to Many

`pgr_edgeDisjointPaths`([Edges SQL](#), **start vid**, **end vids**, [directed])  
Returns set of (seq, path\_id, path\_seq, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \({11}\) to vertices \({5, 10, 12}\)

```
SELECT * FROM pgr_edgeDisjointPaths(
'SELECT id, source, target, cost, reverse_cost
FROM edges',
11, ARRAY[5, 10, 12]);
```

seq	path_id	path_seq	end_vid	node	edge	cost	agg_cost
1	1	1	5	11	8	1	0
2	1	2	5	7	4	1	1
3	1	3	5	6	1	1	2
4	1	4	5	5	-1	0	3
5	2	1	10	11	9	1	0
6	2	2	10	16	16	1	1
7	2	3	10	15	3	1	2
8	2	4	10	10	-1	0	3
9	3	1	12	11	8	1	0
10	3	2	12	7	10	1	1
11	3	3	12	8	12	1	2
12	3	4	12	12	-1	0	3
13	4	1	12	11	11	1	0
14	4	2	12	12	-1	0	1

(14 rows)

#### Many to One

`pgr_edgeDisjointPaths`([Edges SQL](#), **start vids**, **end vid**, [directed])  
Returns set of (seq, path\_id, path\_seq, start\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \({11, 3, 17}\) to vertex \({12}\)

```
SELECT * FROM pgr_edgeDisjointPaths(
'SELECT id, source, target, cost, reverse_cost
FROM edges',
ARRAY[11, 3, 17], 12);
```

seq	path_id	path_seq	start_vid	node	edge	cost	agg_cost
1	1	1	3	3	7	1	0
2	1	2	3	7	8	1	1
3	1	3	3	11	11	1	2
4	1	4	3	12	-1	0	3
5	2	1	11	11	8	1	0
6	2	2	11	7	10	1	1
7	2	3	11	8	12	1	2
8	2	4	11	12	-1	0	3
9	3	1	11	11	11	1	0
10	3	2	11	12	-1	0	1
11	4	1	17	17	15	1	0
12	4	2	17	16	9	1	1
13	4	3	17	11	11	1	2
14	4	4	17	12	-1	0	3

(14 rows)

#### Many to Many

`pgr_edgeDisjointPaths`([Edges SQL](#), **start vids**, **end vids**, [directed])  
Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \({11, 3, 17}\) to vertices \({5, 10, 12}\)

```
SELECT * FROM pgr_edgeDisjointPaths(
'SELECT id, source, target, cost, reverse_cost
FROM edges',
ARRAY[11, 3, 17], ARRAY[5, 10, 12]);
```

seq	path_id	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	1	3	5	3	7	1	0
2	1	2	3	5	7	4	1	1
3	1	3	3	5	6	1	1	2
4	1	4	3	5	5	-1	0	3
5	2	1	3	10	3	7	1	0
6	2	2	3	10	7	8	1	1
7	2	3	3	10	11	9	1	2
8	2	4	3	10	16	16	1	3
9	2	5	3	10	15	3	1	4
10	2	6	3	10	10	-1	0	5
11	3	1	3	12	3	7	1	0
12	3	2	3	12	7	8	1	1
13	3	3	3	12	11	11	1	2
14	3	4	3	12	12	-1	0	3
15	4	1	11	5	11	8	1	0
16	4	2	11	5	7	4	1	1
17	4	3	11	5	6	1	1	2
18	4	4	11	5	5	-1	0	3
19	5	1	11	10	11	9	1	0
20	5	2	11	10	16	16	1	1
21	5	3	11	10	15	3	1	2
22	5	4	11	10	10	-1	0	3
23	6	1	11	12	11	8	1	0
24	6	2	11	12	7	10	1	1
25	6	3	11	12	8	12	1	2
26	6	4	11	12	12	-1	0	3
27	7	1	11	12	11	11	1	0
28	7	2	11	12	12	-1	0	1
29	8	1	17	5	17	15	1	0
30	8	2	17	5	16	16	1	1
31	8	3	17	5	15	3	1	2
32	8	4	17	5	10	2	1	3

33	8	5	17	5	6	1	1	4
34	8	6	17	5	5	-1	0	5
35	9	1	17	10	17	15	1	0
36	9	2	17	10	16	16	1	1
37	9	3	17	10	15	3	1	2
38	9	4	17	10	10	-1	0	3
39	10	1	17	12	17	15	1	0
40	10	2	17	12	16	9	1	1
41	10	3	17	12	11	11	1	2
42	10	4	17	12	12	-1	0	3

(42 rows)

Combinations1

pgr\_edgeDisjointPaths([Edges SQL](#), [Combinations SQL](#), [directed])  
Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table, equivalent to calculating result from vertices \(\{5, 6\}\) to vertices \(\{10, 15, 14\}\) on an undirected graph.

The combinations table:

```
SELECT source, target FROM combinations
WHERE target NOT IN (5, 6);
source | target
-----+-----
5 | 10
6 | 15
6 | 14
(3 rows)
```

The query:

```
SELECT * FROM pgr_edgeDisjointPaths(
'SELECT id, source, target, cost, reverse_cost
FROM edges',
'SELECT * FROM combinations WHERE target NOT IN (5, 6)',
directed => false);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 5 | 10 | 5 | 10 | 1 | 0
2 | 1 | 2 | 5 | 10 | 6 | 2 | -1 | 1
3 | 1 | 3 | 5 | 10 | 10 | -1 | 0 | 0
4 | 2 | 1 | 6 | 15 | 6 | 4 | 1 | 0
5 | 2 | 2 | 6 | 15 | 7 | 8 | 1 | 1
6 | 2 | 3 | 6 | 15 | 11 | 9 | 1 | 2
7 | 2 | 4 | 6 | 15 | 16 | 16 | 1 | 3
8 | 2 | 5 | 6 | 15 | 15 | -1 | 0 | 4
9 | 3 | 1 | 6 | 15 | 6 | 2 | -1 | 0
10 | 3 | 2 | 6 | 15 | 10 | 3 | -1 | -1
11 | 3 | 3 | 6 | 15 | 15 | -1 | 0 | -2
(11 rows)
```

Parameters1

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters1

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries1

Edges SQL1

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)

Column	Type	Default	Description
			Weight of the edge (target, source)
reverse_cost	ANY-NUMERICAL	-1	<ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations [SQL](#)

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Set of (seq, path\_id, path\_seq [, start\_vid] [, end\_vid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Path identifier. <ul style="list-style-type: none"> <li>Has value <b>1</b> for the first of a path fromstart_vid to end_vid.</li> </ul>
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex. Returned when multiple starting vetrices are in the query. <ul style="list-style-type: none"> <li><a href="#">Many to One</a></li> <li><a href="#">Many to Many</a></li> <li><a href="#">Combinations</a></li> </ul>
end_vid	BIGINT	Identifier of the ending vertex. Returned when multiple ending vertices are in the query. <ul style="list-style-type: none"> <li><a href="#">One to Many</a></li> <li><a href="#">Many to Many</a></li> <li><a href="#">Combinations</a></li> </ul>
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Additional Examples

Example:

Manually assigned vertex combinations on an undirected graph.

```
SELECT * FROM pgr_edgeDisjointPaths(
'SELECT id, source, target, cost, reverse_cost
FROM edges',
'SELECT * FROM (VALUES (5, 10), (6, 15), (6, 14)) AS t(source, target)',
directed => false);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 5 | 10 | 5 | 1 | 1 | 0
2 | 1 | 2 | 5 | 10 | 6 | 2 | -1 | 1
3 | 1 | 3 | 5 | 10 | 10 | -1 | 0 | 0
4 | 2 | 1 | 6 | 15 | 6 | 4 | 1 | 0
5 | 2 | 2 | 6 | 15 | 7 | 8 | 1 | 1
6 | 2 | 3 | 6 | 15 | 11 | 9 | 1 | 2
7 | 2 | 4 | 6 | 15 | 16 | 16 | 1 | 3
8 | 2 | 5 | 6 | 15 | 15 | -1 | 0 | 4
9 | 3 | 1 | 6 | 15 | 6 | 2 | -1 | 0
10 | 3 | 2 | 6 | 15 | 10 | 3 | -1 | -1
11 | 3 | 3 | 6 | 15 | 15 | -1 | 0 | -2
(11 rows)
```

See Also

- [Flow - Family of functions](#)
- [Boost: Boykov Kolmogorov max flow](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_maxCardinalityMatch

pgr\_maxCardinalityMatch — Calculates a maximum cardinality matching in a graph.

Availability

- Version 3.4.0
  - Use cost and reverse\_cost on the inner query
  - Results are ordered
  - Works for undirected graphs.
  - New signature
    - pgr\_maxCardinalityMatch(text) returns only edge column.
  - Deprecated signature
    - pgr\_maxCardinalityMatch(text,boolean)
      - directed => false when used.
- Version 3.0.0
  - Function promoted to official.
- Version 2.5.0
  - Renamed from pgr\_maximumCardinalityMatching
  - Function promoted to proposed.
- Version 2.3.0
  - New experimental function.

Description

The main characteristics are:

- Works for **undirected** graphs.
- A matching or independent edge set in a graph is a set of edges without common vertices.
- A maximum matching is a matching that contains the largest possible number of edges.
  - There may be many maximum matchings.
  - Calculates one possible maximum cardinality matching in a graph.
- Running time:  $\mathcal{O}(E \cdot V \cdot \alpha(E, V))$ 
  - $\alpha(E, V)$  is the inverse of the [Ackermann function](#).

 Boost Graph Inside

Signatures

pgr\_maxCardinalityMatch([Edges SQL](#))

Returns set of (edge)  
OR EMPTY SET

Example:

Using all edges.

```
SELECT * FROM pgr_maxCardinalityMatch(
'SELECT id, source, target, cost, reverse_cost FROM edges');
edge
-----
 1
 5
 6
13
14
16
17
18
(8 rows)
```

Parameters

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below.

Inner Queries

Edges SQL

SQL query, which should return a set of rows with the following columns:

Column	Type	Default	Description
id	<b>ANY-INTEGER</b>		Identifier of the edge.
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
target	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
cost	<b>ANY-NUMERICAL</b>		A positive value represents the existence of the edge (source, target).
reverse_cost	<b>ANY-NUMERICAL</b>	-1	A positive value represents the existence of the edge (target, source)

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns<sup>¶</sup>

Column	Type	Description
edge	BIGINT	Identifier of the edge in the original query.

See Also<sup>¶</sup>

- [Flow - Family of functions](#)
- [Migration guide](#)
- [Boost: maximum\\_matching](#)
- [https://en.wikipedia.org/wiki/Matching\\_%28graph\\_theory%29](https://en.wikipedia.org/wiki/Matching_%28graph_theory%29)
- [https://en.wikipedia.org/wiki/Ackermann\\_function](https://en.wikipedia.org/wiki/Ackermann_function)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_maxFlowMinCost - Experimental<sup>¶</sup>

pgr\_maxFlowMinCost — Calculates the edges that minimizes the total cost of the maximum flow on a graph

Availability

- Version 3.2.0
  - New experimental signature:
    - pgr\_maxFlowMinCost(Combinations)
- Version 3.0.0
  - New experimental function.

Description<sup>¶</sup>

graph inside.

The main characteristics are:

- The graph is **directed**.
- Process is done only on edges with positive capacities.
- When the maximum flow is 0 then there is no flow and **EMPTY SET** is returned.
  - There is no flow when source has the same value as target.
- Any duplicated values in source or target are ignored.
- Calculates the flow/residual capacity for each edge. In the output
  - Edges with zero flow are omitted.
- Creates
  - a **super source** and edges from it to all the sources,
  - a **super target** and edges from it to all the targetss.
- The maximum flow through the graph is guaranteed to be the value returned by [pgr\\_maxFlow](#) when executed with the same parameters and can be calculated:
  - By aggregation of the outgoing flow from the sources
  - By aggregation of the incoming flow to the targets
- **TODO** check which statement is true:
  - The cost value of all input edges must be nonnegative.
  - Process is done when the cost value of all input edges is nonnegative.



- Process is done on edges with nonnegative cost.
- Running time:  $\mathcal{O}(U * (E + V * \log V))$ 
  - where  $\mathcal{O}(U)$  is the value of the max flow.
  - $\mathcal{O}(U)$  is upper bound on number of iterations. In many real world cases number of iterations is much smaller than  $\mathcal{O}(U)$ .

❑ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Boost Graph Inside

Signatures

Summary

`pgr_maxFlowMinCost(Edges SQL, start vid, end vid)`  
`pgr_maxFlowMinCost(Edges SQL, start vid, end vids)`  
`pgr_maxFlowMinCost(Edges SQL, start vids, end vid)`  
`pgr_maxFlowMinCost(Edges SQL, start vids, end vids)`  
`pgr_maxFlowMinCost(Edges SQL, Combinations SQL)`  
Returns set of (seq, edge, source, target, flow, residual\_capacity, cost, agg\_cost)  
OR EMPTY SET

One to One

`pgr_maxFlowMinCost(Edges SQL, start vid, end vid)`  
Returns set of (seq, edge, source, target, flow, residual\_capacity, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex  $\backslash(11\backslash)$  to vertex  $\backslash(12\backslash)$

```
SELECT * FROM pgr_maxFlowMinCost(
'SELECT id, source, target, capacity, reverse_capacity, cost, reverse_cost
FROM edges',
11, 12);
seq | edge | source | target | flow | residual_capacity | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
 1 | 10 |    7 |    8 | 100 |          30 | 100 |         100
 2 | 12 |    8 |   12 |   0 |          100 | 200 |         200
 3 |  8 |   11 |    7 | 100 |          30 | 100 |         300
 4 | 11 |   11 |   12 | 130 |           0 | 130 |         430
(4 rows)
```

One to Many

`pgr_maxFlowMinCost(Edges SQL, start vid, end vids)`  
Returns set of (seq, edge, source, target, flow, residual\_capacity, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex  $\backslash(11\backslash)$  to vertices  $\backslash(\backslash5, 10, 12\backslash\backslash)$

```
SELECT * FROM pgr_maxFlowMinCost(
'SELECT id, source, target, capacity, reverse_capacity, cost, reverse_cost
FROM edges',
11, ARRAY[5, 10, 12]);
seq | edge | source | target | flow | residual_capacity | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
 1 |  1 |    6 |    5 | 30 |          100 | 30 |          30
 2 |  4 |    7 |    6 | 30 |           20 | 30 |          60
 3 | 10 |    7 |    8 | 100 |          30 | 100 |         160
 4 | 12 |    8 |   12 |   0 |          100 | 260 |         260
 5 |  8 |   11 |    7 | 130 |           0 | 130 |         390
 6 | 11 |   11 |   12 | 130 |           0 | 130 |         520
 7 |  9 |   11 |   16 | 80 |           50 | 80 |         600
 8 |  3 |   15 |   10 | 80 |           50 | 80 |         680
 9 | 16 |   16 |   15 | 80 |           0 | 80 |         760
(9 rows)
```

Many to One

`pgr_maxFlowMinCost(Edges SQL, start vids, end vid)`  
Returns set of (seq, edge, source, target, flow, residual\_capacity, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \(\{11, 3, 17\}\) to vertex \(\{12\}\)

```
SELECT * FROM pgr_maxFlowMinCost(
'SELECT id, source, target, capacity, reverse_capacity, cost, reverse_cost
FROM edges',
ARRAY[11, 3, 17], 12);
seq | edge | source | target | flow | residual_capacity | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 7 | 3 | 7 | 50 | 0 | 50 | 50
2 | 10 | 7 | 8 | 100 | 30 | 100 | 150
3 | 12 | 8 | 12 | 100 | 0 | 100 | 250
4 | 8 | 11 | 7 | 50 | 80 | 50 | 300
5 | 11 | 11 | 12 | 130 | 0 | 130 | 430
(5 rows)
```

Many to Many

pgr\_maxFlowMinCost([Edges SQL](#), start vids, end vids)  
Returns set of (seq, edge, source, target, flow, residual\_capacity, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \(\{11, 3, 17\}\) to vertices \(\{5, 10, 12\}\)

```
SELECT * FROM pgr_maxFlowMinCost(
'SELECT id, source, target, capacity, reverse_capacity, cost, reverse_cost
FROM edges',
ARRAY[11, 3, 17], ARRAY[5, 10, 12]);
seq | edge | source | target | flow | residual_capacity | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 7 | 3 | 7 | 50 | 0 | 50 | 50
2 | 1 | 6 | 5 | 50 | 80 | 50 | 100
3 | 4 | 7 | 6 | 50 | 0 | 50 | 150
4 | 10 | 7 | 8 | 100 | 30 | 100 | 250
5 | 12 | 8 | 12 | 100 | 0 | 100 | 350
6 | 8 | 11 | 7 | 100 | 30 | 100 | 450
7 | 11 | 11 | 12 | 130 | 0 | 130 | 580
8 | 9 | 11 | 16 | 30 | 100 | 30 | 610
9 | 3 | 15 | 10 | 80 | 50 | 80 | 690
10 | 16 | 16 | 15 | 80 | 0 | 80 | 770
11 | 15 | 17 | 16 | 50 | 0 | 50 | 820
(11 rows)
```

Combinations

pgr\_maxFlowMinCost([Edges SQL](#), [Combinations SQL](#))  
Returns set of (seq, edge, source, target, flow, residual\_capacity, cost, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table, equivalent to calculating result from vertices \(\{5, 6\}\) to vertices \(\{10, 15, 14\}\).

The combinations table:

```
SELECT source, target FROM combinations
WHERE target NOT IN (5, 6);
source | target
-----+-----
5 | 10
6 | 15
6 | 14
(3 rows)
```

The query:

```
SELECT * FROM pgr_maxFlowMinCost(
'SELECT id, source, target, capacity, reverse_capacity, cost, reverse_cost
FROM edges',
'SELECT * FROM combinations WHERE target NOT IN (5, 6)');
seq | edge | source | target | flow | residual_capacity | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 4 | 6 | 7 | 80 | 20 | 80 | 80
2 | 8 | 7 | 11 | 80 | 20 | 80 | 160
3 | 9 | 11 | 16 | 80 | 50 | 80 | 240
4 | 16 | 16 | 15 | 80 | 0 | 80 | 320
(4 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.

Column	Type	Default	Description
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
capacity	ANY-INTEGER		Capacity of the edge (source, target) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>
reverse_capacity	ANY-INTEGER	-1	Capacity of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>
cost	ANY-NUMERICAL		Weight of the edge (source, target) if it exist
reverse_cost	ANY-NUMERICAL	\(-1\)	Weight of the edge (target, source) if it exist

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations [SQL](#)

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Column	Type	Description
seq	INT	Sequential value starting from 1.
edge	BIGINT	Identifier of the edge in the original query (edges_sql).
source	BIGINT	Identifier of the first end point vertex of the edge.
target	BIGINT	Identifier of the second end point vertex of the edge.
flow	BIGINT	Flow through the edge in the direction (source, target).
residual_capacity	BIGINT	Residual capacity of the edge in the direction (source, target).
cost	FLOAT	The cost of sending this flow through the edge in the direction (source, target).
agg_cost	FLOAT	The aggregate cost.

Additional Examples

Example:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_maxFlowMinCost(
'SELECT id, source, target, capacity, reverse_capacity, cost, reverse_cost
FROM edges',
'SELECT * FROM (VALUES (5, 10), (6, 15), (6, 14)) AS t(source, target)');
seq | edge | source | target | flow | residual_capacity | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 4 | 6 | 7 | 80 | 20 | 80 | 80
2 | 8 | 7 | 11 | 80 | 20 | 80 | 160
3 | 9 | 11 | 16 | 80 | 50 | 80 | 240
4 | 16 | 16 | 15 | 80 | 0 | 80 | 320
(4 rows)
```

See Also

- [Flow - Family of functions](#)
- [Boost: push relabel max flow](#)

Indices and tables

- [Index](#)
- [Search Page](#)

`pgr_maxFlowMinCost_Cost` - [Experimental](#)

`pgr_maxFlowMinCost_Cost` — Calculates the minimum total cost of the maximum flow on a graph

#### Availability

- Version 3.2.0
  - New experimental signature:
    - `pgr_maxFlowMinCost_Cost(Combinations)`
- Version 3.0.0
  - New experimental function.

#### Description

##### The main characteristics are:

- The graph is **directed**.
- Process is done only on edges with positive capacities.
- When the maximum flow is 0 then there is no flow and **EMPTY SET** is returned.
  - There is no flow when source has the same value as target.
- Any duplicated values in source or target are ignored.
- Calculates the flow/residual capacity for each edge. In the output
  - Edges with zero flow are omitted.
- Creates
  - a **super source** and edges from it to all the sources,
  - a **super target** and edges from it to all the targets.
- The maximum flow through the graph is guaranteed to be the value returned by `pgr_maxFlow` when executed with the same parameters and can be calculated:
  - By aggregation of the outgoing flow from the sources
  - By aggregation of the incoming flow to the targets

##### The main characteristics are:

- The graph is **directed**.
- **The cost value of all input edges must be nonnegative.**
- When the maximum flow is 0 then there is no flow and **0** is returned.
  - There is no flow when source has the same value as target.
- Any duplicated values in source or target are ignored.
- Uses [pgr\\_maxFlowMinCost - Experimental](#).
- Running time:  $\mathcal{O}(U * (E + V * \log V))$ 
  - where  $\mathcal{O}(U)$  is the value of the max flow.
  - $\mathcal{O}(U)$  is upper bound on number of iterations. In many real world cases number of iterations is much smaller than  $\mathcal{O}(U)$ .

☐ Experimental

#### Warning

##### Possible server crash

- These functions might create a server crash

#### Warning

##### Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

☐ Boost Graph Inside

#### Signatures

#### Summary

pgr\_maxFlowMinCost\_Cost([Edges SQL](#), start vid, end vid)  
pgr\_maxFlowMinCost\_Cost([Edges SQL](#), start vid, end vids)  
pgr\_maxFlowMinCost\_Cost([Edges SQL](#), start vids, end vid)  
pgr\_maxFlowMinCost\_Cost([Edges SQL](#), start vids, end vids)  
pgr\_maxFlowMinCost\_Cost([Edges SQL](#), [Combinations SQL](#))  
RETURNS FLOAT

One to One

pgr\_maxFlowMinCost\_Cost([Edges SQL](#), start vid, end vid)  
RETURNS FLOAT

Example:

From vertex \11\ to vertex \12\

```
SELECT * FROM pgr_maxFlowMinCost_Cost(
'SELECT id, source, target, capacity, reverse_capacity, cost, reverse_cost
FROM edges',
11, 12);
pgr_maxflowmincost_cost
-----
430
(1 row)
```

One to Many

pgr\_maxFlowMinCost\_Cost([Edges SQL](#), start vid, end vids)  
RETURNS FLOAT

Example:

From vertex \11\ to vertices \5, 10, 12\

```
SELECT * FROM pgr_maxFlowMinCost_Cost(
'SELECT id, source, target, capacity, reverse_capacity, cost, reverse_cost
FROM edges',
ARRAY[11, 3, 17], 12);
pgr_maxflowmincost_cost
-----
430
(1 row)
```

Many to One

pgr\_maxFlowMinCost\_Cost([Edges SQL](#), start vids, end vid)  
RETURNS FLOAT

Example:

From vertices \11, 3, 17\ to vertex \12\

```
SELECT * FROM pgr_maxFlowMinCost_Cost(
'SELECT id, source, target, capacity, reverse_capacity, cost, reverse_cost
FROM edges',
11, ARRAY[5, 10, 12]);
pgr_maxflowmincost_cost
-----
760
(1 row)
```

Many to Many

pgr\_maxFlowMinCost\_Cost([Edges SQL](#), start vids, end vids)  
RETURNS FLOAT

Example:

From vertices \11, 3, 17\ to vertices \5, 10, 12\

```
SELECT * FROM pgr_maxFlowMinCost_Cost(
'SELECT id, source, target, capacity, reverse_capacity, cost, reverse_cost
FROM edges',
ARRAY[11, 3, 17], ARRAY[5, 10, 12]);
pgr_maxflowmincost_cost
-----
820
(1 row)
```

Combinations

pgr\_maxFlowMinCost\_Cost([Edges SQL](#), [Combinations SQL](#))  
RETURNS FLOAT

Example:

Using a combinations table, equivalent to calculating result from vertices \5, 6\ to vertices \10, 15, 14\.

The combinations table:

```
SELECT source, target FROM combinations
WHERE target NOT IN (5, 6);
source | target
-----+-----
5 | 10
6 | 15
6 | 14
(3 rows)
```

The query:

```
SELECT * FROM pgr_maxFlowMinCost_Cost(
'SELECT id, source, target, capacity, reverse_capacity, cost, reverse_cost
FROM edges',
'SELECT * FROM combinations WHERE target NOT IN (5, 6)');
pgr_maxflowmincost_cost
-----
320
(1 row)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
capacity	ANY-INTEGER		Capacity of the edge (source, target) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>
reverse_capacity	ANY-INTEGER	-1	Capacity of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>
cost	ANY-NUMERICAL		Weight of the edge (source, target) if it exist
reverse_cost	ANY-NUMERICAL	\(-1\)	Weight of the edge (target, source) if it exist

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Return columns

Type	Description
FLOAT	Minimum Cost Maximum Flow possible from the source(s) to the target(s)

Additional Examples

Example:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_maxFlowMinCost_Cost(
'SELECT id, source, target, capacity, reverse_capacity, cost, reverse_cost
FROM edges',
'SELECT * FROM (VALUES (5, 10), (6, 15), (6, 14)) AS t(source, target)');
pgr_maxflowmincost_cost
-----
320
(1 row)
```

See Also

- [Flow - Family of functions](#)
- [Boost: push relabel max flow](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Flow Functions General Information

The main characteristics are:

- The graph is **directed**.
- Process is done only on edges with positive capacities.
- When the maximum flow is 0 then there is no flow and **EMPTY SET** is returned.
  - There is no flow when source has the same value as target.
- Any duplicated values in source or target are ignored.
- Calculates the flow/residual capacity for each edge. In the output
  - Edges with zero flow are omitted.
- Creates
  - a **super source** and edges from it to all the sources,
  - a **super target** and edges from it to all the targetss.
- The maximum flow through the graph is guaranteed to be the value returned by [pgr\\_maxFlow](#) when executed with the same parameters and can be calculated:
  - By aggregation of the outgoing flow from the sources
  - By aggregation of the incoming flow to the targets

[pgr\\_maxFlow](#) is the maximum Flow and that maximum is guaranteed to be the same on the functions [pgr\\_pushRelabel](#), [pgr\\_edmondsKarp](#), [pgr\\_boykovKolmogorov](#), but the actual flow through each edge may vary.

Inner Queries

Edges SQL

Capacity edges

- [pgr\\_pushRelabel](#)
- [pgr\\_edmondsKarp](#)
- [pgr\\_boykovKolmogorov](#)

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
capacity	ANY-INTEGER		Weight of the edge (source, target)
reverse_capacity	ANY-INTEGER	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Capacity-Cost edges

- [pgr\\_maxFlowMinCost - Experimental](#)
- [pgr\\_maxFlowMinCost\\_Cost - Experimental](#)

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
capacity	ANY-INTEGER		Capacity of the edge (source, target) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Column	Type	Default	Description
reverse_capacity	ANY-INTEGER	-1	Capacity of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>
cost	ANY-NUMERICAL		Weight of the edge (source, target) if it exist
reverse_cost	ANY-NUMERICAL	\(-1\)	Weight of the edge (target, source) if it exist

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Cost edges

- [pgr\\_edgeDisjointPaths](#)

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations [SQL](#)

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Used in

- [pgr\\_pushRelabel](#)
- [pgr\\_edmondsKarp](#)
- [pgr\\_boykovKolmogorov](#)

Column	Type	Description
seq	INT	Sequential value starting from 1.
edge	BIGINT	Identifier of the edge in the original query (edges_sql).
start_vid	BIGINT	Identifier of the first end point vertex of the edge.
end_vid	BIGINT	Identifier of the second end point vertex of the edge.
flow	BIGINT	Flow through the edge in the direction (start_vid, end_vid).
residual_capacity	BIGINT	Residual capacity of the edge in the direction (start_vid, end_vid).



For [pgr\\_maxFlowMinCost - Experimental](#)

Column	Type	Description
seq	INT	Sequential value starting from 1.
edge	BIGINT	Identifier of the edge in the original query (edges_sql).
source	BIGINT	Identifier of the first end point vertex of the edge.
target	BIGINT	Identifier of the second end point vertex of the edge.
flow	BIGINT	Flow through the edge in the direction (source, target).
residual_capacity	BIGINT	Residual capacity of the edge in the direction (source, target).
cost	FLOAT	The cost of sending this flow through the edge in the direction (source, target).
agg_cost	FLOAT	The aggregate cost.

Advanced Documentation

A flow network is a directed graph where each edge has a capacity and a flow. The flow through an edge must not exceed the capacity of the edge. Additionally, the incoming and outgoing flow of a node must be equal except for source which only has outgoing flow, and the destination(sink) which only has incoming flow.

Maximum flow algorithms calculate the maximum flow through the graph and the flow of each edge.

The maximum flow through the graph is guaranteed to be the same with all implementations, but the actual flow through each edge may vary.

Given the following query:

`pgr_maxFlow`  $((edges\_sql, source\_vertex, sink\_vertex))$

where  $edges\_sql = ((id\_i, source\_i, target\_i, capacity\_i, reverse\_capacity\_i))$

Graph definition

The weighted directed graph,  $G(V,E)$ , is defined as:

- the set of vertices  $V$ 
  - $(source\_vertex \cup sink\_vertex \bigcup source\_i \bigcup target\_i)$
- the set of edges  $E$ 
  - $E = \begin{cases} \text{ } \end{cases} \{ (source\_i, target\_i, capacity\_i) \mid \text{ } \}$  when  $capacity > 0$   $\wedge$   $\text{ } \mid \text{ } \}$  if  $reverse\_capacity = \varnothing$   $\wedge$   $\text{ } \mid \text{ } \}$   $\wedge$   $\text{ } \mid \text{ } \}$   $\wedge$   $\text{ } \mid \text{ } \}$  when  $capacity > 0$   $\wedge$   $\text{ } \mid \text{ } \}$   $\cup \{ (target\_i, source\_i, reverse\_capacity\_i) \mid \text{ } \}$  when  $reverse\_capacity_i > 0$   $\wedge$   $\text{ } \mid \text{ } \}$  if  $reverse\_capacity \neq \varnothing$   $\wedge$   $\text{ } \mid \text{ } \}$   $\end{cases}$

Maximum flow problem

Given:

- $G(V,E)$
- $source\_vertex \in V$  the source vertex
- $sink\_vertex \in V$  the sink vertex

Then:

- $pgr\_maxFlow(edges\_sql, source, sink) = \Phi$
- $\Phi = \{ (id\_i, edge\_id\_i, source\_i, target\_i, flow\_i, residual\_capacity\_i) \}$

Where:

$\Phi$  is a subset of the original edges with their residual capacity and flow. The maximum flow through the graph can be obtained by aggregating on the source or sink and summing the flow from/to it. In particular:

- $(id\_i = i)$
- $(edge\_id = id\_i)$  in `edges_sql`
- $(residual\_capacity\_i = capacity\_i - flow\_i)$

See Also

- [https://en.wikipedia.org/wiki/Maximum\\_flow\\_problem](https://en.wikipedia.org/wiki/Maximum_flow_problem)

Indices and tables

- [Index](#)
- [Search Page](#)

Kruskal - Family of functions

- [pgr\\_kruskal](#)
- [pgr\\_kruskalBFS](#)
- [pgr\\_kruskalDD](#)
- [pgr\\_kruskalDFS](#)

`pgr_kruskal`

`pgr_kruskal` — Minimum spanning tree of a graph using Kruskal's algorithm.

Availability

- Version 3.0.0
  - New official function.

Description

This algorithm finds the minimum spanning forest in a possibly disconnected graph using Kruskal's algorithm.

The main Characteristics are:

- It's implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.
- The total weight of all the edges in the tree or forest is minimized.
- Kruskal's running time:  $\mathcal{O}(E * \log E)$
- EMPTY SET is returned when there are no edges in the graph.

Signatures

Summary

pgr\_kruskal([Edges SQL](#))  
Returns set of (edge, cost)  
OR EMPTY SET

Example:

```
Minimum spanning forest

SELECT * FROM pgr_kruskal(
'SELECT id, source, target, cost, reverse_cost
FROM edges ORDER BY id'
) ORDER BY edge;
edge | cost
-----+-----
1 | 1
2 | 1
3 | 1
6 | 1
7 | 1
10 | 1
11 | 1
12 | 1
13 | 1
14 | 1
15 | 1
16 | 1
17 | 1
18 | 1
(14 rows)
```

Parameters

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (edge, cost)

Column Type Description

edge BIGINT Identifier of the edge.

cost FLOAT Cost to traverse the edge.

See Also

- Spanning Tree - Category
- Kruskal - Family of functions
- Sample Data
- Boost: Kruskal's algorithm
- Wikipedia: Kruskal's algorithm

Indices and tables

- Index
- Search Page

pgr\_kruskalBFS

pgr\_kruskalBFS — Kruskal's algorithm for Minimum Spanning Tree with breadth First Search ordering.

Availability

Version 3.7.0:

- Standardizing output columns to (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
  - Added pred result columns.

Version 3.0.0:

- New official function.

Description

Visits and extracts the nodes information in Breath First Search ordering of the Minimum Spanning Tree created using Kruskal's algorithm.

The main Characteristics are:

- It's implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.
- The total weight of all the edges in the tree or forest is minimized.
- Kruskal's running time:  $\backslash O(E * \log E)\backslash$
- Returned tree nodes from a root vertex are on Breath First Search order
- Breath First Search Running time:  $\backslash O(E + V)\backslash$

Signatures

pgr\_kruskalBFS(Edges SQL, root vid, [max\_depth])  
pgr\_kruskalBFS(Edges SQL, root vids, [max\_depth])  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Single vertex

pgr\_kruskalBFS(Edges SQL, root vid, [max\_depth])  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

The Minimum Spanning Tree having as root vertex  $\backslash(6)\backslash$

```
SELECT * FROM pgr_kruskalBFS(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
6);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 10 | 15 | 3 | 1 | 2
5 | 3 | 6 | 15 | 16 | 16 | 1 | 3
6 | 4 | 6 | 16 | 17 | 15 | 1 | 4
7 | 5 | 6 | 17 | 12 | 13 | 1 | 5
8 | 6 | 6 | 12 | 11 | 11 | 1 | 6
9 | 6 | 6 | 12 | 8 | 12 | 1 | 6
10 | 7 | 6 | 8 | 7 | 10 | 1 | 7
11 | 7 | 6 | 8 | 9 | 14 | 1 | 7
12 | 8 | 6 | 7 | 3 | 7 | 1 | 8
13 | 9 | 6 | 3 | 1 | 6 | 1 | 9
(13 rows)
```

Multiple vertices

pgr\_kruskalBFS(Edges SQL, root vids, [max\_depth])  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

The Minimum Spanning Tree starting on vertices  $\{9, 6\}$  with  $\text{depth} \leq 3$

```
SELECT * FROM pgr_kruskalBFS(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
ARRAY[9, 6], max_depth => 3);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 10 | 15 | 3 | 1 | 2
5 | 3 | 6 | 15 | 16 | 16 | 1 | 3
6 | 0 | 9 | 9 | 9 | -1 | 0 | 0
7 | 1 | 9 | 9 | 8 | 14 | 1 | 1
8 | 2 | 9 | 8 | 7 | 10 | 1 | 2
9 | 2 | 9 | 8 | 12 | 12 | 1 | 2
10 | 3 | 9 | 7 | 3 | 7 | 1 | 3
11 | 3 | 9 | 12 | 11 | 11 | 1 | 3
12 | 3 | 9 | 12 | 17 | 13 | 1 | 3
(12 rows)
```

Parameters

Parameter	Type	Description
Edges SQL	TEXT	Edges SQL as described below.
Root vid	BIGINT	Identifier of the root vertex of the tree.
Root vids	ARRAY[ANY-INTEGER]	Array of identifiers of the root vertices. <ul style="list-style-type: none"><li><math>\{0\}</math> values are ignored</li><li>For optimization purposes, any duplicated value is ignored.</li></ul>
distance	FLOAT	Upper limit for the inclusion of a node in the result.

Where:

ANY-NUMERIC:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

BFS optional parameters

Parameter	Type	Default	Description
max_depth	BIGINT	$\{9223372036854775807\}$	Upper limit of the depth of the tree. <ul style="list-style-type: none"><li>When negative throws an error.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Parameter	Type	Description
seq	BIGINT	Sequential value starting from $\{1\}$ .

Parameter	Type	Description
		Depth of the node.
depth	BIGINT	<ul style="list-style-type: none"><li><math>\backslash(0)</math> when node = start_vid.</li><li><math>\backslash(\text{depth}-1)</math> is the depth of pred</li></ul>
start_vid	BIGINT	Identifier of the root vertex.
		Predecessor of node.
pred	BIGINT	<ul style="list-style-type: none"><li>When node = start_vid then has the value node.</li></ul>
node	BIGINT	Identifier of node reached using edge.
		Identifier of the edge used to arrive from pred to node.
edge	BIGINT	<ul style="list-style-type: none"><li><math>\backslash(-1)</math> when node = start_vid.</li></ul>
cost	FLOAT	Cost to traverse edge.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

See Also

- [Spanning Tree - Category](#)
- [Kruskal - Family of functions](#)
- [Sample Data](#)
- [Boost: Kruskal's algorithm](#)
- [Wikipedia: Kruskal's algorithm](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_kruskalDD

pgr\_kruskalDD — Catchment nodes using Kruskal's algorithm.

Availability

Version 3.7.0:

- Standardizing output columns to (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
  - Added pred result columns.

Version 3.0.0:

- New official function.

Description

Using Kruskal's algorithm, extracts the nodes that have aggregate costs less than or equal to **distance** from a **root** vertex (or vertices) within the calculated minimum spanning tree.

The main Characteristics are:

- It's implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.
- The total weight of all the edges in the tree or forest is minimized.
- Kruskal's running time:  $\backslash(O(E * \log E)\backslash)$
- Extracts all the nodes that have costs less than or equal to the value distance.
- The edges extracted will conform to the corresponding spanning tree.
- Edge  $\backslash((u, v)\backslash)$  will not be included when:
  - The distance from the **root** to  $\backslash(u)\backslash >$  limit distance.
  - The distance from the **root** to  $\backslash(v)\backslash >$  limit distance.
  - No new nodes are created on the graph, so when is within the limit and is not within the limit, the edge is not included.
- Returned tree nodes from a root vertex are on Depth First Search order.
- Depth First Search running time:  $\backslash(O(E + V)\backslash)$

 Boost Graph Inside

Signatures

pgr\_kruskalDD([Edges SQL](#), **root vid**, **distance**)

pgr\_kruskalDD([Edges SQL](#), **root vids**, **distance**)  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Single vertex<sup>1</sup>

pgr\_kruskalDD([Edges SQL](#), **root vid**, **distance**)  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

The Minimum Spanning Tree starting on vertex \6\ with \distance \leq 3.5\

```
SELECT * FROM pgr_kruskalDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
6, 3.5);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 10 | 15 | 3 | 1 | 2
5 | 3 | 6 | 15 | 16 | 16 | 1 | 3
(5 rows)
```

Multiple vertices<sup>1</sup>

pgr\_kruskalDD([Edges SQL](#), **root vids**, **distance**)  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

The Minimum Spanning Tree starting on vertices \{9, 6\} with \distance \leq 3.5\

```
SELECT * FROM pgr_kruskalDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
ARRAY[9, 6], 3.5);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 10 | 15 | 3 | 1 | 2
5 | 3 | 6 | 15 | 16 | 16 | 1 | 3
6 | 0 | 9 | 9 | 9 | -1 | 0 | 0
7 | 1 | 9 | 9 | 8 | 14 | 1 | 1
8 | 2 | 9 | 8 | 7 | 10 | 1 | 2
9 | 3 | 9 | 7 | 3 | 7 | 1 | 3
10 | 2 | 9 | 8 | 12 | 12 | 1 | 2
11 | 3 | 9 | 12 | 11 | 11 | 1 | 3
12 | 3 | 9 | 12 | 17 | 13 | 1 | 3
(12 rows)
```

Parameters<sup>1</sup>

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	Edges SQL as described below.
<b>Root vid</b>	BIGINT	Identifier of the root vertex of the tree.
<b>Root vids</b>	ARRAY[ANY-INTEGER]	Array of identifiers of the root vertices. <ul style="list-style-type: none"><li>\(0\) values are ignored</li><li>For optimization purposes, any duplicated value is ignored.</li></ul>
<b>distance</b>	FLOAT	Upper limit for the inclusion of a node in the result.

Where:

ANY-NUMERIC:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Inner Queries<sup>1</sup>

Edges SQL<sup>1</sup>

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns¶

Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Parameter	Type	Description
seq	BIGINT	Sequential value starting from \(\backslash\).
depth	BIGINT	Depth of the node. <ul style="list-style-type: none"><li>• \(\backslash\)</li><li>• \(\backslash\)</li></ul>
start_vid	BIGINT	Identifier of the root vertex.
pred	BIGINT	Predecessor of node. <ul style="list-style-type: none"><li>• When node = start_vid then has the value node.</li></ul>
node	BIGINT	Identifier of node reached using edge.
edge	BIGINT	Identifier of the edge used to arrive from pred to node. <ul style="list-style-type: none"><li>• \(\backslash\)</li></ul>
cost	FLOAT	Cost to traverse edge.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

See Also¶

- [Spanning Tree - Category](#)
- [Kruskal - Family of functions](#)
- [Sample Data](#)
- [Boost: Kruskal's algorithm](#)
- [Wikipedia: Kruskal's algorithm](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_kruskalDFS¶

pgr\_kruskalDFS — Kruskal's algorithm for Minimum Spanning Tree with Depth First Search ordering.

Availability

Version 3.7.0:

- Standardizing output columns to (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
  - Added pred result columns.

Version 3.0.0:

- New official function.

Description¶

Visits and extracts the nodes information in Depth First Search ordering of the Minimum Spanning Tree created using Kruskal's algorithm.

The main Characteristics are:

- It's implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.
- The total weight of all the edges in the tree or forest is minimized.
- Kruskal's running time:  $\mathcal{O}(E * \log E)$
- Returned tree nodes from a root vertex are on Depth First Search order
- Depth First Search Running time:  $\mathcal{O}(E + V)$

Boost Graph Inside

Signatures¶

pgr\_kruskalDFS([Edges SQL](#), **root vid**, [max\_depth])  
pgr\_kruskalDFS([Edges SQL](#), **root vids**, [max\_depth])  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Single vertex

pgr\_kruskalDFS([Edges SQL](#), **root vid**, [max\_depth])  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

The Minimum Spanning Tree having as root vertex \6\

```
SELECT * FROM pgr_kruskalDFS(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
6);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
 1 |    0 |         6 |    6 |   6 |   -1 |    0 |         0
 2 |    1 |         6 |    6 |   5 |    1 |    1 |         1
 3 |    1 |         6 |    6 |  10 |    2 |    1 |         1
 4 |    2 |         6 |   10 |  15 |    3 |    1 |         2
 5 |    3 |         6 |   15 |  16 |   16 |    1 |         3
 6 |    4 |         6 |   16 |  17 |   15 |    1 |         4
 7 |    5 |         6 |   17 |  12 |   13 |    1 |         5
 8 |    6 |         6 |   12 |  11 |   11 |    1 |         6
 9 |    6 |         6 |   12 |   8 |   12 |    1 |         6
10 |    7 |         6 |    8 |   7 |   10 |    1 |         7
11 |    8 |         6 |    7 |   3 |    7 |    1 |         8
12 |    9 |         6 |    3 |   1 |    6 |    1 |         9
13 |    7 |         6 |    8 |   9 |   14 |    1 |         7
(13 rows)
```

Multiple vertices

pgr\_kruskalDFS([Edges SQL](#), **root vids**, [max\_depth])  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

The Minimum Spanning Tree starting on vertices \9, 6\ with \depth <= 3\

```
SELECT * FROM pgr_kruskalDFS(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
ARRAY[9, 6], max_depth => 3);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
 1 |    0 |         6 |    6 |   6 |   -1 |    0 |         0
 2 |    1 |         6 |    6 |   5 |    1 |    1 |         1
 3 |    1 |         6 |    6 |  10 |    2 |    1 |         1
 4 |    2 |         6 |   10 |  15 |    3 |    1 |         2
 5 |    3 |         6 |   15 |  16 |   16 |    1 |         3
 6 |    0 |         9 |    9 |   9 |   -1 |    0 |         0
 7 |    1 |         9 |    9 |   8 |   14 |    1 |         1
 8 |    2 |         9 |    8 |   7 |   10 |    1 |         2
 9 |    3 |         9 |    7 |   3 |    7 |    1 |         3
10 |    2 |         9 |    8 |  12 |   12 |    1 |         2
11 |    3 |         9 |   12 |  11 |   11 |    1 |         3
12 |    3 |         9 |   12 |  17 |   13 |    1 |         3
(12 rows)
```

Parameters

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	Edges SQL as described below.
<b>Root vid</b>	BIGINT	Identifier of the root vertex of the tree.
<b>Root vids</b>	ARRAY[ANY-INTEGER]	Array of identifiers of the root vertices. <ul style="list-style-type: none"><li>\0\ values are ignored</li><li>For optimization purposes, any duplicated value is ignored.</li></ul>

**distance**      FLOAT      Upper limit for the inclusion of a node in the result.

Where:

ANY-NUMERIC:  
SMALLINT, INTEGER, BIGINT, REAL, FLOAT

DFS optional parameters

Parameter	Type	Default	Description
			Upper limit of the depth of the tree.
max_depth	BIGINT	\9223372036854775807\	<ul style="list-style-type: none"><li>When negative throws an error.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.



Column	Type	Default	Description
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns<sup>1</sup>

Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Parameter	Type	Description
seq	BIGINT	Sequential value starting from $\backslash(1\backslash)$ .
depth	BIGINT	Depth of the node. <ul style="list-style-type: none"> <li><math>\backslash(0\backslash)</math> when node = start_vid.</li> <li><math>\backslash(\text{depth}-1\backslash)</math> is the depth of <math>\text{pred}</math></li> </ul>
start_vid	BIGINT	Identifier of the root vertex.
pred	BIGINT	Predecessor of node. <ul style="list-style-type: none"> <li>When node = start_vid then has the value node.</li> </ul>
node	BIGINT	Identifier of node reached using edge.
edge	BIGINT	Identifier of the edge used to arrive from pred to node. <ul style="list-style-type: none"> <li><math>\backslash(-1\backslash)</math> when node = start_vid.</li> </ul>
cost	FLOAT	Cost to traverse edge.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

See Also<sup>1</sup>

- [Spanning Tree - Category](#)
- [Kruskal - Family of functions](#)
- [Sample Data](#)
- [Boost: Kruskal's algorithm](#)
- [Wikipedia: Kruskal's algorithm](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Description<sup>1</sup>

Kruskal's algorithm is a greedy minimum spanning tree algorithm that in each cycle finds and adds the edge of the least possible weight that connects any two trees in the forest.

**The main Characteristics are:**

- It's implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.
- The total weight of all the edges in the tree or forest is minimized.
- Kruskal's running time:  $\backslash(O(E * \log E)\backslash)$

Inner Queries<sup>1</sup>

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.

Column	Type	Default	Description
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
target	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
cost	<b>ANY-NUMERICAL</b>		Weight of the edge (source, target)
reverse_cost	<b>ANY-NUMERICAL</b>	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

See Also

- [Spanning Tree - Category](#)
- [Boost: Kruskal's algorithm](#)
- [Boost: Prim's algorithm](#)
- [Wikipedia: Kruskal's algorithm](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Metrics - Family of functions

- [pgr\\_degree](#) - Returns a set of vertices and corresponding count of incident edges to the vertex.

☐ Experimental

☐ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting
- [pgr\\_betweennessCentrality - Experimental](#) - Calculates relative betweenness centrality using Brandes Algorithm

[pgr\\_degree](#)

[pgr\\_degree](#) — For each vertex in an undirected graph, return the count of edges incident to the vertex.

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)

- Functionality might not change. (But still can)
- pgTap tests have being done. But might need more.
- Documentation might need refinement.

#### Availability

#### Version 3.8.0

- Error messages adjustment.
- New signature with only Edges SQL.
- Function promoted to official.

#### Version 3.4.0

- New proposed function.

#### Description

Calculates the degree of the vertices of an undirected graph

The degree (or valency) of a vertex of a graph is the number of edges that are incident to the vertex.

- Works for **undirected** graphs.
- A loop contributes 2 to a vertex's degree.
- A vertex with degree 0 is called an isolated vertex.
  - Isolated vertex is not part of the result
- Vertex not participating on the subgraph is considered and isolated vertex.
- There can be a dryrun execution and the code used to get the answer will be shown in a PostgreSQLNOTICE.
  - The code can be used as base code for the particular application requirements.
- No ordering is performed.

#### Signatures

```
pgr_degree(Edges SQL , [dryrun])
pgr_degree(Edges SQL , Vertex SQL, [dryrun])
RETURNS SETOF (node, degree)
OR EMPTY SET
```

#### Edges

```
pgr_degree(Edges SQL , [dryrun])
RETURNS SETOF (node, degree)
OR EMPTY SET
```

example:

Get the degree of the vertices defined on the edges table

```
SELECT * FROM pgr_degree($$SELECT id, source, target FROM edges$$)
ORDER BY node;
node | degree
```

```
-----+-----
1 | 1
2 | 1
3 | 2
4 | 1
5 | 1
6 | 3
7 | 4
8 | 3
9 | 1
10 | 3
11 | 4
12 | 3
13 | 1
14 | 1
15 | 2
16 | 3
17 | 2
(17 rows)
```

#### Edges and Vertices

```
pgr_degree(Edges SQL , Vertex SQL, [dryrun])
RETURNS SETOF (node, degree)
OR EMPTY SET
```

Example:

Extracting the vertex information

pgr\_degree can use [pgr\\_extractVertices](#) embedded in the call.

For decent size networks, it is best to prepare your vertices table before hand and use it on pgr\_degree calls. (See [Using a vertex table](#))

Calculate the degree of the nodes:

```
SELECT * FROM pgr_degree(
  $$SELECT id FROM edges$$,
  $$SELECT id, in_edges, out_edges
    FROM pgr_extractVertices('SELECT id, geom FROM edges')$$);
node | degree
```

```
-----+-----
1 | 1
2 | 1
3 | 2
4 | 1
5 | 1
6 | 3
7 | 4
8 | 3
9 | 1
10 | 3
11 | 4
12 | 3
13 | 1
14 | 1
15 | 2
```

16 | 3  
17 | 2  
(17 rows)

Parameters¶

Parameter Type Description

Edges SQL TEXT Edges SQL as described below

Vertex SQL TEXT Vertex SQL as described below

Optional parameters¶

Parameter	Type	Default	Description
dryrun	BOOLEAN	false	<ul style="list-style-type: none"><li>When true do not process and get in a NOTICE the resulting query.</li></ul>

Inner Queries¶

- Edges SQL
- Vertex SQL

Edges SQL¶

For the Edges and Vertices signature:

Column Type Description

id BIGINT Identifier of the edge.

For the Edges signature:

Column Type Description

id BIGINT Identifier of the edge.

source BIGINT Identifier of the first end point vertex of the edge.

target BIGINT Identifier of the second end point vertex of the edge.

Vertex SQL¶

For the Edges and Vertices signature:

Column Type Description

id BIGINT Identifier of the first end point vertex of the edge.

in\_edges BIGINT[] Array of identifiers of the edges that have the vertexid as *first end point*.

- When missing, out\_edges must exist.

out\_edges BIGINT[] Array of identifiers of the edges that have the vertexid as *second end point*.

- When missing, in\_edges must exist.

Result columns¶

Column Type Description

node BIGINT Vertex identifier

degree BIGINT Number of edges that are incident to the vertex id

Additional Examples¶

- Degree of a loop
- Degree of a sub graph
- Using a vertex table
- Dry run execution
- Finding dead ends
- Finding linear vertices

Degree of a loop¶

A loop contributes 2 to a vertex's degree.

Using the [Edges](#) signature.

```
SELECT * from pgr_degree('SELECT 1 as id, 2 as source, 2 as target');
node | degree
-----+-----
2 | 2
(1 row)
```

Using the [Edges and Vertices](#) signature.

```
SELECT * FROM pgr_degree(
  $$SELECT 1 AS id$$,
  $$SELECT id, in_edges, out_edges
    FROM pgr_extractVertices('SELECT 1 as id, 2 as source, 2 as target')$$);
node | degree
-----+-----
2 | 2
(1 row)
```

Degree of a sub graph¶

For the following is a subgraph of the [Sample Data](#):

- $E = \{(1, 5 \rightarrow 6), (1, 6 \rightarrow 10)\}$
- $V = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17\}$

The vertices not participating on the edge are considered isolated

- their degree is 0 in the subgraph and
- their degree is not shown in the output.

Using the [Edges](#) signature.

```
SELECT * FROM pgr_degree($$SELECT * FROM edges WHERE id IN (1, 2)$$);
node | degree
-----+-----
10 | 1
6 | 2
5 | 1
(3 rows)
```

Using the [Edges and Vertices](#) signature.

```
SELECT * FROM pgr_degree(
  $$SELECT * FROM edges WHERE id IN (1, 2)$$,
  $$SELECT id, in_edges, out_edges FROM vertices$$);
node | degree
-----+-----
5 | 1
6 | 2
10 | 1
(3 rows)
```

Using a vertex table¶

For decent size networks, it is best to prepare your vertices table before hand and use it on `pgr_degree` calls.

Extract the vertex information and save into a table:

```
CREATE TABLE vertices AS
SELECT id, in_edges, out_edges
FROM pgr_extractVertices('SELECT id, geom FROM edges');
SELECT 17
```

Calculate the degree of the nodes:

```
SELECT * FROM pgr_degree(
  $$SELECT id FROM edges$$,
  $$SELECT id, in_edges, out_edges FROM vertices$$);
node | degree
-----+-----
1 | 1
2 | 1
3 | 2
4 | 1
5 | 1
6 | 3
7 | 4
8 | 3
9 | 1
10 | 3
11 | 4
12 | 3
13 | 1
14 | 1
15 | 2
16 | 3
17 | 2
(17 rows)
```

Dry run execution¶

To get the query generated used to get the vertex information, use `dryrun => true`.

The results can be used as base code to make a refinement based on the backend development needs.

```
SELECT * FROM pgr_degree(
  $$SELECT id FROM edges WHERE id < 17$$,
  $$SELECT id, in_edges, out_edges FROM vertices$$,
  dryrun => true);
NOTICE:
  WITH
    -- a sub set of edges of the graph goes here
    g_edges AS (
      SELECT id FROM edges WHERE id < 17
    ),
    -- sub set of vertices of the graph goes here
```

```

all_vertices AS (
  SELECT id, in_edges, out_edges FROM vertices
),

g_vertices AS (
  SELECT id,
    unnest(
      coalesce(in_edges::BIGINT[], '{}':BIGINT[])
    ||
      coalesce(out_edges::BIGINT[], '{}':BIGINT[])
    ) AS eid
  FROM all_vertices
),

totals AS (
  SELECT v.id, count(*)
  FROM g_vertices v
  JOIN g_edges e ON (v.eid = e.id) GROUP BY v.id
)

SELECT id::BIGINT, count::BIGINT FROM all_vertices JOIN totals USING (id)
;
node | degree
-----+-----
(0 rows)

```

#### Finding dead ends¶

If there is a vertices table already built using `pgr_extractVertices` and want the degree of the whole graph rather than a subset, it can be forgo using `pgr_degree` and work with the `in_edges` and `out_edges` columns directly.

The degree of a dead end is 1.

To get the dead ends:

```


SELECT id FROM vertices
WHERE array_length(in_edges || out_edges, 1) = 1;
id
----
1
2
5
(3 rows)

```

A dead end happens when

- The vertex is the limit of a cul-de-sac, a no-through road or a no-exit road.
- The vertex is on the limit of the imported graph.
  - If a larger graph is imported then the vertex might not be a dead end

Node `\(4\)`, is a dead end on the query, even that it visually looks like an end point of 3 edges.



Is node `\(4\)` a dead end or not?

The answer to that question will depend on the application.

- Is there such a small curb:
  - That does not allow a vehicle to use that visual intersection?
  - Is the application for pedestrians and therefore the pedestrian can easily walk on the small curb?
  - Is the application for the electricity and the electrical lines than can easily be extended on top of the small curb?
- Is there a big cliff and from eagles view look like the dead end is close to the segment?

Depending on the answer, modification of the data might be needed.

When there are many dead ends, to speed up processing, the [Contraction - Family of functions](#) functions can be used to contract the graph.

#### Finding linear vertices¶

The degree of a linear vertex is 2.

If there is a vertices table already built using the `pgr_extractVertices`

To get the linear edges:

```

SELECT id FROM vertices
WHERE array_length(in_edges || out_edges, 1) = 2;
id
----
3
9
13
15
16
(5 rows)

```

These linear vertices are correct, for example, when those the vertices are speed bumps, stop signals and the application is taking them into account.

When there are many linear vertices, that need not to be taken into account, to speed up the processing, the [Contraction - Family of functions](#) functions can be used to contract the problem.

#### See Also¶

- [Topology - Family of Functions](#)
- [pgr\\_extractVertices](#)

Indices and tables

- [Index](#)

- [Search Page](#)

pgr\_betweennessCentrality - Experimental

pgr\_betweennessCentrality - Calculates the relative betweenness centrality using Brandes Algorithm

Availability

- Version 3.7.0
  - New experimental function.

Description

The Brandes Algorithm takes advantage of the sparse graphs for evaluating the betweenness centrality score of all vertices.

Betweenness centrality measures the extent to which a vertex lies on the shortest paths between all other pairs of vertices. Vertices with a high betweenness centrality score may have considerable influence in a network by the virtue of their control over the shortest paths passing between them.

The removal of these vertices will affect the network by disrupting the it, as most of the shortest paths between vertices pass through them.

This implementation work for both directed and undirected graphs.

- Running time:  $\Theta(V E)$
- Running space:  $\Theta(V E)$
- Throws when there are no edges in the graph

 Boost Graph Inside

Signatures

Summary

pgr\_betweennessCentrality([Edges SQL](#), [directed])

Returns set of (vid, centrality)

Example:

For a directed graph with edges  $\{(1, 2, 3, 4)\}$ .

```
SELECT * FROM pgr_betweennessCentrality(
'SELECT id, source, target, cost, reverse_cost
FROM edges where id < 5'
) ORDER BY vid;
vid | centrality
-----+-----
5 | 0
6 | 0.5
7 | 0
10 | 0.25
15 | 0
(5 rows)
```

Explanation

- The betweenness centrality are between parenthesis.
- The leaf vertices have betweenness centrality  $\{0\}$ .
- Betweenness centrality of vertex  $\{6\}$  is higher than of vertex  $\{10\}$ .
  - Removing vertex  $\{6\}$  will create three graph components.
  - Removing vertex  $\{10\}$  will create two graph components.

Parameters

Parameter	Type	Default	Description
<a href="#">Edges SQL</a>	TEXT		<a href="#">Edges SQL</a> as described below.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)

Column	Type	Default	Description
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Column	Type	Description
vid	BIGINT	Identifier of the vertex.
centrality	FLOAT	Relative betweenness centrality score of the vertex (will be in range [0,1])

See Also

- [Sample Data](#)
- [Boost: betweenness centrality](#)

Indices and tables

- [Index](#)
- [Search Page](#)

See Also

Indices and tables

- [Index](#)
- [Search Page](#)

Prim - Family of functions

- [pgr\\_prim](#)
- [pgr\\_primBFS](#)
- [pgr\\_primDD](#)
- [pgr\\_primDFS](#)

pgr\_prim

pgr\_prim — Minimum spanning forest of a graph using Prim's algorithm.

Availability

- Version 3.0.0
  - New official function.

Description

This algorithm finds the minimum spanning forest in a possibly disconnected graph using Prim's algorithm.

The main characteristics are:

- It's implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.
- Prim's running time:  $\mathcal{O}(E * \log V)$
- EMPTY SET is returned when there are no edges in the graph.

 Boost Graph Inside

Signatures

Summary

pgr\_prim([Edges SQL](#))  
Returns set of (edge, cost)  
OR EMPTY SET

Example:

Minimum spanning forest of a subgraph



```
SELECT edge, cost FROM pgr_prim(
  'SELECT id, source, target, cost, reverse_cost
  FROM edges WHERE id < 14'
) ORDER BY edge;
edge | cost
-----+-----
1 | 1
2 | 1
3 | 1
4 | 1
6 | 1
7 | 1
8 | 1
9 | 1
10 | 1
12 | 1
13 | 1
(11 rows)
```

Parameters1

Parameter	Type	Description
-----------	------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below.

Inner Queries1

Edges SQL1

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns1

Returns set of (edge, cost)

Column	Type	Description
edge	BIGINT	Identifier of the edge.
cost	FLOAT	Cost to traverse the edge.

See Also1

- [Spanning Tree - Category](#)
- [Prim - Family of functions](#)
- [Sample Data](#)
- [Boost: Prim's algorithm documentation](#)
- [Wikipedia: Prim's algorithm](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_primBFS1

pgr\_primBFS — Prim's algorithm for Minimum Spanning Tree with Depth First Search ordering.

Availability

Version 3.7.0:

- Standardizing output columns to (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
  - Added pred result columns.

Version 3.0.0:

- New official function.

Description

Visits and extracts the nodes information in Breath First Search ordering of the Minimum Spanning Tree created using Prim's algorithm.

The main Characteristics are:

- It's implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.
- Prim's running time:  $\backslash(O(E * \log V))$
- Returned tree nodes from a root vertex are on Breath First Search order
- Breath First Search Running time:  $\backslash(O(E + V))$

Boost Graph Inside

Signatures

pgr\_primBFS([Edges SQL](#), **root vid**, [max\_depth])  
pgr\_primBFS([Edges SQL](#), **root vids**, [max\_depth])  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Single vertex

pgr\_primBFS([Edges SQL](#), **root vid**, [max\_depth])  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

The Minimum Spanning Tree having as root vertex  $\backslash(6)$

```
SELECT * FROM pgr_primBFS(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
6);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 6 | 10 | 2 | 1 | 1
4 | 1 | 6 | 6 | 7 | 4 | 1 | 1
5 | 2 | 6 | 10 | 15 | 3 | 1 | 2
6 | 2 | 6 | 10 | 11 | 5 | 1 | 2
7 | 2 | 6 | 7 | 3 | 7 | 1 | 2
8 | 2 | 6 | 7 | 8 | 10 | 1 | 2
9 | 3 | 6 | 11 | 16 | 9 | 1 | 3
10 | 3 | 6 | 11 | 12 | 11 | 1 | 3
11 | 3 | 6 | 3 | 1 | 6 | 1 | 3
12 | 3 | 6 | 8 | 9 | 14 | 1 | 3
13 | 4 | 6 | 12 | 17 | 13 | 1 | 4
(13 rows)
```

Multiple vertices

pgr\_primBFS([Edges SQL](#), **root vids**, [max\_depth])  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

The Minimum Spanning Tree starting on vertices  $\backslash(\{9, 6\})$  with  $\backslash(\text{depth} \leq 3)$

```
SELECT * FROM pgr_primBFS(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
ARRAY[9, 6], max_depth => 3);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 6 | 10 | 2 | 1 | 1
4 | 1 | 6 | 6 | 7 | 4 | 1 | 1
5 | 2 | 6 | 10 | 15 | 3 | 1 | 2
6 | 2 | 6 | 10 | 11 | 5 | 1 | 2
7 | 2 | 6 | 7 | 3 | 7 | 1 | 2
8 | 2 | 6 | 7 | 8 | 10 | 1 | 2
9 | 3 | 6 | 11 | 16 | 9 | 1 | 3
10 | 3 | 6 | 11 | 12 | 11 | 1 | 3
11 | 3 | 6 | 3 | 1 | 6 | 1 | 3
12 | 3 | 6 | 8 | 9 | 14 | 1 | 3
13 | 0 | 9 | 9 | 9 | -1 | 0 | 0
14 | 1 | 9 | 9 | 8 | 14 | 1 | 1
15 | 2 | 9 | 8 | 7 | 10 | 1 | 2
16 | 3 | 9 | 7 | 6 | 4 | 1 | 3
17 | 3 | 9 | 7 | 3 | 7 | 1 | 3
(17 rows)
```

Parameters

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	Edges SQL as described below.
Root vid	BIGINT	Identifier of the root vertex of the tree.
Root vids	ARRAY[ANY-INTEGER]	Array of identifiers of the root vertices. <ul style="list-style-type: none"><li>• <math>\backslash(0)</math> values are ignored</li><li>• For optimization purposes, any duplicated value is ignored.</li></ul>
distance	FLOAT	Upper limit for the inclusion of a node in the result.

Where:

ANY-NUMERIC:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

BFS optional parameters

Parameter	Type	Default	Description
max_depth	BIGINT	\\(9223372036854775807\\)	Upper limit of the depth of the tree. <ul style="list-style-type: none"><li>When negative throws an error.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Parameter	Type	Description
seq	BIGINT	Sequential value starting from \\(1\\).
depth	BIGINT	Depth of the node. <ul style="list-style-type: none"><li>\\(0\\) when node = start_vid.</li><li>\\(depth-1\\) is the depth of pred</li></ul>
start_vid	BIGINT	Identifier of the root vertex.
pred	BIGINT	Predecessor of node. <ul style="list-style-type: none"><li>When node = start_vid then has the value node.</li></ul>
node	BIGINT	Identifier of node reached using edge.
edge	BIGINT	Identifier of the edge used to arrive from pred to node. <ul style="list-style-type: none"><li>\\(-1\\) when node = start_vid.</li></ul>
cost	FLOAT	Cost to traverse edge.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

See Also

- Spanning Tree - Category
- Prim - Family of functions
- Sample Data
- Boost: Prim's algorithm
- Wikipedia: Prim's algorithm

Indices and tables

- Index
- Search Page

pgr\_primDD — Catchment nodes using Prim’s algorithm.

Availability

Version 3.7.0

- Standardizing output columns to (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
  - Added pred result columns.

Version 3.0.0

- New official function.

Description¶

Using Prim’s algorithm, extracts the nodes that have aggregate costs less than or equal to a distance from a root vertex (or vertices) within the calculated minimum spanning tree.

The main Characteristics are:

- It’s implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.
- Prim’s running time:  $\backslash(O(E * \log V))$
- Extracts all the nodes that have costs less than or equal to the value distance.
- The edges extracted will conform to the corresponding spanning tree.
- Edge  $\backslash((u, v))$  will not be included when:
  - The distance from the **root** to  $\backslash(u) >$  limit distance.
  - The distance from the **root** to  $\backslash(v) >$  limit distance.
  - No new nodes are created on the graph, so when is within the limit and is not within the limit, the edge is not included.
- Returned tree nodes from a root vertex are on Depth First Search order.
- Depth First Search running time:  $\backslash(O(E + V))$

 Boost Graph Inside

Signatures¶

pgr\_primDD([Edges SQL](#), **root vid**, **distance**)  
pgr\_primDD([Edges SQL](#), **root vids**, **distance**)  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Single vertex¶

pgr\_primDD([Edges SQL](#), **root vid**, **distance**)  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

The Minimum Spanning Tree starting on vertex  $\backslash(6)$  with  $\backslash(\text{distance} \leq 3.5)$

```
SELECT * FROM pgr_primDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
6, 3.5);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
 1 |    0 |         6 |    6 |    6 |   -1 |    0 |         0
 2 |    1 |         6 |    6 |    5 |    1 |    1 |         1
 3 |    1 |         6 |    6 |   10 |    2 |    1 |         1
 4 |    2 |         6 |   10 |   15 |    3 |    1 |         2
 5 |    2 |         6 |   10 |   11 |    5 |    1 |         2
 6 |    3 |         6 |   11 |   16 |    9 |    1 |         3
 7 |    3 |         6 |   11 |   12 |   11 |    1 |         3
 8 |    1 |         6 |    6 |    7 |    4 |    1 |         1
 9 |    2 |         6 |    7 |    3 |    7 |    1 |         2
10 |    3 |         6 |    3 |    1 |    6 |    1 |         3
11 |    2 |         6 |    7 |   10 |    8 |    1 |         2
12 |    3 |         6 |    8 |    9 |   14 |    1 |         3
(12 rows)
```

Multiple vertices¶

pgr\_primDD([Edges SQL](#), **root vids**, **distance**)  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

The Minimum Spanning Tree starting on vertices  $\backslash(\{9, 6\})$  with  $\backslash(\text{distance} \leq 3.5)$

```
SELECT * FROM pgr_primDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
ARRAY[9, 6], 3.5);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
 1 |    0 |         6 |    6 |    6 |   -1 |    0 |         0
 2 |    1 |         6 |    6 |    5 |    1 |    1 |         1
 3 |    1 |         6 |    6 |   10 |    2 |    1 |         1
 4 |    2 |         6 |   10 |   15 |    3 |    1 |         2
 5 |    2 |         6 |   10 |   11 |    5 |    1 |         2
 6 |    3 |         6 |   11 |   16 |    9 |    1 |         3
 7 |    3 |         6 |   11 |   12 |   11 |    1 |         3
 8 |    1 |         6 |    6 |    7 |    4 |    1 |         1
 9 |    2 |         6 |    7 |    3 |    7 |    1 |         2
10 |    3 |         6 |    3 |    1 |    6 |    1 |         3
11 |    2 |         6 |    7 |   10 |    8 |    1 |         2
12 |    3 |         6 |    8 |    9 |   14 |    1 |         3
13 |    0 |         9 |    9 |    9 |   -1 |    0 |         0
14 |    1 |         9 |    9 |    8 |   14 |    1 |         1
15 |    2 |         9 |    8 |   10 |    7 |    1 |         2
```

16		3		9		7		6		4		1		3
17		3		9		7		3		7		1		3
(17 rows)														

Parameters1

Parameter	Type	Description
Edges SQL	TEXT	Edges SQL as described below.
Root vid	BIGINT	Identifier of the root vertex of the tree.
Root vids	ARRAY[ANY-INTEGER]	Array of identifiers of the root vertices.
		• \0 values are ignored
		• For optimization purposes, any duplicated value is ignored.
distance	FLOAT	Upper limit for the inclusion of a node in the result.

Where:

ANY-NUMERIC:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Inner Queries1

Edges SQL1

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source)
			• When negative: edge (target, source) does not exist, therefore it's not part of the graph.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns1

Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Parameter	Type	Description
seq	BIGINT	Sequential value starting from \1).
depth	BIGINT	Depth of the node.
		• \0 when node = start_vid.
		• \depth-1 is the depth of pred
start_vid	BIGINT	Identifier of the root vertex.
pred	BIGINT	Predecessor of node.
		• When node = start_vid then has the value node.
node	BIGINT	Identifier of node reached using edge.
edge	BIGINT	Identifier of the edge used to arrive from pred to node.
		• \-1 when node = start_vid.
cost	FLOAT	Cost to traverse edge.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

See Also1

- Spanning Tree - Category

- [Prim - Family of functions](#)
- [Sample Data](#)
- [Boost: Prim's algorithm](#)
- [Wikipedia: Prim's algorithm](#)

Indices and tables

- [Index](#)
- [Search Page](#)

`pgr_primDFS`

`pgr_primDFS` — Prim algorithm for Minimum Spanning Tree with Depth First Search ordering.

Availability

Version 3.7.0:

- Standardizing output columns to (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
  - Added `pred` result columns.

Version 3.0.0:

- New official function.

Description

Visits and extracts the nodes information in Depth First Search ordering of the Minimum Spanning Tree created using Prim's algorithm.

The main Characteristics are:

- It's implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.
- Prim's running time:  $\backslash(O(E * \log V))$
- Returned tree nodes from a root vertex are on Depth First Search order
- Depth First Search Running time:  $\backslash(O(E + V))$

`Boost Graph Inside`

Signatures

`pgr_primDFS(Edges SQL, root vid, [max_depth])`  
`pgr_primDFS(Edges SQL, root vids, [max_depth])`  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Single vertex

`pgr_primDFS(Edges SQL, root vid, [max_depth])`  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

The Minimum Spanning Tree having as root vertex  $\backslash(6)$

```
SELECT * FROM pgr_primDFS(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
6);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 10 | 15 | 3 | 1 | 2
5 | 2 | 6 | 10 | 11 | 5 | 1 | 2
6 | 3 | 6 | 11 | 16 | 9 | 1 | 3
7 | 3 | 6 | 11 | 12 | 11 | 1 | 3
8 | 4 | 6 | 12 | 17 | 13 | 1 | 4
9 | 1 | 6 | 6 | 7 | 4 | 1 | 1
10 | 2 | 6 | 7 | 3 | 7 | 1 | 2
11 | 3 | 6 | 3 | 1 | 6 | 1 | 3
12 | 2 | 6 | 7 | 8 | 10 | 1 | 2
13 | 3 | 6 | 8 | 9 | 14 | 1 | 3
(13 rows)
```

Multiple vertices

`pgr_primDFS(Edges SQL, root vids, [max_depth])`  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Example:

The Minimum Spanning Tree starting on vertices  $\backslash(\{9, 6\})$  with  $\backslash(\text{depth} \leq 3)$

```
SELECT * FROM pgr_primDFS(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
ARRAY[9, 6], max_depth => 3);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 10 | 15 | 3 | 1 | 2
5 | 2 | 6 | 10 | 11 | 5 | 1 | 2
6 | 3 | 6 | 11 | 16 | 9 | 1 | 3
7 | 3 | 6 | 11 | 12 | 11 | 1 | 3
8 | 1 | 6 | 6 | 7 | 4 | 1 | 1
9 | 2 | 6 | 7 | 3 | 7 | 1 | 2
```

10	3	6	3	1	6	1	3
11	2	6	7	8	10	1	2
12	3	6	8	9	14	1	3
13	0	9	9	9	-1	0	0
14	1	9	9	8	14	1	1
15	2	9	8	7	10	1	2
16	3	9	7	6	4	1	3
17	3	9	7	3	7	1	3

(17 rows)

Parameters1

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	Edges SQL as described below.
<b>Root vid</b>	BIGINT	Identifier of the root vertex of the tree.
<b>Root vids</b>	ARRAY[ANY-INTEGER]	Array of identifiers of the root vertices.
		<ul style="list-style-type: none"><li>• \0 values are ignored</li></ul>
		<ul style="list-style-type: none"><li>• For optimization purposes, any duplicated value is ignored.</li></ul>
<b>distance</b>	FLOAT	Upper limit for the inclusion of a node in the result.

Where:

ANY-NUMERIC:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

DFS optional parameters1

Parameter	Type	Default	Description
max_depth	BIGINT	\9223372036854775807\)	Upper limit of the depth of the tree.
			<ul style="list-style-type: none"><li>• When negative throws an error.</li></ul>

Inner Queries1

Edges SQL1

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source)
			<ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns1

Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Parameter	Type	Description
seq	BIGINT	Sequential value starting from \1\).
depth	BIGINT	Depth of the node.
		<ul style="list-style-type: none"><li>• \0\ when node = start_vid.</li></ul>
		<ul style="list-style-type: none"><li>• \depth-1\ is the depth of pred</li></ul>
start_vid	BIGINT	Identifier of the root vertex.
pred	BIGINT	Predecessor of node.
		<ul style="list-style-type: none"><li>• When node = start_vid then has the value node.</li></ul>

Parameter	Type	Description
node	BIGINT	Identifier of node reached using edge.
edge	BIGINT	Identifier of the edge used to arrive from pred to node. <ul style="list-style-type: none"><li><math>\backslash(-1)</math> when node = start_vid.</li></ul>
cost	FLOAT	Cost to traverse edge.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

See Also

- [Spanning Tree - Category](#)
- [Prim - Family of functions](#)
- [Sample Data](#)
- [Boost: Prim's algorithm](#)
- [Wikipedia: Prim's algorithm](#)

Indices and tables

- [Index](#)
- [Search Page](#)

**Description**

The prim algorithm was developed in 1930 by Czech mathematician Vojtěch Jarník. It is a greedy algorithm that finds a minimum spanning tree for a weighted undirected graph. This means it finds a subset of the edges that forms a tree that includes every vertex, where the total weight of all the edges in the tree is minimized. The algorithm operates by building this tree one vertex at a time, from an arbitrary starting vertex, at each step adding the cheapest possible connection from the tree to another vertex.

This algorithms find the minimum spanning forest in a possibly disconnected graph; in contrast, the most basic form of Prim's algorithm only finds minimum spanning trees in connected graphs. However, running Prim's algorithm separately for each connected component of the graph, then it is called minimum spanning forest.

The main characteristics are:

- It's implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.
- Prim's running time:  $\backslash(O(E * \log V))$

Note

From boost Graph: "The algorithm as implemented in Boost.Graph does not produce correct results on graphs with parallel edges."

Inner Queries

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

See Also

- [Spanning Tree - Category](#)
- Boost: [Prim's algorithm](#)
- Wikipedia: [Prim's algorithm](#)

Indices and tables

- [Index](#)
- [Search Page](#)



Reference¶

- [pgr\\_version](#)
- [pgr\\_full\\_version](#)

pgr\_version¶

pgr\_version — Query for pgRouting version information.

Availability

- Version 3.0.0
  - Breaking change on result columns
  - Support for old signature ends
- Version 2.0.0
  - Official function.

Description¶

Returns pgRouting version information.

Signature¶

pgr\_version()  
RETURNS TEXT

Example:

pgRouting Version for this documentation

```
SELECT pgr_version();
pgr_version
-----
3.8.0
(1 row)
```

Result columns¶

Type      Description

TEXT   pgRouting version

See Also¶

- [Reference](#)
- [pgr\\_full\\_version](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_full\_version¶

pgr\_full\_version — Get the details of pgRouting version information.

Availability

- Version 3.0.0
  - Official function.

Description¶

Get complete details of pgRouting version information

Boost Graph Inside

Signatures¶

pgr\_full\_version()  
RETURNS (version, build\_type, compile\_date, library, system, PostgreSQL, compiler, boost, hash)

Example:

Information about when this documentation was built

```
SELECT version, library FROM pgr_full_version();
version | library
-----+-----
3.8.0   | pgrouting-3.8.0
(1 row)
```

Result columns¶

Column    Type      Description

version    TEXT   pgRouting version

build\_type    TEXT   The Build type

Column	Type	Description
--------	------	-------------

compile_date	TEXT	Compilation date
--------------	------	------------------

library	TEXT	Library name and version
---------	------	--------------------------

system	TEXT	Operative system
--------	------	------------------

postgreSQL	TEXT	pgsql used
------------	------	------------

compiler	TEXT	Compiler and version
----------	------	----------------------

boost	TEXT	Boost version
-------	------	---------------

hash	TEXT	Git hash of pgRouting build
------	------	-----------------------------

See Also¶

- [Reference](#)
- [pgr\\_version](#)

Indices and tables

- [Index](#)
- [Search Page](#)

See Also¶

Indices and tables

- [Index](#)
- [Search Page](#)

## Topology - Family of Functions¶

The pgRouting's topology of a network represented with a graph in form of two tables: and edge table and a vertex table.

Attributes associated to the tables help to indicate if the graph is directed or undirected, if an edge is one way on a directed graph, and depending on the final application needs, suitable topology(s) need to be created.

The following functions modify the database directly therefore the user must have special permissions given by the administrators to use them.

- [pgr\\_createTopology - Deprecated since v3.8.0](#) - create a topology based on the geometry.
- [pgr\\_createVerticesTable - Deprecated since 3.8.0](#) - reconstruct the vertices table based on the source and target information.
- [pgr\\_analyzeGraph – Deprecated since 3.8.0](#) - to analyze the edges and vertices of the edge table.
- [pgr\\_analyzeOneWay - Deprecated since 3.8.0](#) - to analyze directionality of the edges.
- [pgr\\_nodeNetwork - Deprecated since 3.8.0](#) - to create nodes to a not noded edge table.

## Utility functions¶

- [pgr\\_extractVertices](#) - Extracts vertex information based on the edge table information.
- [pgr\\_findCloseEdges](#) - Finds close edges of points on the fly
- [pgr\\_separateCrossing](#) - Breaks geometries that cross each other.
- [pgr\\_separateTouching](#) - Breaks geometries that (almost) touch each other.

❏ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

[pgr\\_createTopology - Deprecated since v3.8.0¶](#)

[pgr\\_createTopology](#) — Builds a network topology based on the geometry information.

Availability

- Version 3.8.0
  - Deprecated function.
- Version 2.0.0

- Official function.
- Renamed from version 1.x

#### Migration of pgr\_createTopology

Starting from [v3.8.0](#)

**Before Deprecation:** The following was calculated:

- A table with `<edges>_vertices_pgr` was created.

**After Deprecation:** The user is responsible to create the complete topology.

#### Build a routing topology

The basic information to use the majority of the pgRouting functions `id, source, target, cost, [reverse_cost]` is what in pgRouting is called the routing topology.

`reverse_cost` is optional but strongly recommended to have in order to reduce the size of the database due to the size of the geometry columns. Having said that, in this documentation `reverse_cost` is used in this documentation.

When the data comes with geometries and there is no routing topology, then this step is needed.

All the start and end vertices of the geometries need an identifier that is to be stored in `source` and `target` columns of the table of the data. Likewise, `cost` and `reverse_cost` need to have the value of traversing the edge in both directions.

If the columns do not exist they need to be added to the table in question. (see [ALTER TABLE](#))

The function [pgr\\_extractVertices](#) is used to create a vertices table based on the edge identifier and the geometry of the edge of the graph.

```
SELECT * INTO vertices
FROM pgr_extractVertices('SELECT id, geom FROM edges ORDER BY id');
SELECT 18
```

Finally using the data stored on the vertices tables the `source` and `target` are filled up.

```
/* -- set the source information */
UPDATE edges AS e
SET source = v.id, x1 = x, y1 = y
FROM vertices AS v
WHERE ST_StartPoint(e.geom) = v.geom;
UPDATE 24
/* -- set the target information */
UPDATE edges AS e
SET target = v.id, x2 = x, y2 = y
FROM vertices AS v
WHERE ST_EndPoint(e.geom) = v.geom;
UPDATE 24
```

#### Description

The function returns:

- OK after the network topology has been built and the vertices table created.
- FAIL when the network topology was not built due to an error.

Boost Graph Inside

#### Signatures

`pgr_createTopology(edge_table, tolerance, [options])`  
**options:** [the\_geom, id, source, target, rows\_where, clean]  
 RETURNS VARCHAR

#### Parameters

The topology creation function accepts the following parameters:

`edge_table:`

text Network table name. (may contain the schema name as well)

`tolerance:`

float8 Snapping tolerance of disconnected edges. (in projection unit)

`the_geom:`

text Geometry column name of the network table. Default value is `the_geom`.

`id:`

text Primary key column name of the network table. Default value is `id`.

`source:`

text Source column name of the network table. Default value is `source`.

`target:`

text Target column name of the network table. Default value is `target`.

`rows_where:`

text Condition to SELECT a subset or rows. Default value is `true` to indicate all rows that where `source` or `target` have a null value, otherwise the condition is used.

`clean:`

boolean Clean any previous topology. Default value is `false`.

#### Warning

The `edge_table` will be affected

- The `source` column values will change.
- The `target` column values will change.
- An index will be created, if it doesn't exists, to speed up the process to the following columns:
  - `id`
  - `the_geom`
  - `source`

- target

The function returns:

- OK after the network topology has been built.
  - Creates a vertices table: <edge\_table>\_vertices\_pgr.
  - Fills id and the \_geom columns of the vertices table.
  - Fills the source and target columns of the edge table referencing the id of the vertices table.
- FAIL when the network topology was not built due to an error:
  - A required column of the Network table is not found or is not of the appropriate type.
  - The condition is not well formed.
  - The names of source , target or id are the same.
  - The SRID of the geometry could not be determined.

The Vertices Table

The structure of the vertices table is:

id:

bigint Identifier of the vertex.

cnt:

integer Number of vertices in the edge\_table that reference this vertex.

chk:

integer Indicator that the vertex might have a problem.

ein:

integer Number of vertices in the edge\_table that reference this vertex AS incoming.

eout:

integer Number of vertices in the edge\_table that reference this vertex AS outgoing.

the\_geom:

geometry Point geometry of the vertex.

Usage when the edge table's columns MATCH the default values:

The simplest way to use pgr\_createTopology is:

```
SELECT pgr_createTopology('edges', 0.001, 'geom');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createTopology('edges', 0.001, 'geom', 'id', 'source', 'target', rows_where := 'true', clean := f)
NOTICE: Performing checks, please wait .....
NOTICE: Creating Topology. Please wait...
NOTICE: -----> TOPOLOGY CREATED FOR 18 edges
NOTICE: Rows with NULL geometry or NULL id: 0
NOTICE: Vertices table for table public.edges is: public.edges_vertices_pgr
NOTICE: -----
pgr_createTopology
-----
OK
(1 row)
```

When the arguments are given in the order described in the parameters:

We get the same result as the simplest way to use the function.

```
SELECT pgr_createTopology('edges', 0.001,
  'geom', 'id', 'source', 'target');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createTopology('edges', 0.001, 'geom', 'id', 'source', 'target', rows_where := 'true', clean := f)
NOTICE: Performing checks, please wait .....
NOTICE: Creating Topology. Please wait...
NOTICE: -----> TOPOLOGY CREATED FOR 18 edges
NOTICE: Rows with NULL geometry or NULL id: 0
NOTICE: Vertices table for table public.edges is: public.edges_vertices_pgr
NOTICE: -----
pgr_createTopology
-----
OK
(1 row)
```

Warning

An error would occur when the arguments are not given in the appropriate order: In this example, the column id of the table edge\_table is passed to the function as the geometry column, and the geometry column the\_geom is passed to the function as the id column.

```
SELECT pgr_createTopology('edges', 0.001,
  'id', 'geom');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createTopology('edges', 0.001, 'id', 'geom', 'source', 'target', rows_where := 'true', clean := f)
NOTICE: Performing checks, please wait .....
NOTICE: -----> PGR ERROR in pgr_createTopology: Wrong type of Column id:geom
HINT: -----> Expected type of geom is integer,smallint or bigint but USER-DEFINED was found
NOTICE: Unexpected error raise_exception
pgr_createTopology
-----
FAIL
(1 row)
```

When using the named notation

Parameters defined with a default value can be omitted, as long as the value matches the default And The order of the parameters would not matter.

```
SELECT pgr_createTopology('edges', 0.001,
  the_geom:= 'geom', id:= 'id', source:= 'source', target:= 'target');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
-----
OK
(1 row)
```

```
SELECT pgr_createTopology('edges', 0.001,
    source:=source', id:=id', target:=target', the_geom:=geom');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
-----
OK
(1 row)
```

```
SELECT pgr_createTopology('edges', 0.001, 'geom', source:=source');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
-----
OK
(1 row)
```

Selecting rows using rows\_where parameter

Selecting rows based on the id.

```
SELECT pgr_createTopology('edges', 0.001, 'geom', rows_where:=id < 10');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
-----
OK
(1 row)
```

Selecting the rows where the geometry is near the geometry of row withid = 5.

```
SELECT pgr_createTopology('edges', 0.001, 'geom',
    rows_where:=geom && (SELECT st_buffer(geom, 0.05) FROM edges WHERE id=5));
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
-----
OK
(1 row)
```

Selecting the rows where the geometry is near the geometry of the row withgid =100 of the tableothertable.

```
CREATE TABLE oterTable AS (SELECT 100 AS gid, st_point(2.5, 2.5) AS other_geom);
SELECT 1
SELECT pgr_createTopology('edges', 0.001, 'geom',
    rows_where:=geom && (SELECT st_buffer(other_geom, 1) FROM oterTable WHERE gid=100));
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
-----
OK
(1 row)
```

Usage when the edge table's columns DO NOT MATCH the default values

For the following table

```
CREATE TABLE mytable AS (SELECT id AS gid, geom AS mygeom, source AS src , target AS tgt FROM edges) ;
SELECT 18
```

Using positional notation:

The arguments need to be given in the order described in the parameters.

Note that this example uses clean flag. So it recreates the whole vertices table.

```
SELECT pgr_createTopology('mytable', 0.001, 'mygeom', 'gid', 'src', 'tgt', clean := TRUE);
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
-----
OK
(1 row)
```

Warning

An error would occur when the arguments are not given in the appropriate order: In this example, the columngid of the tablemytable is passed to the function AS the geometry column, and the geometry column mygeom is passed to the function AS the id column.

```
SELECT pgr_createTopology('mytable', 0.001, 'gid', 'mygeom', 'src', 'tgt');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createTopology('mytable', 0.001, 'gid', 'mygeom', 'src', 'tgt', rows_where := 'true', clean := f)
NOTICE: Performing checks, please wait .....
NOTICE: ----> PGR ERROR in pgr_createTopology: Wrong type of Column id:mygeom
HINT: ----> Expected type of mygeom is integer,smallint or bigint but USER-DEFINED was found
NOTICE: Unexpected error raise_exception
pgr_createTopology
-----
FAIL
(1 row)
```

When using the named notation

In this scenario omitting a parameter would create an error because the default values for the column names do not match the column names of the table. The order of the parameters do not matter:

```
SELECT pgr_createTopology('mytable', 0.001, the_geom:=mygeom', id:=gid', source:=src', target:=tgt');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
-----
OK
(1 row)
```

```
SELECT pgr_createTopology('mytable', 0.001, source:=src', id:=gid', target:=tgt', the_geom:=mygeom');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
-----
OK
(1 row)
```

Selecting rows using rows\_where parameter

Based on id:

```
SELECT pgr_createTopology('mytable', 0.001, 'mygeom', 'gid', 'src', 'tgt', rows_where:=gid < 10');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
-----
OK
(1 row)
```

```
SELECT pgr_createTopology('mytable', 0.001, source:=src', id:=gid', target:=tgt', the_geom:=mygeom', rows_where:=gid < 10');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
```

Selecting the rows where the geometry is near the geometry of the row with `gid=100` of the table `othertable`.

```
SELECT pgr_createTopology('mytable', 0.001, 'mygeom', 'gid', 'src', 'tgt',
  rows_where:='mygeom && (SELECT st_buffer(other_geom, 1) FROM otherTable WHERE gid=100)');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
-----
OK
(1 row)

SELECT pgr_createTopology('mytable', 0.001, source:='src', id:='gid', target:='tgt', the_geom:='mygeom',
  rows_where:='mygeom && (SELECT st_buffer(other_geom, 1) FROM otherTable WHERE gid=100)');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
pgr_createTopology
-----
OK
(1 row)
```

### Additional Examples

- Create a routing topology
  - Make sure the database does not have the vertices table
  - Clean up the columns of the routing topology to be created
  - Create the vertices table
  - Inspect the vertices table
  - Create the routing topology on the edge table
  - Inspect the routing topology
- With full output

## Create a routing topology

An alternate method to create a routing topology use `pgr_extractVertices`

**Make sure the database does not have the vertices table**

```
DROP TABLE IF EXISTS vertices_table;
NOTICE: table "vertices_table" does not exist, skipping
DROP TABLE
```

### Clean up the columns of the routing topology to be created

```
UPDATE edges
SET source = NULL, target = NULL,
    x1 = NULL, y1 = NULL,
    x2 = NULL, y2 = NULL;
UPDATE 18
```

### Create the vertices table

- When the LINESTRING has a SRID then use `geom::geometry(POINT, <SRID>)`
- For big edge tables that are been prepared,
  - Create it as `UNLOGGED` and
  - After the table is created `ALTER TABLE .. SET LOGGED`

```
SELECT * INTO vertices_table
FROM pgr_extractVertices('SELECT id, geom FROM edges ORDER BY id');
SELECT 17
```

Inspect the vertices table

[illegible]

Create the routing topology on the edge table

WITH  
out going AS (

```
SELECT id AS vid, unnest(out_edges) AS eid, x, y
FROM vertices_table
)
UPDATE edges
SET source = vid, x1 = x, y1 = y
FROM out_going WHERE id = eid;
UPDATE 18
```


Updating the target information

```
WITH
in_coming AS (
SELECT id AS vid, unnest(in_edges) AS eid, x, y
FROM vertices_table
)
UPDATE edges
SET target = vid, x2 = x, y2 = y
FROM in_coming WHERE id = eid;
UPDATE 18
```

[Inspect the routing topology¶](#)

```
SELECT id, source, target, x1, y1, x2, y2
FROM edges ORDER BY id;
id | source | target | x1 | y1 | x2 | y2
-----+-----+-----+-----+-----+-----+-----
1 | 5 | 6 | 2 | 0 | 2 | 1
2 | 6 | 10 | 2 | 1 | 3 | 1
3 | 10 | 15 | 3 | 1 | 4 | 1
4 | 6 | 7 | 2 | 1 | 2 | 2
5 | 10 | 11 | 3 | 1 | 3 | 2
6 | 1 | 3 | 0 | 2 | 1 | 2
7 | 3 | 7 | 1 | 2 | 2 | 2
8 | 7 | 11 | 2 | 2 | 3 | 2
9 | 11 | 16 | 3 | 2 | 4 | 2
10 | 7 | 8 | 2 | 2 | 2 | 3
11 | 11 | 12 | 3 | 2 | 3 | 3
12 | 8 | 12 | 2 | 3 | 3 | 3
13 | 12 | 17 | 3 | 3 | 4 | 3
14 | 8 | 9 | 2 | 3 | 2 | 4
15 | 16 | 17 | 4 | 2 | 4 | 3
16 | 15 | 16 | 4 | 1 | 4 | 2
17 | 2 | 4 | 0.5 | 3.5 | 1.9999999999999999 | 3.5
18 | 13 | 14 | 3.5 | 2.3 | 3.5 | 4
(18 rows)
```

[\\_images/Fig1-originalData.png](#)



Generated topology¶

[With full output¶](#)

This example start a clean topology, with 5 edges, and then its incremented to the rest of the edges.

```
SELECT pgr_createTopology('edges', 0.001, 'geom', rows_where:='id < 6', clean := true);
WARNING: pgr_createTopology(text,double precision,text,text,text,text,text,boolean) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createTopology('edges', 0.001, 'geom', 'id', 'source', 'target', rows_where := 'id < 6', clean := t)
NOTICE: Performing checks, please wait .....
NOTICE: Creating Topology, Please wait...
NOTICE: -----> TOPOLOGY CREATED FOR 5 edges
NOTICE: Rows with NULL geometry or NULL id: 0
NOTICE: Vertices table for table public.edges is: public.edges_vertices_pgr
NOTICE: -----
pgr_createTopology
-----
OK
(1 row)

SELECT pgr_createTopology('edges', 0.001, 'geom');
WARNING: pgr_createTopology(text,double precision,text,text,text,text,text,boolean) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createTopology('edges', 0.001, 'geom', 'id', 'source', 'target', rows_where := 'true', clean := f)
NOTICE: Performing checks, please wait .....
NOTICE: Creating Topology, Please wait...
NOTICE: -----> TOPOLOGY CREATED FOR 13 edges
NOTICE: Rows with NULL geometry or NULL id: 0
NOTICE: Vertices table for table public.edges is: public.edges_vertices_pgr
NOTICE: -----
pgr_createTopology
-----
OK
(1 row)
```

See Also¶

- [Sample Data](#)
- [Topology - Family of Functions](#)

Indices and tables

- [Index](#)
- [Search Page](#)

`pgr_createVerticesTable` - [Deprecated since 3.8.0](#)

`pgr_createVerticesTable` — Reconstructs the vertices table based on the source and target information.

Availability

- Version 3.8.0
  - Deprecated function.
- Version 2.0.0
  - Official function.
  - Renamed from version 1.x

Migration of `pgr_createVerticesTable`

Starting from [v3.8.0](#)

**Before Deprecation:** The following was calculated:

- A table with `<edges>_vertices_pgr` was created.

**After Deprecation:** The user is responsible to create the vertices table, indexes, etc. They may use [pgr\\_extractVertices](#) for that purpose.

```
SELECT * INTO vertices
FROM pgr_extractVertices('SELECT id, geom FROM edges ORDER BY id');
SELECT 17
```

Description

The function returns:

- OK after the vertices table has been reconstructed.
- FAIL when the vertices table was not reconstructed due to an error.

 Boost Graph Inside

Signatures

`pgr_createVerticesTable(edge_table, [the_geom, source, target, rows_where])`  
RETURNS VARCHAR

Parameters

The reconstruction of the vertices table function accepts the following parameters:

`edge_table`:

text Network table name. (may contain the schema name as well)

`the_geom`:

text Geometry column name of the network table. Default value is `the_geom`.

`source`:

text Source column name of the network table. Default value is `source`.

`target`:

text Target column name of the network table. Default value is `target`.

`rows_where`:

text Condition to SELECT a subset or rows. Default value is `true` to indicate all rows.

Warning

The `edge_table` will be affected

- An index will be created, if it doesn't exists, to speed up the process to the following columns:
  - `the_geom`
  - `source`
  - `target`

The function returns:

- OK after the vertices table has been reconstructed.
  - Creates a vertices table: `<edge_table>_vertices_pgr`.
  - Fills `id` and `the_geom` columns of the vertices table based on the source and target columns of the edge table.
- FAIL when the vertices table was not reconstructed due to an error.
  - A required column of the Network table is not found or is not of the appropriate type.
  - The condition is not well formed.
  - The names of source, target are the same.
  - The SRID of the geometry could not be determined.

The Vertices Table

The structure of the vertices table is:

`id`:

bigint Identifier of the vertex.

`cnt`:

integer Number of vertices in the `edge_table` that reference this vertex.

`chk`:

integer Indicator that the vertex might have a problem.

`ein`:



integer Number of vertices in the edge\_table that reference this vertex as incoming.

cout:

integer Number of vertices in the edge\_table that reference this vertex as outgoing.

the\_geom:

geometry Point geometry of the vertex.

Example 1:

The simplest way to use pgr\_createVerticesTable

```
SELECT pgr_createVerticesTable('edges', 'geom');
WARNING: pgr_createvertexestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('edges','geom','source','target','true')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.edges_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 17 VERTICES
NOTICE:                FOR 18 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 18
NOTICE: Vertices table for table public.edges is: public.edges_vertices_pgr
NOTICE: -----
pgr_createvertexestable
-----
OK
(1 row)
```

Additional Examples1

Example 2:

When the arguments are given in the order described in the parameters:

```
SELECT pgr_createVerticesTable('edges', 'geom', 'source', 'target');
WARNING: pgr_createvertexestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('edges','geom','source','target','true')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.edges_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 17 VERTICES
NOTICE:                FOR 18 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 18
NOTICE: Vertices table for table public.edges is: public.edges_vertices_pgr
NOTICE: -----
pgr_createvertexestable
-----
OK
(1 row)
```

We get the same result as the simplest way to use the function.

Warning

An error would occur when the arguments are not given in the appropriate order: In this example, the column source columnsource of the table mytable is passed to the function as the geometry column, and the geometry column the\_geom is passed to the function as the source column.

```
SELECT pgr_createVerticesTable('edges', 'source', 'geom', 'target');
WARNING: pgr_createvertexestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('edges','source','geom','target','true')
NOTICE: Performing checks, please wait .....
NOTICE: -----> PGR ERROR in pgr_createVerticesTable: Wrong type of Column source: geom
HINT: -----> Expected type of geom is integer, smallint or bigint but USER-DEFINED was found
NOTICE: Unexpected error raise_exception
pgr_createvertexestable
-----
FAIL
(1 row)
```

When using the named notation

Example 3:

The order of the parameters do not matter:

```
SELECT pgr_createVerticesTable('edges', the_geom:='geom', source:='source', target:='target');
WARNING: pgr_createvertexestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('edges','geom','source','target','true')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.edges_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 17 VERTICES
NOTICE:                FOR 18 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 18
NOTICE: Vertices table for table public.edges is: public.edges_vertices_pgr
NOTICE: -----
pgr_createvertexestable
-----
OK
(1 row)
```

Example 4:

Using a different ordering

```
SELECT pgr_createVerticesTable('edges', source:='source', target:='target', the_geom:='geom');
WARNING: pgr_createvertexestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('edges','geom','source','target','true')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.edges_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 17 VERTICES
NOTICE:                FOR 18 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 18
NOTICE: Vertices table for table public.edges is: public.edges_vertices_pgr
NOTICE: -----
pgr_createvertexestable
-----
OK
(1 row)
```

Example 5:

Parameters defined with a default value can be omitted, as long as the value matches the default:

```

SELECT pgr_createVerticesTable('edges', 'geom', source:='source');
WARNING: pgr_createverticestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('edges','geom','source','target','true')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.edges_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 17 VERTICES
NOTICE: FOR 18 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 18
NOTICE: Vertices table for table public.edges is: public.edges_vertices_pgr
NOTICE: -----
pgr_createverticestable
-----
OK
(1 row)

```

Selecting rows using rows\_where parameter

Example 6:

Selecting rows based on the id.

```

SELECT pgr_createVerticesTable('edges', 'geom', rows_where:='id < 10');
WARNING: pgr_createverticestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('edges','geom','source','target','id < 10')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.edges_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 9 VERTICES
NOTICE: FOR 10 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 10
NOTICE: Vertices table for table public.edges is: public.edges_vertices_pgr
NOTICE: -----
pgr_createverticestable
-----
OK
(1 row)

```

Example 7:

Selecting the rows where the geometry is near the geometry of row withid =5 .

```

SELECT pgr_createVerticesTable('edges', 'geom',
rows_where:='geom && (select st_buffer(geom,0.5) FROM edges WHERE id=5)');
WARNING: pgr_createverticestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('edges','geom','source','target','geom && (select st_buffer(geom,0.5) FROM edges WHERE id=5)')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.edges_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 9 VERTICES
NOTICE: FOR 9 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 9
NOTICE: Vertices table for table public.edges is: public.edges_vertices_pgr
NOTICE: -----
pgr_createverticestable
-----
OK
(1 row)

```

Example 8:

Selecting the rows where the geometry is near the geometry of the row withgid =100 of the table othertable.

```

DROP TABLE IF EXISTS otherTable;
NOTICE: table "othertable" does not exist, skipping
DROP TABLE
CREATE TABLE otherTable AS (SELECT 100 AS gid, st_point(2.5,2.5) AS other_geom) ;
SELECT 1
SELECT pgr_createVerticesTable('edges', 'geom',
rows_where:='geom && (select st_buffer(other_geom,0.5) FROM otherTable WHERE gid=100)');
WARNING: pgr_createverticestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('edges','geom','source','target','geom && (select st_buffer(other_geom,0.5) FROM otherTable WHERE gid=100)')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.edges_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 10 VERTICES
NOTICE: FOR 12 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 12
NOTICE: Vertices table for table public.edges is: public.edges_vertices_pgr
NOTICE: -----
pgr_createverticestable
-----
OK
(1 row)

```

Usage when the edge table's columns DO NOT MATCH the default values:

Using the following table

```

DROP TABLE IF EXISTS mytable;
NOTICE: table "mytable" does not exist, skipping
DROP TABLE
CREATE TABLE mytable AS (SELECT id AS gid, geom AS mygeom, source AS src ,target AS tgt FROM edges) ;
SELECT 18

```

Using positional notation:

Example 9:

The arguments need to be given in the order described in the parameters:

```

SELECT pgr_createVerticesTable('mytable', 'mygeom', 'src', 'tgt');
WARNING: pgr_createverticestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('mytable','mygeom','src','tgt','true')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.mytable_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 17 VERTICES
NOTICE: FOR 18 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 18
NOTICE: Vertices table for table public.mytable is: public.mytable_vertices_pgr
NOTICE: -----
pgr_createverticestable
-----
OK
(1 row)

```

## Warning

An error would occur when the arguments are not given in the appropriate order: In this example, the column `src` of the table `mytable` is passed to the function as the geometry column, and the geometry column `mygeom` is passed to the function as the source column.

```
SELECT pgr_createVerticesTable('mytable', 'src', 'mygeom', 'tgt');
WARNING: pgr_createverticesTable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('mytable','src','mygeom','tgt','true')
NOTICE: Performing checks, please wait .....
NOTICE: -----> PGR ERROR in pgr_createVerticesTable: Wrong type of Column source: mygeom
HINT: -----> Expected type of mygeom is integer, smallint or bigint but USER-DEFINED was found
NOTICE: Unexpected error raise_exception
pgr_createverticesTable
-----
FAIL
(1 row)
```

When using the named notation

Example 10:

The order of the parameters do not matter:

```
SELECT pgr_createVerticesTable('mytable',the_geom:='mygeom',source:='src',target:='tgt');
WARNING: pgr_createverticesTable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('mytable','mygeom','src','tgt','true')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.mytable_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 17 VERTICES
NOTICE: FOR 18 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 18
NOTICE: Vertices table for table public.mytable is: public.mytable_vertices_pgr
NOTICE: -----
pgr_createverticesTable
-----
OK
(1 row)
```

Example 11:

Using a different ordering

In this scenario omitting a parameter would create an error because the default values for the column names do not match the column names of the table.

```
SELECT pgr_createVerticesTable(
  'mytable', source:='src', target:='tgt',
  the_geom:='mygeom');
WARNING: pgr_createverticesTable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('mytable','mygeom','src','tgt','true')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.mytable_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 17 VERTICES
NOTICE: FOR 18 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 18
NOTICE: Vertices table for table public.mytable is: public.mytable_vertices_pgr
NOTICE: -----
pgr_createverticesTable
-----
OK
(1 row)
```

Selecting rows using `rows_where` parameter

Example 12:

Selecting rows based on the gid. (positional notation)

```
SELECT pgr_createVerticesTable(
  'mytable', 'mygeom', 'src', 'tgt',
  rows_where:='gid < 10');
WARNING: pgr_createverticesTable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('mytable','mygeom','src','tgt','gid < 10')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.mytable_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 9 VERTICES
NOTICE: FOR 10 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 10
NOTICE: Vertices table for table public.mytable is: public.mytable_vertices_pgr
NOTICE: -----
pgr_createverticesTable
-----
OK
(1 row)
```

Example 13:

Selecting rows based on the gid. (named notation)

```
SELECT pgr_createVerticesTable(
  'mytable', source:='src', target:='tgt', the_geom:='mygeom',
  rows_where:='gid < 10');
WARNING: pgr_createverticesTable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('mytable','mygeom','src','tgt','gid < 10')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.mytable_vertices_pgr, please wait...
NOTICE: -----> VERTICES TABLE CREATED WITH 9 VERTICES
NOTICE: FOR 10 EDGES
NOTICE: Edges with NULL geometry,source or target: 0
NOTICE: Edges processed: 10
NOTICE: Vertices table for table public.mytable is: public.mytable_vertices_pgr
NOTICE: -----
pgr_createverticesTable
-----
OK
(1 row)
```

Example 14:

Selecting the rows where the geometry is near the geometry of row with `gid = 5`.

```
SELECT pgr_createVerticesTable(
  'mytable', 'mygeom', 'src', 'tgt',
  rows_where := 'the_geom && (SELECT st_buffer(mygeom,0.5) FROM mytable WHERE gid=5)');
WARNING: pgr_createverticesTable(text,text,text,text,text) deprecated function on v3.8.0
```

```
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('mytable','mygeom','src','tgt','the_geom' && (SELECT st_buffer(mygeom,0.5) FROM mytable WHERE gid=5))
NOTICE: Performing checks, please wait .....
NOTICE: Got column "the_geom" does not exist
NOTICE: ERROR: Condition is not correct, please execute the following query to test your condition
NOTICE: select * from public.mytable WHERE true AND (the_geom && (SELECT st_buffer(mygeom,0.5) FROM mytable WHERE gid=5)) limit 1
pgr_createverticestable
-----
FAIL
(1 row)
```

Example 15:

```
TBD

SELECT pgr_createVerticesTable(
    'mytable', source:='src', target:='tgt', the_geom:='mygeom',
    rows_where:='mygeom && (SELECT st_buffer(mygeom,0.5) FROM mytable WHERE id=5)');
WARNING: pgr_createverticestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('mytable','mygeom','src','tgt','mygeom' && (SELECT st_buffer(mygeom,0.5) FROM mytable WHERE id=5))
NOTICE: Performing checks, please wait .....
NOTICE: Got column "id" does not exist
NOTICE: ERROR: Condition is not correct, please execute the following query to test your condition
NOTICE: select * from public.mytable WHERE true AND (mygeom && (SELECT st_buffer(mygeom,0.5) FROM mytable WHERE id=5)) limit 1
pgr_createverticestable
-----
FAIL
(1 row)
```

Example 16:

```

    Selecting the rows where the geometry is near the geometry of the row withgid =100 of the table othertable.

DROP TABLE IF EXISTS otherTable;
DROP TABLE
CREATE TABLE otherTable AS (SELECT 100 AS gid, st_point(2.5,2.5) AS other_geom) ;
SELECT 1

SELECT pgr_createVerticesTable(
    'mytable','mygeom','src','tgt',
    rows_where:='the_geom && (SELECT st_buffer(othergeom,0.5) FROM otherTable WHERE gid=100)');
WARNING: pgr_createverticestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('mytable','mygeom','src','tgt','the_geom' && (SELECT st_buffer(othergeom,0.5) FROM otherTable WHERE gid=100))
NOTICE: Performing checks, please wait .....
NOTICE: Got column "the_geom" does not exist
NOTICE: ERROR: Condition is not correct, please execute the following query to test your condition
NOTICE: select * from public.mytable WHERE true AND (the_geom && (SELECT st_buffer(othergeom,0.5) FROM otherTable WHERE gid=100)) limit 1
pgr_createverticestable
-----
FAIL
(1 row)
```

Example 17:

```
TBD

SELECT pgr_createVerticesTable(
    'mytable', source:='src', target:='tgt', the_geom:='mygeom',
    rows_where:='the_geom && (SELECT st_buffer(othergeom,0.5) FROM otherTable WHERE gid=100)');
WARNING: pgr_createverticestable(text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('mytable','mygeom','src','tgt','the_geom' && (SELECT st_buffer(othergeom,0.5) FROM otherTable WHERE gid=100))
NOTICE: Performing checks, please wait .....
NOTICE: Got column "the_geom" does not exist
NOTICE: ERROR: Condition is not correct, please execute the following query to test your condition
NOTICE: select * from public.mytable WHERE true AND (the_geom && (SELECT st_buffer(othergeom,0.5) FROM otherTable WHERE gid=100)) limit 1
pgr_createverticestable
-----
FAIL
(1 row)
```

See Also

- [Sample Data](#)
- [Topology - Family of Functions](#) for an overview of a topology for routing algorithms.

Indices and tables

- [Index](#)
- [Search Page](#)

**pgr\_analyzeGraph** – **Deprecated since 3.8.0**

pgr\_analyzeGraph — Analyzes the network topology.

Availability

- Version 3.8.0
  - Deprecated function.
- Version 2.0.0
  - Official function.

**Migration of pgr\_analyzeGraph**

Starting from [v3.8.0](#)

**Before Deprecation:** The following was calculated:

- Number of isolated segments.
- Number of dead ends.
- Number of potential gaps found near dead ends.
- Number of intersections. (between 2 edges)

WHERE

Graph component:

A connected subgraph that is not part of any larger connected subgraph.

Isolated segment:

A graph component with only one segment.

Dead ends:

A vertex that participates in only one edge.

gaps:

Space between two geometries.

Intersection:

Is a topological relationship between two geometries.

Migration.

Components.

Instead of counting only isolated segments, determine all the components of the graph.

Depending of the final application requirements use:

- [pgr\\_connectedComponents](#)
- [pgr\\_strongComponents](#)
- [pgr\\_biconnectedComponents](#)

For example:

```
SELECT *
FROM pgr_connectedComponents(
  'SELECT id, source, target, cost, reverse_cost FROM edges'
);
```

seq | component | node

1	1	1
2	1	3
3	1	5
4	1	6
5	1	7
6	1	8
7	1	9
8	1	10
9	1	11
10	1	12
11	1	15
12	1	16
13	1	17
14	2	2
15	2	4
16	13	13
17	13	14

(17 rows)

Dead ends.

Instead of counting the dead ends, determine all the dead ends of the graph using [pgr\\_degree](#).

For example:

```
SELECT *
FROM pgr_degree(($$SELECT id, source, target FROM edges$$)
WHERE degree = 1;
```

node | degree

9	1
5	1
4	1
14	1
13	1
2	1
1	1

(7 rows)

Potential gaps near dead ends.

Instead of counting potential gaps between geometries, determine the geometric gaps in the graph using [pgr\\_findCloseEdges](#).

For example:

```
WITH
deadends AS (
  SELECT id,geom, (in_edges || out_edges)[1] as inhere
  FROM vertices where array_length(in_edges || out_edges, 1) = 1),
results AS (
  SELECT (pgr_findCloseEdges('SELECT id, geom FROM edges WHERE id != ' || inhere , geom, 0.001)).*
  FROM deadends)
SELECT d.id, edge_id, distance, st_AsText(geom) AS point, st_AsText(edge) edge
FROM results JOIN deadends d USING (geom);
```

id	edge_id	distance	point	edge
4	14	1.00008890058e-12	POINT(1.999999999999 3.5)	LINESTRING(1.999999999999 3.5,2 3.5)

(1 row)

Topological relationships.

Instead of counting intersections, determine topological relationships between geometries.

Several PostGIS functions can be used: [ST\\_Intersects](#), [ST\\_Crosses](#), [ST\\_Overlaps](#), etc.

For example:

```
SELECT e1.id AS id1, e2.id AS id2
FROM edges e1, edges e2 WHERE e1 < e2 AND st_crosses(e1.geom, e2.geom);
```

id1	id2
13	18

(1 row)

Description

The function returns:

- OK after the analysis has finished.

- FAIL when the analysis was not completed due to an error.

pgr\_analyzeGraph(**edge\_table**,**tolerance**,**[options]**)  
**options:** [the\_geom, id, source, target, rows\_where]  
RETURNS VARCHAR

Prerequisites

The edge table to be analyzed must contain a source column and a target column filled with id's of the vertices of the segments and the corresponding vertices table <edge\_table>\_vertices\_pgr that stores the vertices information.

Parameters

The analyze graph function accepts the following parameters:

edge\_table:

text Network table name. (may contain the schema name as well)

tolerance:

float8 Snapping tolerance of disconnected edges. (in projection unit)

the\_geom:

text Geometry column name of the network table. Default value isthe\_geom.

id:

text Primary key column name of the network table. Default value isid.

source:

text Source column name of the network table. Default value issource.

target:

text Target column name of the network table. Default value istarget.

rows\_where:

text Condition to select a subset or rows. Default value istrue to indicate all rows.

The function returns:

- OK after the analysis has finished.
  - Uses the vertices table: <edge\_table>\_vertices\_pgr.
  - Fills completely the cnt and chk columns of the vertices table.
  - Returns the analysis of the section of the network defined byrows\_where
- FAIL when the analysis was not completed due to an error.
  - The vertices table is not found.
  - A required column of the Network table is not found or is not of the appropriate type.
  - The condition is not well formed.
  - The names of source , target or id are the same.
  - The SRID of the geometry could not be determined.

The Vertices Table

The structure of the vertices table is:

id:

bigint Identifier of the vertex.

cnt:

integer Number of vertices in the edge\_table that reference this vertex.

chk:

integer Indicator that the vertex might have a problem.

ein:

integer Number of vertices in the edge\_table that reference this vertex as incoming.

eout:

integer Number of vertices in the edge\_table that reference this vertex as outgoing.

the\_geom:

geometry Point geometry of the vertex.

Usage when the edge table's columns MATCH the default values

The simplest way to use pgr\_analyzeGraph is:

```
SELECT id, geom AS the_geom, 0 AS cnt, 0 AS chk INTO edges_vertices_pgr FROM vertices;
SELECT 17
SELECT pgr_analyzeGraph('edges',0.001,'geom');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edges',0.001,'geom','id','source','target','true')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 7
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

Arguments are given in the order described in the parameters:

```
SELECT pgr_analyzeGraph('edges',0.001,'geom','id','source','target');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph("edges",0.001,'geom','id','source','target','true')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 7
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

We get the same result as the simplest way to use the function.

Warning

An error would occur when

the arguments are not given in the appropriate order:

In this example, the column `id` of the table `mytable` is passed to the function as the geometry column, and the geometry column `the_geom` is passed to the function as the `id` column.

```
SELECT pgr_analyzeGraph('edges',0.001,'id','geom','source','target');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph("edges",0.001,'id','geom','source','target','true')
NOTICE: Performing checks, please wait ...
NOTICE: Got function st_srid(bigint) does not exist
NOTICE: ERROR: something went wrong when checking for SRID of id in table public.edges
pgr_analyzegraph
-----
FAIL
(1 row)
```

When using the named notation

The order of the parameters do not matter:

```
SELECT pgr_analyzeGraph('edges',0.001,the_geom:='geom',id:='id',source:='source',target:='target');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph("edges",0.001,'geom','id','source','target','true')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 7
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

```
SELECT pgr_analyzeGraph('edges',0.001,source:='source',id:='id',target:='target',the_geom:='geom');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph("edges",0.001,'geom','id','source','target','true')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 7
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

Parameters defined with a default value can be omitted, as long as the value matches the default:

```
SELECT pgr_analyzeGraph('edges',0.001, 'geom', source:='source');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph("edges",0.001,'geom','id','source','target','true')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 7
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

Selecting rows using `rows_where` parameter

Selecting rows based on the `id`. Displays the analysis a the section of the network.

```
SELECT pgr_analyzeGraph('edges',0.001, 'geom', rows_where:='id < 10');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph("edges",0.001,'geom','id','source','target','id < 10')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
```

```
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 0
NOTICE: Dead ends: 4
NOTICE: Potential gaps found near dead ends: 0
NOTICE: Intersections detected: 0
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

Selecting the rows where the geometry is near the geometry of row withid = 5

```
SELECT pgr_analyzeGraph('edges',0.001, 'geom', rows_where:='geom && (SELECT st_buffer(geom,0.05) FROM edge_table WHERE id=5)');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edges',0.001,'geom','id','source','target','geom && (SELECT st_buffer(geom,0.05) FROM edge_table WHERE id=5)')
NOTICE: Performing checks, please wait ...
NOTICE: Got relation "edge_table" does not exist
NOTICE: ERROR: Condition is not correct. Please execute the following query to test your condition
NOTICE: select count(*) from public.edges WHERE true AND (geom && (SELECT st_buffer(geom,0.05) FROM edge_table WHERE id=5))
pgr_analyzegraph
-----
FAIL
(1 row)
```

Selecting the rows where the geometry is near the geometry of the row withgid =100 of the table othertable.

```
CREATE TABLE otherTable AS (SELECT 100 AS gid, st_point(2.5,2.5) AS other_geom) ;
SELECT 1
SELECT pgr_analyzeGraph('edges',0.001, 'geom', rows_where:='geom && (SELECT st_buffer(geom,1) FROM otherTable WHERE gid=100)');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edges',0.001,'geom','id','source','target','geom && (SELECT st_buffer(geom,1) FROM otherTable WHERE gid=100)')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 7
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

Usage when the edge table's columns DO NOT MATCH the default values:

For the following table

```
CREATE TABLE mytable AS (SELECT id AS gid, source AS src , target AS tgt , geom AS mygeom FROM edges);
SELECT 18
SELECT id, geom AS the_geom, 0 AS cnt, 0 AS chk INTO mytable_vertices_pgr FROM vertices;
SELECT 17
```

Using positional notation:

The arguments need to be given in the order described in the parameters:

```
SELECT pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt','true')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 7
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

Warning

An error would occur when the arguments are not given in the appropriate order: In this example, the columngid of the table mytable is passed to the function as the geometry column, and the geometry column mygeom is passed to the function as the id column.

```
SELECT pgr_analyzeGraph('mytable',0.0001,'gid','mygeom','src','tgt');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('mytable',0.0001,'gid','mygeom','src','tgt','true')
NOTICE: Performing checks, please wait ...
NOTICE: Got function st_srid(bigint) does not exist
NOTICE: ERROR: something went wrong when checking for SRID of gid in table public.mytable
pgr_analyzegraph
-----
FAIL
(1 row)
```

When using the named notation

The order of the parameters do not matter:

```
SELECT pgr_analyzeGraph('mytable',0.001,the_geom:='mygeom',id:='gid',source:='src',target:='tgt');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt','true')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 7
```



```
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)

SELECT pgr_analyzeGraph('mytable',0.001,source:='src',id:='gid',target:='tgt',the_geom:='mygeom');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt',true)
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 7
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

In this scenario omitting a parameter would create an error because the default values for the column names do not match the column names of the table.

Selecting rows using rows\_where parameter

Selecting rows based on the id.

```
SELECT pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt',rows_where:='gid < 10');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt','gid < 10')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 0
NOTICE: Dead ends: 4
NOTICE: Potential gaps found near dead ends: 0
NOTICE: Intersections detected: 0
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)

SELECT pgr_analyzeGraph('mytable',0.001,source:='src',id:='gid',target:='tgt',the_geom:='mygeom',rows_where:='gid < 10');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt','gid < 10')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 0
NOTICE: Dead ends: 4
NOTICE: Potential gaps found near dead ends: 0
NOTICE: Intersections detected: 0
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

Selecting the rows WHERE the geometry is near the geometry of row withid =5 .

```
SELECT pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt',
rows_where:='mygeom && (SELECT st_buffer(mygeom,1) FROM mytable WHERE gid=5)');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt','mygeom && (SELECT st_buffer(mygeom,1) FROM mytable WHERE gid=5)')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 1
NOTICE: Dead ends: 5
NOTICE: Potential gaps found near dead ends: 0
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)

SELECT pgr_analyzeGraph('mytable',0.001,source:='src',id:='gid',target:='tgt',the_geom:='mygeom',
rows_where:='mygeom && (SELECT st_buffer(mygeom,1) FROM mytable WHERE gid=5)');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt','mygeom && (SELECT st_buffer(mygeom,1) FROM mytable WHERE gid=5)')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 1
NOTICE: Dead ends: 5
NOTICE: Potential gaps found near dead ends: 0
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

Selecting the rows WHERE the geometry is near the place='myhouse' of the tableothertable. (note the use of quote\_literal)

```
DROP TABLE IF EXISTS otherTable;
DROP TABLE
CREATE TABLE otherTable AS (SELECT 'myhouse':text AS place, st_point(2.5,2.5) AS other_geom) ;
SELECT 1
SELECT pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt',
  rows_where='mygeom && (SELECT st_buffer(other_geom,1) FROM otherTable WHERE place=||quote_literal('myhouse'))||');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt','mygeom && (SELECT st_buffer(other_geom,1) FROM otherTable WHERE place='myhouse')')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 10
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

```
SELECT pgr_analyzeGraph('mytable',0.001,source:='src',id:='gid',target:='tgt',the_geom:='mygeom',
  rows_where='mygeom && (SELECT st_buffer(other_geom,1) FROM otherTable WHERE place=||quote_literal('myhouse'))||');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt','mygeom && (SELECT st_buffer(other_geom,1) FROM otherTable WHERE place='myhouse')')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 10
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

#### Additional Examples

```
UPDATE edges_vertices_pgr SET (cnt,chk) = (0,0);
UPDATE 17
SELECT pgr_analyzeGraph('edges', 0.001, 'geom');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edges',0.001,'geom','id','source','target','true')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 7
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

```
SELECT pgr_analyzeGraph('edges',0.001,'geom', rows_where='id < 10');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edges',0.001,'geom','id','source','target','id < 10')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 0
NOTICE: Dead ends: 4
NOTICE: Potential gaps found near dead ends: 0
NOTICE: Intersections detected: 0
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

```
SELECT pgr_analyzeGraph('edges',0.001,'geom', rows_where='id >= 10');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edges',0.001,'geom','id','source','target','id >= 10')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 2
NOTICE: Dead ends: 8
NOTICE: Potential gaps found near dead ends: 1
NOTICE: Intersections detected: 1
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

```
SELECT pgr_analyzeGraph('edges',0.001,'geom', rows_where='id < 17');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edges',0.001,'geom','id','source','target','id < 17')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
```

```
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 0
NOTICE: Dead ends: 3
NOTICE: Potential gaps found near dead ends: 0
NOTICE: Intersections detected: 0
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)

DELETE FROM edges WHERE id >= 17;
DELETE 2
UPDATE edges_vertices_pgr SET (cnt,chk) = (0,0);
UPDATE 17

SELECT pgr_analyzeGraph('edges', 0.001, 'geom');
WARNING: pgr_analyzegraph(text,double precision,text,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edges',0.001,'geom','id','source','target','true')
NOTICE: Performing checks, please wait ...
NOTICE: Analyzing for dead ends. Please wait...
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 0
NOTICE: Dead ends: 3
NOTICE: Potential gaps found near dead ends: 0
NOTICE: Intersections detected: 0
NOTICE: Ring geometries: 0
pgr_analyzegraph
-----
OK
(1 row)
```

See Also

- [Sample Data](#)
- [Topology - Family of Functions](#)
- [pgr\\_analyzeOneWay - Deprecated since 3.8.0](#)

Indices and tables

- [Index](#)
- [Search Page](#)

**pgr\_analyzeOneWay - Deprecated since 3.8.0**

pgr\_analyzeOneWay — Analyzes oneway Sstreets and identifies flipped segments.

This function analyzes oneway streets in a graph and identifies any flipped segments.

Availability

- Version 3.8.0
  - Deprecated function.
- Version 2.0.0
  - Official function.

**Migration of pgr\_analyzeOneWay**

Starting from [v3.8.0](#)

**Before Deprecation:** The following was calculated:

- Number of potential problems in directionality

WHERE

Directionality problems were calculated based on codes.

Dead ends.

A routing problem can arise when from a vertex there is only a way on or a way out but not both:

Either saving or using directly [pgr\\_extractVertices](#) get the dead ends information and determine if the adjacent edge is one way or not.

In this example [pgr\\_extractVertices](#) has already been applied.

```
WITH
deadends AS (
  SELECT (in_edges || out_edges)[1] as id
  FROM vertices where array_length(in_edges || out_edges, 1) = 1)
SELECT * FROM edges JOIN deadends USING (id)
WHERE cost < 0 OR reverse_cost < 0;
id | source | target | cost | reverse_cost | capacity | reverse_capacity | x1 | y1 | x2 | y2 | geom
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
(0 rows)
```

Bridges.

Another routing problem can arise when there is an edge of an undirected graph whose deletion increases its number of connected components, and the bridge is only one way.

To determine if the bridges are or not one way.

```
SELECT id, cost < 0 OR reverse_cost<0 AS is_OneWay
FROM pgr_bridges('SELECT id, source, target, cost, reverse_cost FROM edges')
JOIN edges ON (edge = id);
id | is_oneway
-----+-----
1 | f
6 | f
7 | f
14 | f
17 | f
18 | f
(6 rows)
```

Description

The analyses of one way segments is pretty simple but can be a powerful tools to identifying some the potential problems created by setting the direction of a segment the wrong way. A node is a *source* if it has edges the exit from that node and no edges enter that node. Conversely, a node is *asink* if all edges enter the node but none exit that node. For *asource* type node it is logically impossible to exist because no vehicle can exit the node if no vehicle and enter the node. Likewise, if you had a *sink* node you would have an infinite number of vehicle piling up on this node because you can enter it but not leave it.

So why do we care if the are not feasible? Well if the direction of an edge was reversed by mistake we could generate exactly these conditions. Think about a divided highway and on the north bound lane one segment got entered wrong or maybe a sequence of multiple segments got entered wrong or maybe this happened on a round-about. The result would be potentially a *source* and/or a *sink* node.

So by counting the number of edges entering and exiting each node we can identify both *source* and *sink* nodes so that you can look at those areas of your network to make repairs and/or report the problem back to your data vendor.

Prerequisites

The edge table to be analyzed must contain a source column and a target column filled with id's of the vertices of the segments and the corresponding vertices table <edge\_table>\_vertices\_pgr that stores the vertices information.

Boost Graph Inside

Signatures

pgr\_analyzeOneWay(**geom\_table**, **s\_in\_rules**, **s\_out\_rules**, **t\_in\_rules**, **t\_out\_rules**, [options])  
**options:** [oneway, source, target, two\_way\_if\_null]  
RETURNS TEXT

Parameters

edge\_table:

text Network table name. (may contain the schema name as well)

s\_in\_rules:

text[] source node **in** rules

s\_out\_rules:

text[] source node **out** rules

t\_in\_rules:

text[] target node **in** rules

t\_out\_rules:

text[] target node **out** rules

oneway:

text oneway column name name of the network table. Default value is oneway.

source:

text Source column name of the network table. Default value is source.

target:

text Target column name of the network table. Default value is target.

two\_way\_if\_null:

boolean flag to treat oneway NULL values as bi-directional. Default value is true.

The function returns:

- OK after the analysis has finished.
  - Uses the vertices table: <edge\_table>\_vertices\_pgr.
  - Fills completely the ein and eout columns of the vertices table.
- FAIL when the analysis was not completed due to an error.
  - The vertices table is not found.
  - A required column of the Network table is not found or is not of the appropriate type.
  - The names of source , target or oneway are the same.

The rules are defined as an array of text strings that if match the oneway value would be counted as true for the source or target **in** or **out** condition.

The Vertices Table

The structure of the vertices table is:

id:

bigint Identifier of the vertex.

cnt:

integer Number of vertices in the edge\_table that reference this vertex.

chk:

integer Indicator that the vertex might have a problem.

ein:

integer Number of vertices in the edge\_table that reference this vertex as incoming.

eout:

integer Number of vertices in the edge\_table that reference this vertex as outgoing.

the\_geom:

geometry Point geometry of the vertex.

Additional Examples

```
ALTER TABLE edges ADD COLUMN dir TEXT;
ALTER TABLE
SELECT *, NULL::INTEGER ein, NULL::INTEGER eout INTO edges_vertices_pgr FROM vertices;
SELECT 17
UPDATE edges SET
```

```
dir = CASE WHEN (cost>0 AND reverse_cost>0) THEN 'B' /* both ways */
        WHEN (cost>0 AND reverse_cost<0) THEN 'FT' /* direction of the LINESSTRING */
        WHEN (cost<0 AND reverse_cost>0) THEN 'TF' /* reverse direction of the LINESSTRING */
        ELSE '' END;
UPDATE 18
/* unknown */
SELECT pgr_analyzeOneWay('edges',
    ARRAY['', 'B', 'TF'],
    ARRAY['', 'B', 'FT'],
    ARRAY['', 'B', 'FT'],
    ARRAY['', 'B', 'TF'],
    oneway:= 'dir');
WARNING: pgr_analyzeoneway(text,text[],text[],text[],text[],boolean,text,text,text) deprecated function on v3.8.0
NOTICE: PROCESSING.
NOTICE: pgr_analyzeOneway('edges',{'','B,TF'},{'','B,FT'},{'','B,FT'},{'','B,TF'},'dir','source','target',t)
NOTICE: Analyzing graph for one way street errors.
NOTICE: Analysis 25% complete ...
NOTICE: Analysis 50% complete ...
NOTICE: Analysis 75% complete ...
NOTICE: Analysis 100% complete ...
NOTICE: Found 0 potential problems in directionality
pgr_analyzeoneway
-----
OK
(1 row)
```

See Also

- [Topology - Family of Functions](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_nodeNetwork - Deprecaded since 3.8.0

pgr\_nodeNetwork - Nodes an network edge table.

Author:

Nicolas Ribot

Copyright:

Nicolas Ribot, The source code is released under the MIT-X license.

The function reads edges from a not "noded" network table and writes the "noded" edges into a new table.

```
| pgr_nodenetwork(edge_table, tolerance, [options])
| options: [id, text the_geom, table_ending, rows_where, outall]

| RETURNS TEXT
```

Availability

- Version 3.8.0
  - Deprecaded function.
  - Not checking and not creating indexes.
  - Using pgr\_separateTouching and pgr\_separateCrossing.
  - Created table with BIGINT.
- Version 2.0.0
  - Official function.

Migration of pgr\_nodeNetwork

Starting from [v3.8.0](#)

**Before Deprecation:** A table with *<edges>\_noded* was created. with split edges.

Migration

Use [pgr\\_separateTouching](#) and/or use [pgr\\_separateCrossing](#)

Description

**The main characteristics are:**

A common problem associated with bringing GIS data into pgRouting is the fact that the data is often not "noded" correctly. This will create invalid topologies, which will result in routes that are incorrect.

What we mean by "noded" is that at every intersection in the road network all the edges will be broken into separate road segments. There are cases like an over-pass and under-pass intersection where you can not traverse from the over-pass to the under-pass, but this function does not have the ability to detect and accommodate those situations.

This function reads the edge\_table table, that has a primary key column id and geometry column named the\_geom and intersect all the segments in it against all the other segments and then creates a table edge\_table\_noded. It uses the tolerance for deciding that multiple nodes within the tolerance are considered the same node.

Parameters

edge\_table:

text Network table name. (may contain the schema name as well)

tolerance:

float8 tolerance for coincident points (in projection unit)

id:

text Primary key column name of the network table. Default value is id.

the\_geom:

text Geometry column name of the network table. Default value is the\_geom.

table\_ending:

text Suffix for the new table's. Default value is noded.

The output table will have for edge\_table\_noded

id:

bigint Unique identifier for the table

old\_id:

bigint Identifier of the edge in original table

sub\_id:

integer Segment number of the original edge

source:

bigint Empty source column

target:

bigint Empty target column

the geom:

geometry Geometry column of the noded network

#### Examples¶

Create the topology for the data in [Sample Data](#)

```
SELECT * INTO vertices
FROM pgr_extractVertices('SELECT id, geom FROM edges ORDER BY id');
SELECT 17
```

```
/* -- set the source information */
UPDATE edges AS e
SET source = v.id, x1 = x, y1 = y
FROM vertices AS v
WHERE ST_StartPoint(e.geom) = v.geom;
UPDATE 18
/* -- set the target information */
UPDATE edges AS e
SET target = v.id, x2 = x, y2 = y
FROM vertices AS v
WHERE ST_EndPoint(e.geom) = v.geom;
UPDATE 18
```

Analyze the network for intersections.

```
SELECT e1.id id1, e2.id id2, ST_ASTEXT(ST_Intersection(e1.geom, e2.geom)) AS point
FROM edges e1, edges e2
WHERE e1.id < e2.id AND ST_Crosses(e1.geom, e2.geom);
id1 | id2 | point
-----+-----+-----
13 | 18 | POINT(3.5 3)
(1 row)
```

Analyze the network for gaps.

```
WITH
data AS (
  SELECT id, geom, (in_edges || out_edges)[1] as inhere
  FROM vertices where array_length(in_edges || out_edges, 1) = 1
),
results AS (
  SELECT d.id, d.inhere,
  (pgr_findCloseEdges('SELECT id, geom FROM edges WHERE id != ' || inhere , geom, 0.001)).*
  FROM data d
)
SELECT
  id, inhere, edge_id, fraction,
  ST_AsText(geom) AS geom, ST_AsText(edge) AS edge
FROM results;
id | inhere | edge_id | fraction | geom | edge
-----+-----+-----+-----+-----+-----
4 | 17 | 14 | 0.5 | POINT(1.9999999999999999 3.5) | LINESTRING(1.9999999999999999 3.5,2 3.5)
(1 row)
```

The analysis tell us that the network has a gap and an intersection.

#### Fixing an intersection¶

Storing the intersections.

```
SELECT e1.id id1, e2.id id2, st_intersection(e1.geom, e2.geom) AS point
INTO intersections
FROM edges e1, edges e2
WHERE e1.id < e2.id AND st_crosses(e1.geom, e2.geom);
SELECT 1
```

Calling pgr\_nodeNetwork.

```
SELECT pgr_nodeNetwork('edges', 0.001, the_geom => 'geom', rows_where=>'id in ('||id1||','||id2||')')
FROM intersections;
WARNING: pgr_nodenetwork(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
NOTICE: PROCESSING:
NOTICE: id: id
NOTICE: the_geom: geom
NOTICE: table_ending: noded
NOTICE: rows_where: id in (13,18)
NOTICE: outall: f
NOTICE: pgr_nodeNetwork('edges', 0.001, 'id', 'geom', 'noded', 'id in (13,18)', f)
NOTICE: Performing checks, please wait .....
NOTICE: Processing, please wait .....
NOTICE: Split Edges: 2
NOTICE: Untouched Edges: 0
NOTICE: Total original Edges: 2
NOTICE: Edges generated: 4
NOTICE: Untouched Edges: 0
NOTICE: Total New segments: 4
NOTICE: New Table: public.edges_noded
NOTICE: -----
pgr_nodenetwork
-----
OK
(1 row)
```

Inspecting the generated table, we can see that edges 13 and 18 have been segmented.

```
SELECT old_id, ST_AsText(geom) FROM edges_noded ORDER BY old_id, sub_id;
```

```
old_id | st_astext
-----+-----
13 | LINESTRING(3 3,3.5 3)
13 | LINESTRING(3.5 3,4 3)
18 | LINESTRING(3.5 2,3,3.5 3)
18 | LINESTRING(3.5 3,3.5 4)
(4 rows)
```

Update the topology

Add new segments to the edges table.

```
INSERT INTO edges (cost, reverse_cost, geom)
WITH
the_fractions AS (
  SELECT e1.id, id, st_lineLocatePoint(e1.geom, point) AS fraction
  FROM intersections, edges e1, edges e2 WHERE e1.id = id1 AND e2.id = id2
  UNION
  SELECT e2.id, st_lineLocatePoint(e2.geom, point)
  FROM intersections, edges e1, edges e2 WHERE e1.id = id1 AND e2.id = id2
)
SELECT
CASE WHEN sub_id = 1
  THEN cost*fraction ELSE cost*(1-fraction) END as cost,
CASE WHEN sub_id = 1
  THEN reverse_cost*(1-fraction) ELSE reverse_cost*(fraction) END AS reverse_cost,
segments.geom
FROM edges as edges
JOIN edges_noded as segments ON (edges.id = segments.old_id)
JOIN the_fractions f ON (segments.old_id = f.id);
INSERT 0 4
```

Insert the intersection as new vertices.

```
INSERT INTO vertices (id, geom)
SELECT row_number() over() + 100, point
FROM intersections;
INSERT 0 1
```

Update source and target information on the edges table.

```
UPDATE edges e SET source = v.id FROM
vertices v where source IS NULL AND (st_startPoint(e.geom) = v.geom);
UPDATE 4
UPDATE edges e SET target = v.id FROM
vertices v where target IS NULL AND (st_endPoint(e.geom) = v.geom);
UPDATE 4
```

Delete original edge.

```
DELETE FROM edges
WHERE id IN (
  SELECT id1 FROM intersections
  UNION
  SELECT id2 FROM intersections);
DELETE 2
```

Update the vertex topology

```
WITH data AS (
  select p.id, p.in_edges, p.out_edges
  FROM pgr_extractVertices(select id, source, target from edges) p)
UPDATE vertices v
SET (in_edges,out_edges) = (d.in_edges,d.out_edges)
FROM data d where d.id = v.id;
UPDATE 18
```

Analyze the network for intersections.

```
SELECT e1.id, e2.id
FROM edges_noded e1, edges e2
WHERE e1.id < e2.id AND st_crosses(e1.geom, e2.geom);
id | id
----+----
(0 rows)
```

Fixing a gap

Store the deadends

```
WITH
data AS (
  SELECT id, geom, ((in_edges || out_edges)[1] as inhere
  FROM vertices where array_length(in_edges || out_edges, 1) = 1)
SELECT
  d.id, d.inhere,
  (pgr_findCloseEdges(SELECT id, geom FROM edges WHERE id != ' || inhere , geom, 0.001)).*
INTO deadends
FROM data d;
SELECT 1
```

Calling pgr\_nodeNetwork.

```
SELECT pgr_nodeNetwork('edges', 0.001, the_geom => 'geom', rows_where=>"id in ('||inhere||','||edge_id||')")
FROM deadends;
WARNING:  pgr_nodenetwork(text,double precision,text,text,text,text,boolean) deprecated function on v3.8.0
NOTICE:  PROCESSING:
NOTICE:  id: id
NOTICE:  the_geom: geom
NOTICE:  table_ending: noded
NOTICE:  rows_where: id in (17,14)
NOTICE:  outall: f
NOTICE:  pgr_nodeNetwork('edges', 0.001, 'id', 'geom', 'noded', 'id in (17,14)', f)
NOTICE:  Performing checks, please wait .....
NOTICE:  Processing, please wait .....
NOTICE:    Split Edges: 1
NOTICE:    Untouched Edges: 1
NOTICE:      Total original Edges: 2
NOTICE:      Edges generated: 2
NOTICE:    Untouched Edges: 1
NOTICE:      Total New segments: 3
NOTICE:    New Table: public.edges_noded
NOTICE:  -----
pgr_nodenetwork
-----
OK
(1 row)
```

Inspecting the generated table, we can see that edge 14 has been segmented.

```
SELECT old_id, ST_AsText(geom) FROM edges_noded ORDER BY old_id, sub_id;
old_id | st_astext
-----+-----
14 | LINESTRING(2 3,1.9999999999999999 3.5)
14 | LINESTRING(1.9999999999999999 3.5,2 4)
```

```
17 | LINESTRING(0.5 3.5,1.999999999999 3.5)
(3 rows)
```

## Update the topology

Add new segments to the edges table.

```

INSERT INTO edges (cost, reverse_cost, geom)
SELECT
  CASE WHEN sub_id = 1 THEN cost*fracton ELSE cost*(1-fracton) END as cost,
  CASE WHEN sub_id = 1 THEN reverse_cost*(1-fracton) ELSE reverse_cost*fracton END as reverse_cost, en,geom
FROM deadends r JOIN edges_noded en ON (old_id = edge_id) JOIN edges e ON (old_id = e.id)
UNION
SELECT 0,0,edge FROM deadends;
INSERT 0 3

```

Insert the intersection as new vertices.

```

/* Update the vertices table */
INSERT INTO vertices (id, geom)
select row_number() over() + 200, st_endpoint(edge) FROM deadends;
INSERT 0 1

```

Update source and target information on the edges table.

```
UPDATE edges e SET source = v.id FROM
vertices v where source IS NULL AND (st_startPoint(e.geom) = v.geom);
UPDATE 3
UPDATE edges e SET target = v.id FROM
vertices v where target IS NULL AND (st_endPoint(e.geom) = v.geom);
UPDATE 3
```

Delete original edge.

```
DELETE FROM edges WHERE id IN (SELECT edge_id FROM deadends);
DELETE 1
```

### Update the vertex topology

```
WITH data AS (
  select p.id, p.in_edges, p.out_edges
  FROM pgr_extractVertices('select id, source, target from edges') p)
UPDATE vertices v
SET (in_edges,out_edges) = (d.in_edges,d.out_edges)
FROM data d where d.id = v.id;
UPDATE 19
```

## Analyze the network for gaps.

```
WITH  
data AS (  
SELECT id, geom, (in_edges || out_edges)[1] as inhre  
FROM vertices where array_length(in_edges || out_edges, 1) = 1),  
results AS (  
SELECT (pgr_findCloseEdges(  
    "SELECT id, geom FROM edges WHERE id != ' || inhre , geom, 0.001))",*,  
d.id, d.inhre  
FROM data d  
)  
SELECT * FROM results;  
  
edge_id | fraction | side | distance | geom | edge | id | inhre
```

17		1		1.0000898900582341e-12		010101000000000000000000000040		010200000002020000000000000000000040		C4068EEFFFFFFFFF3F000000000000C40
----	--	---	--	------------------------	--	--------------------------------	--	--------------------------------------	--	-----------------------------------

(1 row)

### See Also

[Topology - Family of Functions](#) for an overview of a topology for routing algorithms.

## Indices and tables

- [Index](#)
- [Search Page](#)

**pqr extractVertices**

`pgr_extractVertices` — Extracts the vertices information

## Availability

Version 3.8.0

- Error messages adjustment.
- Function promoted to official.

Version 3.3.0

- Function promoted to proposed.

Version 3.0.0

- New experimental function.

### Description

This is an auxiliary function for extracting the vertex information of the set of edges of a graph.

- When the edge identifier is given, then it will also calculate the in and out edges

Boost Graph Inside

## Signatures

```
pgr_extractVertices(Edges SQL, [dryrun])
RETURNS SETOF (id, in_edges, out_edges, x, y, geom)
OR EMPTY SET
```

Example:

## Extracting the vertex information

[illegible]



2 | | {17} | 0.5 | 3.5 | 0101000000000000000000E03F000000000000C40  
3 | {6} | {7} | 1 | 2 | 0101000000000000000000F03F0000000000000040  
4 | {17} | | 1.999999999999 | 3.5 | 010100000068EEFFFFFFFFF3F0000000000000C40  
5 | | {1} | 2 | 0 | 010100000000000000000000004000000000000000  
6 | {1} | {2,4} | 2 | 1 | 0101000000000000000000000040000000000000F03F  
7 | {4,7} | {8,10} | 2 | 2 | 01010000000000000000000000400000000000000040  
8 | {10} | {12,14} | 2 | 3 | 01010000000000000000000000400000000000000840  
9 | {14} | | 2 | 4 | 010100000000000000000000004000000000000001040  
10 | {2} | {3,5} | 3 | 1 | 01010000000000000000000000840000000000000F03F  
11 | {5,8} | {9,11} | 3 | 2 | 010100000000000000000000008400000000000000040  
12 | {11,12} | {13} | 3 | 3 | 010100000000000000000000008400000000000000840  
13 | | {18} | 3.5 | 2.3 | 01010000000000000000000000C406666666666660240  
14 | {18} | | 3.5 | 4 | 01010000000000000000000000C4000000000000001040  
15 | {3} | {16} | 4 | 1 | 010100000000000000000000001040000000000000F03F  
16 | {9,16} | {15} | 4 | 2 | 01010000000000000000000010400000000000000040  
17 | {13,15} | | 4 | 3 | 010100000000000000000000104000000000000000840  
(17 rows)

Parameters¶

Parameter Type Description

Edges SQL TEXT Edges SQL as described below

Optional parameters¶

Parameter Type Default Description

dryrun BOOLEAN false • When true do not process and get in a NOTICE the resulting query.

Inner Queries¶

- Edges SQL
  - When line geometry is known
  - When vertex geometry is known
  - When identifiers of vertices are known

Edges SQL¶

When line geometry is known¶

Column Type Description

id BIGINT (Optional) identifier of the edge.

geom LINESTRING Geometry of the edge.

This inner query takes precedence over the next two inner query, therefore other columns are ignored whergeom column appears.

- Ignored columns:
  - startpoint
  - endpoint
  - source
  - target

When vertex geometry is known¶

To use this inner query the column geom should not be part of the set of columns.

Column Type Description

id BIGINT (Optional) identifier of the edge.

startpoint POINT POINT geometry of the starting vertex.

endpoint POINT POINT geometry of the ending vertex.

This inner query takes precedence over the next inner query, therefore other columns are ignored wherstartpoint and endpoint columns appears.

- Ignored columns:
  - source
  - target

When identifiers of vertices are known¶

To use this inner query the columns geom, startpoint and endpoint should not be part of the set of columns.

Column Type Description

id BIGINT (Optional) identifier of the edge.

Column	Type	Description
source	ANY-INTEGER	Identifier of the first end point vertex of the edge.
target	ANY-INTEGER	Identifier of the second end point vertex of the edge.

Result columns¶

Column	Type	Description
id	BIGINT	Vertex identifier
in_edges	BIGINT[]	Array of identifiers of the edges that have the vertexid as <i>first end point</i> . <ul style="list-style-type: none"><li>• NULL When the id is not part of the inner query</li></ul>
out_edges	BIGINT[]	Array of identifiers of the edges that have the vertexid as <i>second end point</i> . <ul style="list-style-type: none"><li>• NULL When the id is not part of the inner query</li></ul>
x	FLOAT	X value of the point geometry <ul style="list-style-type: none"><li>• NULL When no geometry is provided</li></ul>
y	FLOAT	X value of the point geometry <ul style="list-style-type: none"><li>• NULL When no geometry is provided</li></ul>
geom	POINT	Geometry of the point <ul style="list-style-type: none"><li>• NULL When no geometry is provided</li></ul>

Additional Examples¶

- [Dry run execution](#)
- [Create a routing topology](#)
  - [Make sure the database does not have the vertices table](#)
  - [Clean up the columns of the routing topology to be created](#)
  - [Create the vertices table](#)
  - [Inspect the vertices table](#)
  - [Create the routing topology on the edge table](#)
  - [Inspect the routing topology](#)

Dry run execution¶

To get the query generated used to get the vertex information, use `dryrun := true`.  
The results can be used as base code to make a refinement based on the backend development needs.

```
SELECT * FROM pgr_extractVertices(
'SELECT id, geom FROM edges',
dryrun => true);
NOTICE:
WITH

main_sql AS (
SELECT id, geom FROM edges
),

the_out AS (
SELECT id::BIGINT AS out_edge, ST_StartPoint(geom) AS geom
FROM main_sql
),

agg_out AS (
SELECT array_agg(out_edge ORDER BY out_edge) AS out_edges, ST_X(geom) AS x, ST_Y(geom) AS y, geom
FROM the_out
GROUP BY geom
),

the_in AS (
SELECT id::BIGINT AS in_edge, ST_EndPoint(geom) AS geom
FROM main_sql
),

agg_in AS (
SELECT array_agg(in_edge ORDER BY in_edge) AS in_edges, ST_X(geom) AS x, ST_Y(geom) AS y, geom
FROM the_in
GROUP BY geom
),

the_points AS (
SELECT in_edges, out_edges, coalesce(agg_out.geom, agg_in.geom) AS geom
FROM agg_out
FULL OUTER JOIN agg_in USING (x, y)
)

SELECT row_number() over(ORDER BY ST_X(geom), ST_Y(geom)) AS id, in_edges, out_edges, ST_X(geom), ST_Y(geom), geom
FROM the_points;
id | in_edges | out_edges | x | y | geom
-----+-----+-----+---+---+-----
(0 rows)
```

Create a routing topology¶

Make sure the database does not have the vertices table¶

```
DROP TABLE IF EXISTS vertices_table;
NOTICE: table "vertices_table" does not exist, skipping
DROP TABLE
```

**Clean up the columns of the routing topology to be created**

```
UPDATE edges
SET source = NULL, target = NULL,
x1 = NULL, y1 = NULL,
x2 = NULL, y2 = NULL;
UPDATE 18
```

## Create the vertices table

- When the LINESTRING has a SRID then use `geom::geometry(POINT, <SRID>)`
- For big edge tables that are been prepared,
  - Create it as UNLOGGED and
  - After the table is created `ALTER TABLE .. SET LOGGED`

```
SELECT * INTO vertices_table
FROM pgr_extractVertices('SELECT id, geom FROM edges ORDER BY id');
SELECT 17
```

### Inspect the vertices table¶

[illegible]

Create the routing topology on the edge table

```
WITH
out_going AS (
  SELECT id AS vid, unnest(out_edges) AS eid, x, y
  FROM vertices_table
)
UPDATE edges
SET source = vid, x1 = x, y1 = y
FROM out_going WHERE id = eid;
UPDATE 18
```


### Updating the target information

```
WITH
in_coming AS (
  SELECT id AS vid, unnest(in_edges) AS eid, x, y
  FROM vertices_table
)
UPDATE edges
SET target = vid, x2 = x, y2 = y
FROM in_coming WHERE id = eid;
UPDATE 18
```

### Inspect the routing topology

SELECT id, source, target, x1, y1, x2, y2						
FROM edges ORDER BY id;						
id	source	target	x1	y1	x2	y2
1	5	6	2	0	2	1
2	6	10	2	1	3	1
3	10	15	3	1	4	1
4	6	7	2	1	2	2
5	10	11	3	1	3	2
6	1	3	0	2	1	2
7	3	7	1	1	2	2
8	7	11	2	2	3	2
9	11	16	3	2	4	2
10	7	8	2	2	2	3
11	11	12	2	2	3	3
12	8	12	2	3	3	3
13	12	17	3	3	4	3
14	8	9	2	3	2	4
15	16	17	4	2	4	3
16	15	16	4	1	4	2
17	2	4	0.5	3.5	1.9999999999999999	3.5
18	13	14	3.5	2.3	3.5	4

[\\_images/Fig1-originalData.png](#)



Generated topology

See Also

- [Topology - Family of Functions](#)

Indices and tables

- [Index](#)
- [Search Page](#)

`pgr_findCloseEdges`

`pgr_findCloseEdges` - Finds the close edges to a point geometry.

Availability

Version 3.8.0

- Error messages adjustment.
- `partial` option is removed.
- Function promoted to official.

Version 3.4.0

- New proposed function.

Description

`pgr_findCloseEdges` - An utility function that finds the closest edge to a point geometry.

- The geometries must be in the same coordinate system (have the same SRID).
- The code to do the calculations can be obtained for further specific adjustments needed by the application.
- EMPTY SET is returned on dryrun executions

☐ Boost Graph Inside

Signatures

Summary

`pgr_findCloseEdges`([Edges SQL](#), **point**, **tolerance**, [**options**])  
`pgr_findCloseEdges`([Edges SQL](#), **points**, **tolerance**, [**options**])  
**options**: [`cap`, `dryrun`]  
Returns set of (edge\_id, fraction, side, distance, geom, edge)  
OR EMPTY SET

One point

`pgr_findCloseEdges`([Edges SQL](#), **point**, **tolerance**, [**options**])  
**options**: [`cap`, `dryrun`]  
Returns set of (edge\_id, fraction, side, distance, geom, edge)  
OR EMPTY SET

Example:

Get two close edges to points of interest with(pid = 5)

- `cap => 2`

```
SELECT
  edge_id, fraction, side, distance,
  distance, ST_AsText(geom) AS point, ST_AsText(edge) AS edge
FROM pgr_findCloseEdges(
  $$SELECT id, geom FROM edges$,
  (SELECT geom FROM pointsOfInterest WHERE pid = 5),
  0.5, cap => 2);
edge_id | fraction | side | distance | distance | point | edge
-----+-----+-----+-----+-----+-----+-----
5 | 0.8 | l | 0.1 | 0.1 | POINT(2.9 1.8) | LINESTRING(2.9 1.8,3 1.8)
8 | 0.9 | r | 0.2 | 0.2 | POINT(2.9 1.8) | LINESTRING(2.9 1.8,2.9 2)
(2 rows)
```

Many points

`pgr_findCloseEdges`([Edges SQL](#), **points**, **tolerance**, [**options**])  
**options**: [`cap`, `dryrun`]  
Returns set of (edge\_id, fraction, side, distance, geom, edge)  
OR EMPTY SET

Example:

For each points of interests, find the closest edge.

```
SELECT
edge_id, fraction, side, distance,
ST_AsText(geom) AS point,
ST_AsText(edge) AS edge
FROM pgr_findCloseEdges(
  $$SELECT id, geom FROM edges$$,
  (SELECT array_agg(geom) FROM pointsOfInterest),
  0.5);
edge_id | fraction | side | distance | point | edge
-----+-----+-----+-----+-----+-----
1 | 0.4 | l | 0.2 | POINT(1.8 0.4) | LINESTRING(1.8 0.4,2 0.4)
6 | 0.3 | r | 0.2 | POINT(0.3 1.8) | LINESTRING(0.3 1.8,0.3 2)
12 | 0.6 | l | 0.2 | POINT(2.6 3.2) | LINESTRING(2.6 3.2,2.6 3)
15 | 0.4 | r | 0.2 | POINT(4.2 2.4) | LINESTRING(4.2 2.4,4 2.4)
5 | 0.8 | l | 0.1 | POINT(2.9 1.8) | LINESTRING(2.9 1.8,3 1.8)
4 | 0.7 | r | 0.2 | POINT(2.2 1.7) | LINESTRING(2.2 1.7,2 1.7)
(6 rows)
```

Parameters

Parameter	Type	Description
-----------	------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below.

**point** POINT The point geometry

**points** POINT[] An array of point geometries

**tolerance** FLOAT Max distance between geometries

Optional parameters

Parameter	Type	Default	Description
cap	INTEGER	\(1\)	Limit output rows
dryrun	BOOLEAN	false	<ul style="list-style-type: none"><li>When false calculations are performed.</li><li>When true calculations are not performed and the query to do the calculations is exposed in a PostgreSQL NOTICE.</li></ul>

Inner Queries

Edges SQL

Column	Type	Description
--------	------	-------------

**id** **ANY-INTEGER** Identifier of the edge.

**geom** geometry The LINESTRING geometry of the edge.

Result columns

Returns set of (edge\_id, fraction, side, distance, geom, edge)

Column	Type	Description
edge_id	BIGINT	Identifier of the edge. <ul style="list-style-type: none"><li>When \(\text{cap} = 1\), it is the closest edge.</li></ul>
fraction	FLOAT	Value in $<0,1>$ that indicates the relative position from the first end-point of the edge.
side	CHAR	Value in $\{r, l\}$ indicating if the point is: <ul style="list-style-type: none"><li>At the right <math>r</math> of the segment.<ul style="list-style-type: none"><li>When the point is on the line it is considered to be on the right.</li></ul></li><li>At the left <math>l</math> of the segment.</li></ul>
distance	FLOAT	Distance from the point to the edge.
geom	geometry	Original POINT geometry.
edge	geometry	LINESTRING geometry that connects the original <b>point</b> to the closest point of the edge with identifier <code>edge_id</code>

Additional Examples

- [One point in an edge](#)
- [One point dry run execution](#)
- [Many points in an edge](#)
- [Many points dry run execution](#)

- ### One point in an edge¶

- ### One point dry run execution

### Many points in an edge

[illegible]

```
ST_MakeLine(point, ST_ClosestPoint(geom, point)) AS new_line
FROM edges_sql, point_sql
WHERE ST_DWithin(geom, point, 5)
ORDER BY geom <-> point),
prepare_cap AS (
SELECT row_number() OVER (PARTITION BY point ORDER BY point, distance) AS m, *
FROM results)
SELECT edge_id, fraction, side, distance, point, new_line
FROM prepare_cap
WHERE m <= 1

edge_id | fraction | side | distance | geom | edge
-----+-----+-----+-----+-----+-----
(0 rows)
```

Find at most two routes to a given point¶

Using [pgr\\_withPoints](#) - Proposed

```
SELECT * FROM pgr_withPoints(
  $$ SELECT * FROM edges $$,
  $$ SELECT edge_id, fraction, side
    FROM pgr_findCloseEdges(
      $$ SELECT id, geom FROM edges $$,
      (SELECT geom FROM pointsOfInterest WHERE pid = 5),
      0.5, cap => 2)
  $$,
  1, ARRAY[-1, -2]);
seq | path_seq | end_pid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | -2 | 1 | 6 | 1 | 0
2 | 2 | -2 | 3 | 7 | 1 | 1
3 | 3 | -2 | 7 | 8 | 0.9 | 2
4 | 4 | -2 | -1 | 0 | 2.9
5 | 1 | -1 | 1 | 6 | 1 | 0
6 | 2 | -1 | 3 | 7 | 1 | 1
7 | 3 | -1 | 7 | 8 | 1 | 2
8 | 4 | -1 | 11 | 9 | 1 | 3
9 | 5 | -1 | 16 | 16 | 1 | 4
10 | 6 | -1 | 15 | 3 | 1 | 5
11 | 7 | -1 | 10 | 5 | 0.8 | 6
12 | 8 | -1 | -1 | -1 | 0 | 6.8
(12 rows)
```

A point of interest table¶

Handling points outside the graph.

Points of Interest¶

Some times the applications work "on the fly" starting from a location that is not a vertex in the graph. Those locations, in pgRouting are called points of interest.

The information needed in the points of interest ispid, edge\_id, side, fraction.

On this documentation there will be some 6 fixed points of interest and they will be stored on a table.

Column	Description
pid	A unique identifier.
edge_id	Identifier of the nearest segment.
side	Is it on the left, right or both sides of the segmentedge_id.
fraction	Where in the segment is the point located.
geom	The geometry of the points.
distance	The distance between geom and the segmentedge_id.
edge	A segment that connects thegeom of the point to the closest point on the segment edge_id.
newPoint	A point on segment edge_id that is the closest to geom.

```
CREATE TABLE pointsOfInterest(
  pid BIGSERIAL PRIMARY KEY,
  edge_id BIGINT,
  side CHAR,
  fraction FLOAT,
  distance FLOAT,
  edge geometry,
  newPoint geometry,
  geom geometry);
IF v > 3.4 THEN
```

Points of interest fill up¶

Inserting the points of interest.

```
INSERT INTO pointsOfInterest (geom) VALUES
(ST_Point(1.8, 0.4)),
(ST_Point(4.2, 2.4)),
(ST_Point(2.6, 3.2)),
(ST_Point(0.3, 1.8)),
(ST_Point(2.9, 1.8)),
(ST_Point(2.2, 1.7));
```

Filling the rest of the table.

```
UPDATE pointsofinterest SET
  edge_id = poi.edge_id,
  side = poi.side,
  fraction = round(poi.fraction::numeric, 2),
  distance = round(poi.distance::numeric, 2),
  edge = poi.edge,
  newPoint = ST_EndPoint(poi.edge)
FROM (
  SELECT *
  FROM pgr_findCloseEdges(
```

```
$$SELECT id, geom FROM edges$$,(SELECT array_agg(geom) FROM pointsOfInterest), 0.5) ) AS poi
WHERE pointsOfInterest.geom = poi.geom;
```

Any other additional modification: In this manual, point\{6\} can be reached from both sides.

```
UPDATE pointsOfInterest SET side = 'b' WHERE pid = 6;
```

The points of interest:

```
SELECT
pid, ST_AsText(geom) geom,
edge_id, fraction AS frac, side, distance AS dist,
ST_AsText(edge) edge, ST_AsText(newPoint) newPoint
FROM pointsOfInterest;
pid | geom | edge_id | frac | side | dist | edge | newpoint
-----+-----+-----+-----+-----+-----+-----+-----
1 | POINT(1.8 0.4) | 1 | 0.4 | l | 0.2 | LINESTRING(1.8 0.4,2 0.4) | POINT(2 0.4)
4 | POINT(0.3 1.8) | 6 | 0.3 | r | 0.2 | LINESTRING(0.3 1.8,0.3 2) | POINT(0.3 2)
3 | POINT(2.6 3.2) | 12 | 0.6 | l | 0.2 | LINESTRING(2.6 3.2,2.6 3) | POINT(2.6 3)
2 | POINT(4.2 2.4) | 15 | 0.4 | r | 0.2 | LINESTRING(4.2 2.4,4 2.4) | POINT(4 2.4)
5 | POINT(2.9 1.8) | 5 | 0.8 | l | 0.1 | LINESTRING(2.9 1.8,3 1.8) | POINT(3 1.8)
6 | POINT(2.2 1.7) | 4 | 0.7 | b | 0.2 | LINESTRING(2.2 1.7,2 1.7) | POINT(2 1.7)
(6 rows)
```

See Also

- [withPoints - Category](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_separateCrossing

pgr\_separateCrossing - From crossing geometries generates geometries that do not cross.

Availability

Version 3.8.0

- Function promoted to official.
- Proposed function.

Description

This is an auxiliary function for separating crossing edges.

Signature

```
pgr_separateCrossing(Edges SQL, [tolerance, dryrun])
RETURNS (seq_id,sub_id,geom)
```

Example:

Get the segments of the crossing geometries

```
SELECT id, sub_id, ST_AsText(geom)
FROM pgr_separateCrossing('SELECT id, geom FROM edges')
ORDER BY id, sub_id;
id | sub_id | st_astext
-----+-----+-----
13 | 1 | LINESTRING(3 3,3.5 3)
13 | 2 | LINESTRING(3.5 3,4 3)
18 | 1 | LINESTRING(3.5 2.3,3.5 3)
18 | 2 | LINESTRING(3.5 3,3.5 4)
(4 rows)
```

Parameters

Parameter	Type	Description
-----------	------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below

Optional parameters

Parameter	Type	Default	Description
tolerance	FLOAT	0.01	Used in ST_Snap before ST_Split
dryrun	BOOLEAN	false	<ul style="list-style-type: none"><li>When true do not process and get in a NOTICE the resulting query.</li></ul>

Inner Queries

Edges SQL

Column	Type	Description
id	ANY-INTEGER	(Optional) identifier of the edge.
geom	LINESTRING	Geometry of the edge.

Examples



- [Get the code for further refinement.](#)
- [Fixing an intersection](#)

#### [Get the code for further refinement.¶](#)

When there are special details that need to be taken care of because of the final application or the quality of the data, the code can be obtained On a PostgreSQLOTICE using the `dryrun` flag.

```
SELECT *
FROM pgr_separateCrossing('SELECT id, geom FROM edges', dryrun => true);
NOTICE:
WITH
  edges_table AS (
    SELECT id, geom FROM edges
  ),

get_crossings AS (
  SELECT e1.id id1, e2.id id2, e1.geom AS g1, e2.geom AS g2, ST_Intersection(e1.geom, e2.geom) AS point
  FROM edges_table e1, edges_table e2
  WHERE e1.id < e2.id AND ST_Crosses(e1.geom, e2.geom)
),

crossings AS (
  SELECT id1, g1, point FROM get_crossings
  UNION
  SELECT id2, g2, point FROM get_crossings
),

blades AS (
  SELECT id1, g1, ST_UnaryUnion(ST_Collect(point)) AS blade
  FROM crossings
  GROUP BY id1, g1
),

collection AS (
  SELECT id1, (st_dump(st_split(st_snap(g1, blade, 0.01), blade))) *
  FROM blades
)

SELECT row_number() over()::INTEGER AS seq, id1::BIGINT, path[1], geom
FROM collection;
;
seq | id | sub_id | geom
-----+-----+-----+-----
(0 rows)
```

#### [Fixing an intersection¶](#)

In this example the original edge table will be used to store the additional geometries.

An example use without results

Routing from `\(1\)` to `\(18\)` gives no solution.

```
SELECT *
FROM pgr_dijkstra('SELECT id, source, target, cost, reverse_cost FROM edges', 1, 18);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
(0 rows)
```

Analyze the network for intersections.

```
SELECT
  e1.id id1, e2.id id2,
  ST_AsText(ST_Intersection(e1.geom, e2.geom)) AS point
FROM edges e1, edges e2
WHERE e1.id < e2.id AND ST_Crosses(e1.geom, e2.geom);
id1 | id2 | point
-----+-----+-----
13 | 18 | POINT(3.5 3)
(1 row)
```

The analysis tell us that the network has an intersection.

Prepare tables

Additional columns to control the origin of the segments.

```
ALTER TABLE edges ADD old_id BIGINT;
ALTER TABLE
```

Adding new segments.

Calling [pgr\\_separateCrossing](#) and adding the new segments to the edges table.

```
INSERT INTO edges (old_id, geom)
SELECT id, geom
FROM pgr_separateCrossing('SELECT id, geom FROM edges');
INSERT 0 4
```

Update other values

In this example only `cost` and `reverse_cost` are updated, where they are based on the length of the geometry and the directionality is kept using the `sign` function.

```
WITH
costs AS (
  SELECT e2.id, sign(e1.cost) * ST_Length(e2.geom) AS cost,
  sign(e1.reverse_cost) * ST_Length(e2.geom) AS reverse_cost
  FROM edges e1 JOIN edges e2 ON (e1.id = e2.old_id)
)
UPDATE edges e
SET (cost, reverse_cost) = (c.cost, c.reverse_cost)
FROM costs AS c WHERE e.id = c.id;
UPDATE 4
```

Update the topology

Insert the new vertices if any.

```
WITH
new_vertex AS (
  SELECT ev.*
  FROM pgr_extractVertices('SELECT id, geom FROM edges WHERE old_id IS NOT NULL') ev
  LEFT JOIN vertices v using(geom)
  WHERE v IS NULL)
INSERT INTO vertices (in_edges, out_edges,x,y,geom)
SELECT in_edges, out_edges,x,y,geom FROM new_vertex;
INSERT 0 1
```

Update source and target information on the edges table.

```
/* -- set the source information */
UPDATE edges AS e
```

The example has results

```
SELECT *
FROM pgr_dijkstra('SELECT id, source, target, cost, reverse_cost FROM edges', 1, 18);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

## See Also

## Indices and tables

- [Index](#)
- [Search Page](#)

`pgr_separateTouching` - From touching geometries generates geometries that are properly connected at endpoints

Version 3.8.0

- Function promoted to official.
- Proposed function.

This is an auxiliary function for processing geometries that touch but don't share exact endpoints, splitting them at their intersection points to improve network connectivity.

```
pgr_separateTouching(Edges SQL, [tolerance, dryrun])  
RETURNS (seq,id,sub_id,geom)
```

### Get the segments of the crossing geometries

## Parameters

### Optional parameters

## Inner Queries

## Examples

- [Get the code for further refinement.](#)

- [Fixing a gap](#)

[Get the code for further refinement¶](#)

When there are special details that need to be taken care of because of the final application or the quality of the data, the code can be obtained On a PostgreSQLOTICE using the `dryrun` flag.

```
SELECT *
FROM pgr_separateTouching('SELECT id, geom FROM edges', dryrun => true);
NOTICE:
WITH
  edges_table AS (
    SELECT id, geom FROM edges
  ),

  get_touching AS (
    SELECT e1.id id1, e2.id id2, ST_Snap(e1.geom, e2.geom, 0.01) AS geom, e1.geom AS g1, e2.geom AS g2
    FROM edges_table e1, edges_table e2
    WHERE e1.id != e2.id AND ST_DWithin(e1.geom, e2.geom, 0.01) AND NOT(
      ST_StartPoint(e1.geom) = ST_StartPoint(e2.geom) OR ST_StartPoint(e1.geom) = ST_EndPoint(e2.geom)
      OR ST_EndPoint(e1.geom) = ST_StartPoint(e2.geom) OR ST_EndPoint(e1.geom) = ST_EndPoint(e2.geom))
  ),

  touchings AS (
    SELECT id1, g1, g2, st_intersection(geom, g2) AS point
    FROM get_touching
    WHERE NOT (geom = g1) OR
      (ST_touches(g1, g2) AND NOT
        (ST_Intersection(geom, g2) = ST_StartPoint(g1)
        OR ST_Intersection(geom, g2) = ST_EndPoint(g1)))
  ),

  blades AS (
    SELECT id1, g1, ST_UnaryUnion(ST_Collect(point)) AS blade
    FROM touchings
    GROUP BY id1, g1
  ),

  collection AS (
    SELECT id1, (st_dump(st_split(st_snap(g1, blade, 0.01), blade))) *
    FROM blades
  )

  SELECT row_number() over()::INTEGER AS seq, id1::BIGINT, path[1], geom
  FROM collection;
;
seq | id | sub_id | geom
-----+-----+-----+-----
(0 rows)
```

[Fixing a gap¶](#)

In this example the original edge table will be used to store the additional geometries.

An example use without results

Routing from `\(1\)` to `\(2\)` gives no solution.

```
SELECT *
FROM pgr_dijkstra('SELECT id, source, target, cost, reverse_cost FROM edges', 1, 2);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
(0 rows)
```

Analyze the network for gaps.

```
WITH
  deadends AS (
    SELECT id AS vid, (in_edges || out_edges)[1] AS edge, geom AS vgeom
    FROM vertices
    WHERE array_length(in_edges || out_edges, 1) = 1
  )
SELECT id, ST_AsText(geom), vid, ST_AsText(vgeom), ST_Distance(geom, vgeom)
FROM edges, deadends
WHERE id != edge AND ST_Distance(geom, vgeom) < 0.1;
id | st_astext | vid | st_astext | st_distance
-----+-----+-----+-----+-----
14 | LINESTRING(2 3,2 4) | 4 | POINT(1.9999999999999999 3.5) | 1.000088900582341e-12
(1 row)
```

The analysis tell us that the network has a gap.

Prepare tables

Additional columns to control the origin of the segments.

```
ALTER TABLE edges ADD old_id BIGINT;
ALTER TABLE
```

Adding new segments.

Calling [pgr\\_separateTouching](#) and adding the new segments to the edges table.

```
INSERT INTO edges (old_id, geom)
SELECT id, geom
FROM pgr_separateTouching('SELECT id, geom FROM edges');
INSERT 0 2
```

Update other values

In this example only `cost` and `reverse_cost` are updated, where they are based on the length of the geometry and the directionality is kept using the `sign` function.

```
WITH
  costs AS (
    SELECT e2.id,
      sign(e1.cost) * ST_Length(e2.geom) AS cost,
      sign(e1.reverse_cost) * ST_Length(e2.geom) AS reverse_cost
    FROM edges e1
    JOIN edges e2 ON (e1.id = e2.old_id)
  )
UPDATE edges e SET (cost, reverse_cost) = (c.cost, c.reverse_cost)
FROM costs AS c
WHERE e.id = c.id;
UPDATE 2
```

Update the topology

Insert the new vertices if any.

```
WITH new_vertex AS (
  SELECT ev.*
  FROM pgr_extractVertices('SELECT id, geom FROM edges WHERE old_id IS NOT NULL') ev
  LEFT JOIN vertices v using(geom)
  WHERE v IS NULL
```

```
)
INSERT INTO vertices (in_edges, out_edges,x,y,geom)
SELECT in_edges, out_edges,x,y,geom
FROM new_vertex;
INSERT 0 0
```

Update source and target information on the edges table.

```
/* -- set the source information */
UPDATE edges AS e
SET source = v.id, x1 = x, y1 = y
FROM vertices AS v
WHERE source IS NULL AND ST_StartPoint(e.geom) = v.geom;
UPDATE 2
/* -- set the target information */
UPDATE edges AS e
SET target = v.id, x2 = x, y2 = y
FROM vertices AS v
WHERE target IS NULL AND ST_EndPoint(e.geom) = v.geom;
UPDATE 2
```

The example has results

Routing from \{1\} to \{2\} gives a solution.

```
SELECT *
FROM pgr_dijkstra('SELECT id, source, target, cost, reverse_cost FROM edges', 1, 2);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 2 | 1 | 6 | 1 | 0
2 | 2 | 1 | 2 | 2 | 3 | 7 | 1 | 1
3 | 3 | 1 | 2 | 7 | 10 | 1 | 2
4 | 4 | 1 | 2 | 8 | 19 | 0.5 | 3
5 | 5 | 1 | 2 | 4 | 17 | 1 | 3.5
6 | 6 | 1 | 2 | -1 | 0 | 4.5
(6 rows)
```

See Also¶

[Topology - Family of Functions](#) for an overview of a topology for routing algorithms.

Indices and tables

- [Index](#)
- [Search Page](#)

See Also¶

Indices and tables

- [Index](#)
- [Search Page](#)

Traveling Sales Person - Family of functions¶

- [pgr\\_TSP](#) - When input is given as matrix cell information.
- [pgr\\_TSPeuclidean](#) - When input are coordinates.

pgr\_TSP¶

- [pgr\\_TSP](#) - Approximation using *metric* algorithm.

Availability:

- Version 3.2.1
  - Metric Algorithm from [Boost library](#)
  - Simulated Annealing Algorithm no longer supported
    - The Simulated Annealing Algorithm related parameters are ignored: max\_processing\_time, tries\_per\_temperature, max\_changes\_per\_temperature, max\_consecutive\_non\_changes, initial\_temperature, final\_temperature, cooling\_factor, randomize
- Version 2.3.0
  - Signature change
    - Old signature no longer supported
- Version 2.0.0
  - Official function.

Description¶

Problem Definition¶

The travelling salesperson problem (TSP) asks the following question:

*Given a list of cities and the distances between each pair of cities, which is the shortest possible route that visits each city exactly once and returns to the origin city?*

Characteristics¶

- This problem is an NP-hard optimization problem.
- Metric Algorithm is used
- Implementation generates solutions that *are twice as long as the optimal tour in the worst case* when:
  - Graph is undirected
  - Graph is fully connected
  - Graph where traveling costs on edges obey the triangle inequality.
- On an undirected graph:
  - The traveling costs are symmetric:
  - Traveling costs from *u* to *v* are just as much as traveling from *v* to *u*

- Can be Used with [Cost Matrix - Category](#) functions preferably with *directed => false*.
  - With directed => false
    - Will generate a graph that:
      - is undirected
      - is fully connected (As long as the graph has one component)
      - all traveling costs on edges obey the triangle inequality.
    - When start\_vid = 0 OR end\_vid = 0
      - The solutions generated are guaranteed to be *twice as long as the optimal tour in the worst case*
    - When start\_vid != 0 AND end\_vid != 0 AND start\_vid != end\_vid
      - It is **not guaranteed** that the solution will be, in the worst case, twice as long as the optimal tour, due to the fact that *end\_vid* is forced to be in a fixed position.
  - With directed => true
    - It is **not guaranteed** that the solution will be, in the worst case, twice as long as the optimal tour
    - Will generate a graph that:
      - is directed
      - is fully connected (As long as the graph has one component)
      - some (or all) traveling costs on edges might not obey the triangle inequality.
    - As an undirected graph is required, the directed graph is transformed as follows:
      - edges  $(u, v)$  and  $(v, u)$  is considered to be the same edge (denoted  $(u, v)$ )
      - if *agg\_cost* differs between one or more instances of edge  $(u, v)$
      - The minimum value of the *agg\_cost* all instances of edge  $(u, v)$  is going to be considered as the *agg\_cost* of edge  $(u, v)$
      - Some (or all) traveling costs on edges will still might not obey the triangle inequality.
  - When the data is incomplete, but it is a connected graph:
    - the missing values will be calculated with dijkstra algorithm.

Boost Graph Inside

Signatures

Summary

pgr\_TSP([Matrix SQL](#), [start\_id, end\_id])  
Returns set of (seq, node, cost, agg\_cost)  
OR EMPTY SET

Example:

- Using [pgr\\_dijkstraCostMatrix](#) to generate the matrix information
- **Line 4** Vertices  $\{(2, 4, 13, 14)\}$  are not included because they are not connected.

```
1 SELECT * FROM pgr_TSP(
2   $$SELECT * FROM pgr_dijkstraCostMatrix(
3     SELECT id, source, target, cost, reverse_cost FROM edges,
4     (SELECT array_agg(id) FROM vertices WHERE id NOT IN (2, 4, 13, 14)),
5     directed => false) $$);
6 seq | node | cost | agg_cost
7 ----+-----+-----+-----
8 1 | 1 | 0 | 0
9 2 | 3 | 1 | 1
10 3 | 7 | 1 | 2
11 4 | 6 | 1 | 3
12 5 | 5 | 1 | 4
13 6 | 10 | 2 | 6
14 7 | 11 | 1 | 7
15 8 | 12 | 1 | 8
16 9 | 16 | 2 | 10
17 10 | 15 | 1 | 11
18 11 | 17 | 2 | 13
19 12 | 9 | 3 | 16
20 13 | 8 | 1 | 17
21 14 | 1 | 3 | 20
22 (14 rows)
23
```

Parameters

Parameter	Type	Description
<a href="#">Matrix SQL</a>	TEXT	<a href="#">Matrix SQL</a> as described below

TSP optional parameters

Column	Type	Default	Description
start_id	ANY-INTEGER	0	The first visiting vertex
			<ul style="list-style-type: none"><li>• When 0 any vertex can become the first visiting vertex.</li></ul>
end_id	ANY-INTEGER	0	Last visiting vertex before returning to start_vid.
			<ul style="list-style-type: none"><li>• When 0 any vertex can become the last visiting vertex before returning to start_id.</li></ul>
			<ul style="list-style-type: none"><li>• When NOT 0 and start_id = 0 then it is the first and last vertex</li></ul>

Inner Queries

Column	Type	Description
start_vid	ANY-INTEGER	Identifier of the starting vertex.
end_vid	ANY-INTEGER	Identifier of the ending vertex.
agg_cost	ANY-NUMERICAL	Cost for going from start_vid to end_vid

Result columns

Returns SET OF (seq, node, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Row sequence.
node	BIGINT	Identifier of the node/coordinate/point.
cost	FLOAT	Cost to traverse from the current node to the next node in the path sequence. <ul style="list-style-type: none"><li>0 for the last row in the tour sequence.</li></ul>
agg_cost	FLOAT	Aggregate cost from the node at seq = 1 to the current node. <ul style="list-style-type: none"><li>0 for the first row in the tour sequence.</li></ul>

Additional Examples

- Start from vertex \1\
- Using points of interest to generate an asymmetric matrix.
- Connected incomplete data

Start from vertex \1\

- Line 6 start\_vid => 1

```
1 SELECT * FROM pgr_TSP(
2   $$SELECT * FROM pgr_dijkstraCostMatrix(
3     'SELECT id, source, target, cost, reverse_cost FROM edges',
4     (SELECT array_agg(id) FROM vertices WHERE id NOT IN (2, 4, 13, 14)),
5     directed => false) $$,
6   start_vid => 1);
7 seq | node | cost | agg_cost
8 ----+-----+-----+-----
9  1 |  1 |  0 |    0
10  2 |  3 |  1 |    1
11  3 |  7 |  1 |    2
12  4 |  6 |  1 |    3
13  5 |  5 |  1 |    4
14  6 | 10 |  2 |    6
15  7 | 11 |  1 |    7
16  8 | 12 |  1 |    8
17  9 | 16 |  2 |   10
18 10 | 15 |  1 |   11
19 11 | 17 |  2 |   13
20 12 |  9 |  3 |   16
21 13 |  8 |  1 |   17
22 14 |  1 |  3 |   20
23 (14 rows)
24
```

Using points of interest to generate an asymmetric matrix

To generate an asymmetric matrix:

- Line 4 The side information of pointsOfInterest is ignored by not including it in the query
- Line 6 Generating an asymmetric matrix with directed => true
  - \(min(agg\\_cost(u, v), agg\\_cost(v, u))\) is going to be considered as the agg\_cost
  - The solution that can be larger than *twice as long as the optimal tour* because:
    - Triangle inequality might not be satisfied.
    - start\_id != 0 AND end\_id != 0

```
1 SELECT * FROM pgr_TSP(
2   $$SELECT * FROM pgr_withPointsCostMatrix(
3     'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
4     'SELECT pid, edge_id, fraction from pointsOfInterest',
5     array[-1, 10, 7, 11, -6],
6     directed => true) $$,
7   start_id => 7,
8   end_id => 11);
9 seq | node | cost | agg_cost
10 ----+-----+-----+-----
11  1 |  7 |  0 |    0
12  2 | -6 | 0.3 |   0.3
13  3 | -1 | 1.3 |   1.6
14  4 | 10 | 1.6 |   3.2
15  5 | 11 |  1 |   4.2
16  6 |  7 |  1 |   5.2
17 (6 rows)
18
```

Connected incomplete data

Using selected edges \{2, 4, 5, 8, 9, 15\} the matrix is not complete.

```
1 SELECT * FROM pgr_dijkstraCostMatrix(
```

```
2 $q1$SELECT id, source, target, cost, reverse_cost FROM edges WHERE id IN (2, 4, 5, 8, 9, 15)$q1$,
3 (SELECT ARRAY[6, 7, 10, 11, 16, 17]),
4 directed => true);
5 start_vid | end_vid | agg_cost
6 -----+-----+-----
7 6 | 7 | 1
8 6 | 11 | 2
9 6 | 16 | 3
10 6 | 17 | 4
11 7 | 6 | 1
12 7 | 11 | 1
13 7 | 16 | 2
14 7 | 17 | 3
15 10 | 6 | 1
16 10 | 7 | 2
17 10 | 11 | 1
18 10 | 16 | 2
19 10 | 17 | 3
20 11 | 6 | 2
21 11 | 7 | 1
22 11 | 16 | 1
23 11 | 17 | 2
24 16 | 6 | 3
25 16 | 7 | 2
26 16 | 11 | 1
27 16 | 17 | 1
28 17 | 6 | 4
29 17 | 7 | 3
30 17 | 11 | 2
31 17 | 16 | 1
32(25 rows)
33
```

Cost value for (17 → 10) do not exist on the matrix, but the value used is taken from (10 → 17).

```
1 SELECT * FROM pgr_TSP(
2 $$SELECT * FROM pgr_dijkstraCostMatrix(
3 $q1$SELECT id, source, target, cost, reverse_cost FROM edges WHERE id IN (2, 4, 5, 8, 9, 15)$q1$,
4 (SELECT ARRAY[6, 7, 10, 11, 16, 17]),
5 directed => true)$);
6 seq | node | cost | agg_cost
7 -----+-----+-----+-----
8 1 | 6 | 0 | 0
9 2 | 7 | 1 | 1
10 3 | 11 | 1 | 2
11 4 | 16 | 1 | 3
12 5 | 17 | 1 | 4
13 6 | 10 | 3 | 7
14 7 | 6 | 1 | 8
15(7 rows)
16
```

See Also

- [Traveling Sales Person - Family of functions](#)
- [Sample Data](#)
- [Boost: metric TSP approx](#)
- [Wikipedia: Traveling Salesman Problem](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_TSPeuclidean

- pgr\_TSPeuclidean - Approximation using *metric* algorithm.

Availability:

- Version 3.2.1
  - Using [Boost: metric TSP approx](#)
  - Simulated Annealing Algorithm no longer supported
    - The Simulated Annealing Algorithm related parameters are ignored: *max\_processing\_time*, *tries\_per\_temperature*, *max\_changes\_per\_temperature*, *max\_consecutive\_non\_changes*, *initial\_temperature*, *final\_temperature*, *cooling\_factor*, *randomize*
- Version 3.0.0
  - Name change from pgr\_euclidianTSP
- Version 2.3.0
  - New official function.

Description

Problem Definition

The travelling salesperson problem (TSP) asks the following question:

*Given a list of cities and the distances between each pair of cities, which is the shortest possible route that visits each city exactly once and returns to the origin city?*

Characteristics

- This problem is an NP-hard optimization problem.
- Metric Algorithm is used
- Implementation generates solutions that *are twice as long as the optimal tour in the worst case* when:
  - Graph is undirected
  - Graph is fully connected
  - Graph where traveling costs on edges obey the triangle inequality.
- On an undirected graph:
  - The traveling costs are symmetric:
  - Traveling costs from *u* to *v* are just as much as traveling from *v* to *u*

- Any duplicated identifier will be ignored. The coordinates that will be kept is arbitrarily.

- The coordinates are quite similar for the same identifier, for example

```
1, 3.5, 1
1, 3.499999999999 0.9999999
```

- The coordinates are quite different for the same identifier, for example

```
2, 3.5, 1.0
2, 3.6, 1.1
```

 Boost Graph Inside

Signatures

Summary

pgr\_TSPeuclidean([Coordinates SQL](#), [start\_id, end\_id])  
Returns set of (seq, node, cost, agg\_cost)  
OR EMPTY SET

Example:

With default values

```
SELECT * FROM pgr_TSPeuclidean(
$$
SELECT id, st_X(geom) AS x, st_Y(geom) AS y FROM vertices
$$);
seq | node | cost | agg_cost
-----+-----+-----+-----
1 | 1 | 0 | 0
2 | 6 | 2.2360679775 | 2.2360679775
3 | 5 | 1 | 3.2360679775
4 | 10 | 1.41421356237 | 4.65028153987
5 | 7 | 1.41421356237 | 6.06449510225
6 | 2 | 2.12132034356 | 8.18581544581
7 | 9 | 1.58113883008 | 9.76695427589
8 | 4 | 0.5 | 10.2669542759
9 | 14 | 1.58113883009 | 11.848093106
10 | 17 | 1.11803398875 | 12.9661270947
11 | 16 | 1 | 13.9661270947
12 | 15 | 1 | 14.9661270947
13 | 11 | 1.41421356237 | 16.3803406571
14 | 13 | 0.583095189485 | 16.9634358466
15 | 12 | 0.860232526704 | 17.8236683733
16 | 8 | 1 | 18.8236683733
17 | 3 | 1.41421356237 | 20.2378819357
18 | 1 | 1 | 21.2378819357
(18 rows)
```

Parameters

Parameter	Type	Description
-----------	------	-------------

[Coordinates SQL](#) TEXT [Coordinates SQL](#) as described below

TSP optional parameters

Column	Type	Default	Description
start_id	ANY-INTEGER	0	The first visiting vertex
			<ul style="list-style-type: none"><li>When 0 any vertex can become the first visiting vertex.</li></ul>
end_id	ANY-INTEGER	0	Last visiting vertex before returning to start_vid.
			<ul style="list-style-type: none"><li>When 0 any vertex can become the last visiting vertex before returning to start_id.</li></ul>
			<ul style="list-style-type: none"><li>When NOT 0 and start_id = 0 then it is the first and last vertex</li></ul>

Inner Queries

Coordinates SQL

Column	Type	Description
id	ANY-INTEGER	Identifier of the starting vertex.
x	ANY-NUMERICAL	X value of the coordinate.
y	ANY-NUMERICAL	Y value of the coordinate.

Result columns

Returns SET OF (seq, node, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Row sequence.
node	BIGINT	Identifier of the node/coordinate/point.



Column	Type	Description
cost	FLOAT	Cost to traverse from the current node to the next node in the path sequence.
		<ul style="list-style-type: none"><li>0 for the last row in the tour sequence.</li></ul>
agg_cost	FLOAT	Aggregate cost from the node at seq = 1 to the current node.
		<ul style="list-style-type: none"><li>0 for the first row in the tour sequence.</li></ul>

Additional Examples¶

- [Test 29 cities of Western Sahara](#)
  - [Creating a table for the data and storing the data](#)
  - [Adding a geometry \(for visual purposes\)](#)
  - [Total tour cost](#)
  - [Getting a geometry of the tour](#)
  - [Visual results](#)

[Test 29 cities of Western Sahara¶](#)

This example shows how to make performance tests using University of Waterloo's [example data](#) using the 29 cities of [Western Sahara dataset](#)

[Creating a table for the data and storing the data¶](#)

```
CREATE TABLE wi29 (id BIGINT, x FLOAT, y FLOAT, geom geometry);
INSERT INTO wi29 (id, x, y) VALUES
(1,20833.3333,17100.0000),
(2,20900.0000,17066.6667),
(3,21300.0000,13016.6667),
(4,21600.0000,14150.0000),
(5,21600.0000,14966.6667),
(6,21600.0000,16500.0000),
(7,22183.3333,13133.3333),
(8,22583.3333,14300.0000),
(9,22683.3333,12716.6667),
(10,23616.6667,15866.6667),
(11,23700.0000,15933.3333),
(12,23883.3333,14533.3333),
(13,24166.6667,13250.0000),
(14,25149.1667,12365.8333),
(15,26133.3333,14500.0000),
(16,26150.0000,10550.0000),
(17,26283.3333,12766.6667),
(18,26433.3333,13433.3333),
(19,26550.0000,13850.0000),
(20,26733.3333,11683.3333),
(21,27026.1111,13051.9444),
(22,27096.1111,13415.8333),
(23,27153.6111,13203.3333),
(24,27166.6667,9833.3333),
(25,27233.3333,10450.0000),
(26,27233.3333,11783.3333),
(27,27266.6667,10383.3333),
(28,27433.3333,12400.0000),
(29,27462.5000,12992.2222);
```

[Adding a geometry \(for visual purposes\)¶](#)

```
UPDATE wi29 SET geom = ST_makePoint(x,y);
```

[Total tour cost¶](#)

Getting a total cost of the tour, compare the value with the length of an optimal tour is 27603, given on the dataset

```
SELECT *
FROM pgr_TSPeuclidean($$SELECT * FROM wi29$$)
WHERE seq = 30;
seq | node | cost | agg_cost
-----+-----+-----+-----
30 | 1 | 2266.91173136 | 28777.4854127
(1 row)
```

[Getting a geometry of the tour¶](#)

```
WITH
tsp_results AS (SELECT seq, geom FROM pgr_TSPeuclidean($$SELECT * FROM wi29$$) JOIN wi29 ON (node = id))
SELECT ST_MakeLine(ARRAY(SELECT geom FROM tsp_results ORDER BY seq));
```

```
010200000001E000000F085C9545558D4400000000000B3D04000000000069D440107A36ABAAAAD04000000000018D54000000000001DD040107A36AB2A10D7401FF46C5655FDCE40000000000025D740E10B93A9AA1ECF40F085C954D5
(1 row)
```

[Visual results¶](#)

Visually, The first image is the [optimal solution](#) and the second image is the solution obtained with pgr\_TSPeuclidean.



- [Traveling Sales Person - Family of functions](#)
- [Sample Data](#)
- [Boost: metric TSP approx](#)
- [University of Waterloo TSP](#)
- [Wikipedia: Traveling Salesman Problem](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Table of Contents

- [General Information](#)
  - [Problem Definition](#)
  - [Origin](#)
  - [Characteristics](#)
  - [TSP optional parameters](#)
- [See Also](#)

[General Information¶](#)

[Problem Definition¶](#)

The travelling salesperson problem (TSP) asks the following question:

*Given a list of cities and the distances between each pair of cities, which is the shortest possible route that visits each city exactly once and returns to the origin city?*

[Origin¶](#)

The traveling sales person problem was studied in the 18th century by mathematicians **Sir William Rowam Hamilton** and **Thomas Penyngton Kirkman**.

A discussion about the work of Hamilton & Kirkman can be found in the book **Graph Theory (Biggs et al. 1976)**.

- ISBN-13: 978-0198539162
- ISBN-10: 0198539169

It is believed that the general form of the TSP have been first studied by Kair Menger in Vienna and Harvard. The problem was later promoted by Hassler, Whitney & Merrill at Princeton. A detailed description about the connection between Menger & Whitney, and the development of the TSP can be found in [On the history of combinatorial optimization \(till 1960\)](#)

To calculate the number of different tours through  $(n)$  cities:

- Given a starting city,
- There are  $(n-1)$  choices for the second city,
- And  $(n-2)$  choices for the third city, etc.
- Multiplying these together we get  $((n-1)! = (n-1) (n-2) \dots 1)$ .
- Now since the travel costs do not depend on the direction taken around the tour:
  - this number by 2
  - $((n-1)!/2)$ .

[Characteristics¶](#)

- This problem is an NP-hard optimization problem.
- Metric Algorithm is used
- Implementation generates solutions that *are twice as long as the optimal tour in the worst case* when:
  - Graph is undirected
  - Graph is fully connected
  - Graph where traveling costs on edges obey the triangle inequality.
- On an undirected graph:
  - The traveling costs are symmetric:
  - Traveling costs from  $u$  to  $v$  are just as much as traveling from  $v$  to  $u$

[TSP optional parameters¶](#)

Column	Type	Default	Description
start_id	ANY-INTEGER	0	The first visiting vertex
			<ul style="list-style-type: none"><li>• When 0 any vertex can become the first visiting vertex.</li></ul>
end_id	ANY-INTEGER	0	Last visiting vertex before returning to start_vid.
			<ul style="list-style-type: none"><li>• When 0 any vertex can become the last visiting vertex before returning to start_id.</li><li>• When NOT 0 and start_id = 0 then it is the first and last vertex</li></ul>

[See Also¶](#)

References

- [Boost: metric TSP approx](#)
- [University of Waterloo TSP](#)
- [Wikipedia: Traveling Salesman Problem](#)

Indices and tables

- [Index](#)
- [Search Page](#)

BFS - Category

- [pgr\\_kruskalBFS](#)
- [pgr\\_primBFS](#)

Traversal using breadth first search.

- It's implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.

Parameters

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below.
root_vid	BIGINT	Identifier of the root vertex of the tree. <ul style="list-style-type: none"><li>When value is \0 then gets the spanning forest starting in aleatory nodes for each tree in the forest.</li></ul>
root_vids	ARRAY [ <b>ANY-INTEGER</b> ]	Array of identifiers of the root vertices. <ul style="list-style-type: none"><li>\0 values are ignored</li><li>For optimization purposes, any duplicated value is ignored.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERIC:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT, NUMERIC

BFS optional parameters

Parameter	Type	Default	Description
max_depth	BIGINT	\(9223372036854775807\)	Upper limit of the depth of the tree. <ul style="list-style-type: none"><li>When negative throws an error.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	<b>ANY-INTEGER</b>		Identifier of the edge.
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
target	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
cost	<b>ANY-NUMERICAL</b>		Weight of the edge (source, target)
reverse_cost	<b>ANY-NUMERICAL</b>	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, depth, start\_vid, node, edge, cost, agg\_cost)

Parameter	Type	Description
seq	BIGINT	Sequential value starting from \((1\)).
depth	BIGINT	Depth of the node. <ul style="list-style-type: none"><li>\((0\)) when node = start_vid.</li></ul>
start_vid	BIGINT	Identifier of the root vertex.
node	BIGINT	Identifier of node reached using edge.
edge	BIGINT	Identifier of the edge used to arrive to node. <ul style="list-style-type: none"><li>\((-1\)) when node = start_vid.</li></ul>
cost	FLOAT	Cost to traverse edge.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERIC:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT, NUMERIC

See Also

- [Boost: Prim's algorithm](#)
- [Boost: Kruskal's algorithm](#)
- [Wikipedia: Prim's algorithm](#)
- [Wikipedia: Kruskal's algorithm](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Cost - Category

- [pgr\\_aStarCost](#)
- [pgr\\_bdAstarCost](#)
- [pgr\\_dijkstraCost](#)
- [pgr\\_bdDijkstraCost](#)
- [pgr\\_dijkstraNearCost - Proposed](#)

Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.
- [pgr\\_withPointsCost - Proposed](#)

General Information

Characteristics

Each function works as part of the family it belongs to.

The main Characteristics are:

- It does not return a path.
- Returns the sum of the costs of the shortest path of each pair combination of nodes requested.
- Let be the case the values returned are stored in a table, so the unique index would be the pair(start\_vid, end\_vid).
- Depending on the function and its parameters, the results can be symmetric.
  - The **aggregate cost** of \((u, v)\) is the same as for \((v, u)\).
- Any duplicated value in the start or end vertex identifiers are ignored.
- The returned values are ordered:

- start\_vid ascending
- end\_vid ascending

See Also

Indices and tables

- [Index](#)
- [Search Page](#)

Cost Matrix - Category

- [pgr\\_aStarCostMatrix](#)
- [pgr\\_dijkstraCostMatrix](#)
- [pgr\\_bdAstarCostMatrix](#)
- [pgr\\_bdDijkstraCostMatrix](#)

Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.
- [pgr\\_withPointsCostMatrix - proposed](#)

General Information

Synopsis

[Traveling Sales Person - Family of functions](#) needs as input a symmetric cost matrix and no edge  $(u, v)$  must value  $\infty$ .

This collection of functions will return a cost matrix in form of a table.

Characteristics

The main Characteristics are:

- Can be used as input to [pgr\\_TSP](#).
  - Use directly when the resulting matrix is symmetric and there is no  $\infty$  value.
  - It will be the users responsibility to make the matrix symmetric.
    - By using geometric or harmonic average of the non symmetric values.
    - By using max or min the non symmetric values.
    - By setting the upper triangle to be the mirror image of the lower triangle.
    - By setting the lower triangle to be the mirror image of the upper triangle.
  - It is also the users responsibility to fix an  $\infty$  value.
- Each function works as part of the family it belongs to.
- It does not return a path.
- Returns the sum of the costs of the shortest path for pair combination of nodes in the graph.
- Process is done only on edges with positive costs.
- Values are returned when there is a path.
  - When the starting vertex and ending vertex are the same, there is no path.
    - The aggregate cost in the non included values  $(v, v)$  is 0.
  - When the starting vertex and ending vertex are the different and there is no path.
    - The aggregate cost in the non included values  $(u, v)$  is  $\infty$ .
- Let be the case the values returned are stored in a table:
  - The unique index would be the pair: (start\_vid, end\_vid).
- Depending on the function and its parameters, the results can be symmetric.
  - The aggregate cost of  $(u, v)$  is the same as for  $(v, u)$ .
- Any duplicated value in the **start vids** are ignored.
- The returned values are ordered:
  - start\_vid ascending
  - end\_vid ascending

Parameters

Used in:

- [pgr\\_aStarCostMatrix](#)

- [pgr\\_dijkstraCostMatrix](#)

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below

**start vids** ARRAY[BIGINT] Array of identifiers of starting vertices.

Used in:

- [pgr\\_withPointsCostMatrix - proposed](#)

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below

[Points SQL](#) TEXT [Points SQL](#) as described below

**start vids** ARRAY[BIGINT] Array of identifiers of starting vertices.

Optional parameters<sup>¶</sup>

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li></ul>
			<ul style="list-style-type: none"><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries<sup>¶</sup>

Edges SQL<sup>¶</sup>

Used in:

- [pgr\\_withPointsCostMatrix - proposed](#)
- [pgr\\_dijkstraCostMatrix](#)

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Points SQL<sup>¶</sup>

Parameter	Type	Default	Description
pid	ANY-INTEGER	value	Identifier of the point. <ul style="list-style-type: none"><li>Use with positive value, as internally will be converted to negative value</li></ul>
			<ul style="list-style-type: none"><li>If column is present, it can not be NULL.</li></ul>
			<ul style="list-style-type: none"><li>If column is not present, a sequential negative <b>value</b> will be given automatically.</li></ul>
edge_id	ANY-INTEGER		Identifier of the "closest" edge to the point.
fraction	ANY-NUMERICAL		Value in <0,1> that indicates the relative position from the first end point of the edge.

Parameter	Type	Default	Description
			Value in [b, r, l, NULL] indicating if the point is:
side	CHAR	b	<ul style="list-style-type: none"> <li>In the right r,</li> <li>In the left l,</li> <li>In both sides b, NULL</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

See Also

- [Traveling Sales Person - Family of functions](#)

Indices and tables

- [Index](#)
- [Search Page](#)

DFS - Category

Traversal using Depth First Search.

- [pgr\\_kruskalDFS](#)
- [pgr\\_primDFS](#)

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.
- [pgr\\_depthFirstSearch - Proposed](#) - Depth first search traversal of the graph.

In general:

- It's implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.

See Also

- [Boost: Prim's algorithm](#)
- [Boost: Kruskal's algorithm](#)
- [Wikipedia: Prim's algorithm](#)
- [Wikipedia: Kruskal's algorithm](#)

Indices and tables

- [Index](#)
- [Search Page](#)

## Driving Distance - Category 1

- [pgr\\_drivingDistance](#) - Driving Distance based on Dijkstra's algorithm
- [pgr\\_primDD](#) - Driving Distance based on Prim's algorithm
- [pgr\\_kruskalDD](#) - Driving Distance based on Kruskal's algorithm
- Post processing
  - [pgr\\_alphaShape](#) - Alpha shape computation

☐ Proposed

## Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.
- [pgr\\_withPointsDD - Proposed](#) - Driving Distance based on pgr\_withPoints

## pgr\_alphaShape 1

pgr\_alphaShape — Polygon part of an alpha shape.

## Availability

- Version 3.8.0
  - Deprecated function.
- Version 3.0.0
  - Breaking change on signature
  - Old signature no longer supported
  - **Boost 1.54 & Boost 1.55** are supported
  - **Boost 1.56+** is preferable
    - Boost Geometry is stable on Boost 1.56
- Version 2.1.0
  - Added alpha argument with default 0 (use optimal value)
  - Support to return multiple outer/inner ring
- Version 2.0.0
  - New official function.
  - Renamed from version 1.x

## Migration of pgr\_alphaShape 1

Starting from [v3.8.0](#)

**Before Deprecation:** The following was calculated:

- An alphaShape was calculated

**After Deprecation:**

PostGIS has two ways of generating alphaShape.

If you have SFCGAL, which you can install using

```
CREATE EXTENSION postgis_sfcgal
```

- Since PostGIS 3.5+ use [CG\\_AlphaShape](#)
- For PostGIS 3.5+ use the old name ST\_AlphaShape

Other PostGIS options are \* [ST\\_ConvexHull](#) \* [ST\\_ConcaveHull](#)

## Description 1

Returns the polygon part of an alpha shape.

## Characteristics

- Input is a *geometry* and returns a *geometry*
- Uses PostGis ST\_DelaunyTriangles
- Instead of using CGAL's definition of *alpha* it use the `spoon_radius`
  - $\backslash(\text{spoon\_radius} = \backslash\text{sqrt alpha}\backslash)$
- A Triangle area is considered part of the alpha shape when  $\backslash(\text{circumcenter}\backslash \text{radius} < \text{spoon}\backslash\text{radius}\backslash)$
- The alpha parameter is the **spoon radius**
- When the total number of points is less than 3, returns an EMPTY geometry



Signatures

pgr\_alphaShape(**geometry**, [alpha])  
RETURNS geometry

Example:

passing a geometry collection with spoon radius  $\backslash(1.5\backslash)$  using the return variable geom

```
SELECT ST_Area(pgr_alphaShape((SELECT ST_Collect(geom)
FROM vertices), 1.5));
WARNING: pgr_alphashape(geometry,double precision) deprecated function on v3.8.0
st_area
-----
9.75
(1 row)
```

Parameters

Parameter	Type	Default	Description
<b>geometry</b>	geometry		Geometry with at least $\backslash(3\backslash)$ points
alpha	FLOAT	0	The radius of the spoon.

Return Value

Kind of geometry	Description
GEOMETRY COLLECTION	A Geometry collection of Polygons

See Also

- [pgr\\_drivingDistance](#)
- [Sample Data](#)
- [ST\\_ConcaveHull](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Calculate nodes that are within a distance.

- Extracts all the nodes that have costs less than or equal to the value distance.
- The edges extracted will conform to the corresponding spanning tree.
- Edge  $\backslash((u, v)\backslash)$  will not be included when:
  - The distance from the **root** to  $\backslash(u\backslash) >$  limit distance.
  - The distance from the **root** to  $\backslash(v\backslash) >$  limit distance.
  - No new nodes are created on the graph, so when is within the limit and is not within the limit, the edge is not included.

Parameters

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	Edges SQL as described below.
<b>Root vid</b>	BIGINT	Identifier of the root vertex of the tree.  Array of identifiers of the root vertices.
<b>Root vids</b>	ARRAY[ANY-INTEGER]	<ul style="list-style-type: none"><li><math>\backslash(0\backslash)</math> values are ignored</li><li>For optimization purposes, any duplicated value is ignored.</li></ul>
<b>distance</b>	FLOAT	Upper limit for the inclusion of a node in the result.

Where:

ANY-NUMERIC:  
  
SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Inner Queries

Edges SQL

Column	Type	Default	Description
id	<b>ANY-INTEGER</b>		Identifier of the edge.
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.

Column	Type	Default	Description
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns¶

Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Parameter	Type	Description
seq	BIGINT	Sequential value starting from \(\backslash\).
depth	BIGINT	Depth of the node. <ul style="list-style-type: none"> <li>\(0\backslash\) when node = start_vid.</li> <li>\(depth-1\backslash\) is the depth of pred</li> </ul>
start_vid	BIGINT	Identifier of the root vertex.
pred	BIGINT	Predecessor of node. <ul style="list-style-type: none"> <li>When node = start_vid then has the value node.</li> </ul>
node	BIGINT	Identifier of node reached using edge.
edge	BIGINT	Identifier of the edge used to arrive from pred to node. <ul style="list-style-type: none"> <li>\(-1\backslash\) when node = start_vid.</li> </ul>
cost	FLOAT	Cost to traverse edge.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

See Also¶

Indices and tables

- [Index](#)
- [Search Page](#)

K shortest paths - Category¶

- [pgr\\_KSP](#) - Yen's algorithm based on pgr\_dijkstra

❏ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.
- [pgr\\_withPointsKSP - Proposed](#) - Yen's algorithm based on pgr\_withPoints

Indices and tables

- [Index](#)
- [Search Page](#)

Spanning Tree - Category¶

- [Kruskal - Family of functions](#)
- [Prim - Family of functions](#)

A spanning tree of an undirected graph is a tree that includes all the vertices of G with the minimum possible number of edges.

For a disconnected graph, there is no single tree, but a spanning forest, consisting of a spanning tree of each connected component.

Characteristics:

- It's implementation is only on **undirected** graph.
- Process is done only on edges with positive costs.
- When the graph is connected
  - The resulting edges make up a tree
- When the graph is not connected,
  - Finds a minimum spanning tree for each connected component.
  - The resulting edges make up a forest.

See Also

- [Boost: Prim's algorithm](#)
- [Boost: Kruskal's algorithm](#)
- [Wikipedia: Prim's algorithm](#)
- [Wikipedia: Kruskal's algorithm](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Via - Category

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.
- [pgr\\_dijkstraVia - Proposed](#)
- [pgr\\_withPointsVia - Proposed](#)
- [pgr\\_trspVia - Proposed](#)
- [pgr\\_trspVia\\_withPoints - Proposed](#)

General Information

This category intends to solve the general problem:

Given a graph and a list of vertices, find the shortest path between  $(vertex_i)$  and  $(vertex_{i+1})$  for all vertices

In other words, find a continuous route that visits all the vertices in the order given.

path:

represents a section of a **route**.

route:

is a sequence of **paths**

Parameters

Used in:

- [pgr\\_dijkstraVia - Proposed](#)
- [pgr\\_trspVia - Proposed](#)

Parameter	Type	Default	Description
<a href="#">Edges.SQL</a>	TEXT		SQL query as described.
<b>via vertices</b>	ARRAY [ <b>ANY-INTEGER</b> ]		Array of ordered vertices identifiers that are going to be visited.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Used in:

- [pgr\\_withPointsVia - Proposed](#)
- [pgr\\_trspVia\\_withPoints - Proposed](#)

Parameter	Type	Default	Description
<a href="#">Edges SQL</a>	TEXT		SQL query as described.
<a href="#">Points SQL</a>	TEXT		SQL query as described.
<b>via vertices</b>	ARRAY [ <b>ANY-INTEGER</b> ]		Array of ordered vertices identifiers that are going to be visited. <ul style="list-style-type: none"><li>• When positive it is considered a vertex identifier</li><li>• When negative it is considered a point identifier</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Besides the compulsory parameters each function has, there are optional parameters that exist due to the kind of function.

Via optional parameters<sup>1</sup>

Used in all Via functions

Parameter	Type	Default	Description
strict	BOOLEAN	false	<ul style="list-style-type: none"><li>• When true if a path is missing stops and returns <b>EMPTY SET</b></li><li>• When false ignores missing paths returning all paths found</li></ul>
U_turn_on_edge	BOOLEAN	true	<ul style="list-style-type: none"><li>• When true departing from a visited vertex will not try to avoid</li></ul>

Inner Queries<sup>1</sup>

Depending on the function one or more inner queries are needed.

Edges SQL<sup>1</sup>

Used in all Via functions

Column	Type	Default	Description
id	<b>ANY-INTEGER</b>		Identifier of the edge.
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
target	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
cost	<b>ANY-NUMERICAL</b>		Weight of the edge (source, target)
reverse_cost	<b>ANY-NUMERICAL</b>	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Restrictions SQL<sup>1</sup>

Used in

- [pgr\\_trspVia - Proposed](#)

Column	Type	Description
path	ARRAY [ <b>ANY-INTEGER</b> ]	Sequence of edge identifiers that form a path that is not allowed to be taken. - Empty arrays on NULL arrays are ignored. - Arrays that have a NULL element will raise an exception.
Cost	<b>ANY-NUMERICAL</b>	Cost of taking the forbidden path.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Points SQL ¶

Used in

- [pgr\\_withPointsVia - Proposed](#)

Parameter	Type	Default	Description
			Identifier of the point. <ul style="list-style-type: none"><li>• Use with positive value, as internally will be converted to negative value</li><li>• If column is present, it can not be NULL.</li><li>• If column is not present, a sequential negative <b>value</b> will be given automatically.</li></ul>
pid	ANY-INTEGER	value	
edge_id	ANY-INTEGER		Identifier of the "closest" edge to the point.
fraction	ANY-NUMERICAL		Value in <0,1> that indicates the relative position from the first end point of the edge.
			Value in [b, r, l, NULL] indicating if the point is: <ul style="list-style-type: none"><li>• In the right r,</li><li>• In the left l,</li><li>• In both sides b, NULL</li></ul>
side	CHAR	b	

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns ¶

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Identifier of a path. Has value1 for the first path.
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex of the path.
end_vid	BIGINT	Identifier of the ending vertex of the path.
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.
		Identifier of the edge used to go fromnode to the next node in the path sequence. <ul style="list-style-type: none"><li>• -1 for the last node of the path.</li><li>• -2 for the last node of the route.</li></ul>
edge	BIGINT	
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.
route_agg_cost	FLOAT	Total cost from start_vid of seq = 1 to end_vid of the currentseq.

Note

When start\_vid, end\_vid and node columns have negative values, the identifier is for a Point.

See Also ¶

- [pgr\\_dijkstraVia - Proposed](#)
- [pgr\\_trspVia - Proposed](#)
- [pgr\\_withPointsVia - Proposed](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Vehicle Routing Functions - Category ¶

☐ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting
- Pickup and delivery problem
  - [pgr\\_pickDeliver - Experimental](#) - Pickup & Delivery using a Cost Matrix
  - [pgr\\_pickDeliverEuclidean - Experimental](#) - Pickup & Delivery with Euclidean distances
- Distribution problem
  - [pgr\\_vrpOneDepot - Experimental](#) - From a single depot, distributes orders

Contents

- [Vehicle Routing Functions - Category](#)
  - [Introduction](#)
    - [Characteristics](#)
  - [Pick & Delivery](#)
  - [Parameters](#)
    - [Pick & deliver](#)
    - [Pick-Deliver optional parameters](#)
  - [Inner Queries](#)
    - [Orders SQL](#)
    - [Vehicles SQL](#)
    - [Matrix SQL](#)
  - [Result columns](#)
    - [Summary Row](#)
  - [Handling Parameters](#)
    - [Capacity and Demand Units Handling](#)
    - [Locations](#)
    - [Time Handling](#)
    - [Factor handling](#)
  - [See Also](#)

[pgr\\_pickDeliver - Experimental](#)

pgr\_pickDeliver - Pickup and delivery Vehicle Routing Problem

☐ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.

- Might need c/c++ coding.
- May lack documentation.
- Documentation if any might need to be rewritten.
- Documentation examples might need to be automatically generated.
- Might need a lot of feedback from the community.
- Might depend on a proposed function of pgRouting
- Might depend on a deprecated function of pgRouting

#### Availability

- Version 3.0.0
  - New experimental function.

#### Synopsis

Problem: Distribute and optimize the pickup-delivery pairs into a fleet of vehicles.

- Optimization problem is NP-hard.
- pickup and Delivery with time windows.
- All vehicles are equal.
  - Same Starting location.
  - Same Ending location which is the same as Starting location.
  - All vehicles travel at the same speed.
- A customer is for doing a pickup or doing a deliver.
  - has an open time.
  - has a closing time.
  - has a service time.
  - has an (x, y) location.
- There is a customer where to deliver a pickup.
  - travel time between customers is distance / speed
  - pickup and delivery pair is done with the same vehicle.
  - A pickup is done before the delivery.

#### Characteristics

- All trucks depart at time 0.
- No multiple time windows for a location.
- Less vehicle used is considered better.
- Less total duration is better.
- Less wait time is better.
- the algorithm will raise an exception when
  - If there is a pickup-deliver pair than violates time window
  - The speed, max\_cycles, ma\_capacity have illegal values
- Six different initial will be optimized - the best solution found will be result

#### Signature

`pgr_pickDeliver(Orders SQL, Vehicles SQL, Matrix SQL, [options])`

**options:** {factor, max\_cycles, initial\_sol}

Returns set of (seq, vehicle\_number, vehicle\_id, stop, order\_id, stop\_type, cargo, travel\_time, arrival\_time, wait\_time, service\_time, departure\_time)

Example:

Solve the following problem

Given the vehicles:

```
SELECT id, capacity, start_node_id, start_open, start_close
FROM vehicles;
id | capacity | start_node_id | start_open | start_close
-----+-----+-----+-----+-----
1 | 50 | 11 | 0 | 50
2 | 50 | 11 | 0 | 50
(2 rows)
```

and the orders:

```
SELECT id, demand,
       p_node_id, p_open, p_close, p_service,
       d_node_id, d_open, d_close, d_service
FROM orders;
id | demand | p_node_id | p_open | p_close | p_service | d_node_id | d_open | d_close | d_service
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 10 | 10 | 2 | 10 | 3 | 3 | 6 | 15 | 3
2 | 20 | 16 | 4 | 15 | 2 | 15 | 6 | 20 | 3
3 | 30 | 7 | 2 | 10 | 3 | 12 | 3 | 20 | 3
(3 rows)
```

The query:

```
SELECT * FROM pgr_pickDeliver(
  $$SELECT id, demand,
    p_node_id, p_open, p_close, p_service,
    d_node_id, d_open, d_close, d_service
  FROM orders$$,
  $$SELECT id, capacity, start_node_id, start_open, start_close
  FROM vehicles$$,
  $$SELECT * from pgr_dijkstraCostMatrix(
    'SELECT * FROM edges',
```

```
(SELECT array_agg(id) FROM (SELECT p_node_id AS id FROM orders
UNION
SELECT d_node_id FROM orders
UNION
SELECT start_node_id FROM vehicles) a))
$$);
seq | vehicle_seq | vehicle_id | stop_seq | stop_type | stop_id | order_id | cargo | travel_time | arrival_time | wait_time | service_time | departure_time
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | 1 | 11 | -1 | 0 | 0 | 0 | 0 | 0 | 0
2 | 1 | 1 | 1 | 2 | 7 | 3 | 30 | 1 | 1 | 1 | 3 | 5
3 | 1 | 1 | 1 | 3 | 12 | 3 | 0 | 2 | 7 | 0 | 3 | 10
4 | 1 | 1 | 1 | 4 | 2 | 16 | 2 | 20 | 1 | 12 | 0 | 14
5 | 1 | 1 | 1 | 5 | 3 | 15 | 2 | 0 | 1 | 15 | 0 | 18
6 | 1 | 1 | 1 | 6 | 6 | 11 | -1 | 0 | 2 | 20 | 0 | 20
7 | 2 | 2 | 1 | 1 | 11 | -1 | 0 | 0 | 0 | 0 | 0 | 0
8 | 2 | 2 | 2 | 2 | 10 | 1 | 10 | 3 | 3 | 0 | 3 | 6
9 | 2 | 2 | 3 | 3 | 3 | 1 | 0 | 3 | 9 | 0 | 3 | 12
10 | 2 | 2 | 4 | 6 | 11 | -1 | 0 | 2 | 14 | 0 | 0 | 14
11 | -2 | 0 | 0 | -1 | -1 | -1 | -1 | 16 | -1 | 1 | 17 | 34
(11 rows)
```

Parameters

The parameters are:

Column	Type	Description
--------	------	-------------

[Orders SQL](#) TEXT [Orders SQL](#) as described below.

[Vehicles SQL](#) TEXT [Vehicles SQL](#) as described below.

[Matrix SQL](#) TEXT [Matrix SQL](#) as described below.

Pick-Deliver optional parameters

Column	Type	Default	Description
factor	NUMERIC	1	Travel time multiplier. See <a href="#">Factor handling</a>
max_cycles	INTEGER	10	Maximum number of cycles to perform on the optimization.
			Initial solution to be used.
			<ul style="list-style-type: none"><li>1 One order per truck</li><li>2 Push front order.</li><li>3 Push back order.</li></ul>
initial_sol	INTEGER	4	<ul style="list-style-type: none"><li>4 Optimize insert.</li><li>5 Push back order that allows more orders to be inserted at the back</li><li>6 Push front order that allows more orders to be inserted at the front</li></ul>

Orders SQL

A *SELECT* statement that returns the following columns:

id, demand  
p\_node\_id, p\_open, p\_close, [p\_service,]  
d\_node\_id, d\_open, d\_close, [d\_service,]

where:

Column	Type	Description
id	ANY-INTEGER	Identifier of the pick-delivery order pair.
demand	ANY-NUMERICAL	Number of units in the order
p_open	ANY-NUMERICAL	The time, relative to 0, when the pickup location opens.
p_close	ANY-NUMERICAL	The time, relative to 0, when the pickup location closes.
[p_service]	ANY-NUMERICAL	The duration of the loading at the pickup location. <ul style="list-style-type: none"><li>When missing: 0 time units are used</li></ul>
d_open	ANY-NUMERICAL	The time, relative to 0, when the delivery location opens.
d_close	ANY-NUMERICAL	The time, relative to 0, when the delivery location closes.
[d_service]	ANY-NUMERICAL	The duration of the unloading at the delivery location. <ul style="list-style-type: none"><li>When missing: 0 time units are used</li></ul>

Where:



ANY-INTEGGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Column	Type	Description
p_node_id	ANY-INTEGGER	The node identifier of the pickup, must match a vertex identifier in the <a href="#">Matrix SQL</a> .
d_node_id	ANY-INTEGGER	The node identifier of the delivery, must match a vertex identifier in the <a href="#">Matrix SQL</a> .

Where:

ANY-INTEGGER:

SMALLINT, INTEGER, BIGINT

Vehicles SQL

A *SELECT* statement that returns the following columns:

id, capacity  
start\_node\_id, start\_open, start\_close [, start\_service,]  
[end\_node\_id, end\_open, end\_close, end\_service]

where:

Column	Type	Description
id	ANY-NUMERICAL	Identifier of the vehicle.
capacity	ANY-NUMERICAL	Maiximum capacity units
start_open	ANY-NUMERICAL	The time, relative to 0, when the starting location opens.
start_close	ANY-NUMERICAL	The time, relative to 0, when the starting location closes.
[start_service]	ANY-NUMERICAL	The duration of the loading at the starting location. <ul style="list-style-type: none"><li>When missing: A duration of \0\ time units is used.</li></ul>
[end_open]	ANY-NUMERICAL	The time, relative to 0, when the ending location opens. <ul style="list-style-type: none"><li>When missing: The value of start_open is used</li></ul>
[end_close]	ANY-NUMERICAL	The time, relative to 0, when the ending location closes. <ul style="list-style-type: none"><li>When missing: The value of start_close is used</li></ul>
[end_service]	ANY-NUMERICAL	The duration of the loading at the ending location. <ul style="list-style-type: none"><li>When missing: A duration in start_service is used.</li></ul>

Column	Type	Description
start_node_id	ANY-INTEGGER	The node identifier of the start location, must match a vertex identifier in the <a href="#">Matrix SQL</a> .
[end_node_id]	ANY-INTEGGER	The node identifier of the end location, must match a vertex identifier in the <a href="#">Matrix SQL</a> . <ul style="list-style-type: none"><li>When missing: end_node_id is used.</li></ul>

Where:

ANY-INTEGGER:

SMALLINT, INTEGER, BIGINT

Matrix SQL

Where:

ANY-INTEGGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of  
(seq, vehicle\_seq, vehicle\_id, stop\_seq, stop\_type,  
travel\_time, arrival\_time, wait\_time, service\_time, departure\_time)  
UNION  
(summary row)

Column	Type	Description
--------	------	-------------

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
vehicle_seq	INTEGER	Sequential value starting from 1 for current vehicles. The $\backslash(n_{th})$ vehicle in the solution. <ul style="list-style-type: none"> <li>Value <math>\backslash(-2)</math> indicates it is the summary row.</li> </ul>
vehicle_id	BIGINT	Current vehicle identifier. <ul style="list-style-type: none"> <li>Summary row has the <b>total capacity violations</b>. <ul style="list-style-type: none"> <li>A capacity violation happens when overloading or underloading a vehicle.</li> </ul> </li> </ul>
stop_seq	INTEGER	Sequential value starting from 1 for the stops made by the current vehicle. The $\backslash(m_{th})$ stop of the current vehicle. <ul style="list-style-type: none"> <li>Summary row has the <b>total time windows violations</b>. <ul style="list-style-type: none"> <li>A time window violation happens when arriving after the location has closed.</li> </ul> </li> </ul>
stop_type	INTEGER	<ul style="list-style-type: none"> <li>Kind of stop location the vehicle is at <ul style="list-style-type: none"> <li><math>\backslash(-1)</math>: at the solution summary row</li> <li><math>\backslash(1)</math>: Starting location</li> <li><math>\backslash(2)</math>: Pickup location</li> <li><math>\backslash(3)</math>: Delivery location</li> <li><math>\backslash(6)</math>: Ending location and indicates the vehicle's summary row</li> </ul> </li> </ul>
order_id	BIGINT	Pickup-Delivery order pair identifier. <ul style="list-style-type: none"> <li>Value <math>\backslash(-1)</math>: When no order is involved on the current stop location.</li> </ul>
cargo	FLOAT	Cargo units of the vehicle when leaving the stop. <ul style="list-style-type: none"> <li>Value <math>\backslash(-1)</math> on solution summary row.</li> </ul>
travel_time	FLOAT	Travel time from previous stop_seq to current stop_seq. <ul style="list-style-type: none"> <li>Summary has the <b>total traveling time</b>: <ul style="list-style-type: none"> <li>The sum of all the travel_time.</li> </ul> </li> </ul>
arrival_time	FLOAT	Time spent waiting for current location to open. <ul style="list-style-type: none"> <li><math>\backslash(-1)</math>: at the solution summary row.</li> <li><math>\backslash(0)</math>: at the starting location.</li> </ul>
wait_time	FLOAT	Time spent waiting for current location to open. <ul style="list-style-type: none"> <li>Summary row has the <b>total waiting time</b>: <ul style="list-style-type: none"> <li>The sum of all the wait_time.</li> </ul> </li> </ul>
service_time	FLOAT	Service duration at current location. <ul style="list-style-type: none"> <li>Summary row has the <b>total service time</b>: <ul style="list-style-type: none"> <li>The sum of all the service_time.</li> </ul> </li> <li>The time at which the vehicle departs from the stop. <ul style="list-style-type: none"> <li><math>\backslash(arrival\_time + wait\_time + service\_time)</math></li> </ul> </li> </ul>
departure_time	FLOAT	<ul style="list-style-type: none"> <li>The ending location has the <b>total time</b> used by the current vehicle.</li> <li>Summary row has the <b>total solution time</b>: <ul style="list-style-type: none"> <li><math>\backslash(total\ traveling\ time + total\ waiting\ time + total\ service\ time)</math></li> </ul> </li> </ul>

See Also

- [Vehicle Routing Functions - Category](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_pickDeliverEuclidean - Experimental

pgr\_pickDeliverEuclidean - Pickup and delivery Vehicle Routing Problem

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.0.0
  - Replaces pgr\_gsoc\_vrppdtw
  - New experimental function.

Synopsis

Problem: Distribute and optimize the pickup-delivery pairs into a fleet of vehicles.

- Optimization problem is NP-hard.
- Pickup and Delivery:
  - capacitated
  - with time windows.
- The vehicles
  - have (x, y) start and ending locations.
  - have a start and ending service times.
  - have opening and closing times for the start and ending locations.
- An order is for doing a pickup and a a deliver.
  - has (x, y) pickup and delivery locations.
  - has opening and closing times for the pickup and delivery locations.
  - has a pickup and deliver service times.
- There is a customer where to deliver a pickup.
  - travel time between customers is distance / speed
  - pickup and delivery pair is done with the same vehicle.
  - A pickup is done before the delivery.

Characteristics

- No multiple time windows for a location.
- Less vehicle used is considered better.
- Less total duration is better.
- Less wait time is better.
- Six different optional different initial solutions
  - the best solution found will be result

Signature

pgr\_pickDeliverEuclidean([Orders SQL](#), [Vehicles SQL](#), [options](#))

**options:** [factor, max\_cycles, initial\_sol]

Returns set of (seq, vehicle\_number, vehicle\_id, stop, order\_id, stop\_type, cargo, travel\_time, arrival\_time, wait\_time, service\_time, departure\_time)

Example:

Solve the following problem

Given the vehicles:

```
SELECT id, capacity, start_x, start_y, start_open, start_close
FROM vehicles;
id | capacity | start_x | start_y | start_open | start_close
-----+-----+-----+-----+-----+-----
1 | 50 | 3 | 2 | 0 | 50
2 | 50 | 3 | 2 | 0 | 50
(2 rows)
```

and the orders:

```
SELECT id, demand,
       p_x, p_y, p_open, p_close, p_service,
       d_x, d_y, d_open, d_close, d_service
FROM orders;
id | demand | p_x | p_y | p_open | p_close | p_service | d_x | d_y | d_open | d_close | d_service
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----

```

1	10	3	1	2	10	3	1	2	6	15	3
2	20	4	2	4	15	2	4	1	6	20	3
3	30	2	2	2	10	3	3	3	3	20	3

(3 rows)

The query:

```
SELECT * FROM pgr_pickDeliverEuclidean(
  $$SELECT id, demand,
    p_x, p_y, p_open, p_close, p_service,
    d_x, d_y, d_open, d_close, d_service
  FROM orders$$,
  $$SELECT id, capacity, start_x, start_y, start_open, start_close
  FROM vehicles$$);
seq | vehicle_seq | vehicle_id | stop_seq | stop_type | order_id | cargo | travel_time | arrival_time | wait_time | service_time | departure_time
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | 1 | 1 | -1 | 0 | 0 | 0 | 0 | 0 | 0
2 | 1 | 1 | 1 | 2 | 2 | 3 | 30 | 1 | 1 | 1 | 3 | 5
3 | 1 | 1 | 1 | 3 | 3 | 3 | 0 | 1.41421356237 | 6.41421356237 | 0 | 3 | 9.41421356237
4 | 1 | 1 | 1 | 4 | 2 | 2 | 20 | 1.41421356237 | 10.8284271247 | 0 | 2 | 12.8284271247
5 | 1 | 1 | 1 | 5 | 3 | 2 | 0 | 1 | 13.8284271247 | 0 | 3 | 16.8284271247
6 | 1 | 1 | 1 | 6 | 6 | -1 | 0 | 1.41421356237 | 18.2426406871 | 0 | 0 | 18.2426406871
7 | 2 | 2 | 1 | 1 | 1 | -1 | 0 | 0 | 0 | 0 | 0 | 0
8 | 2 | 2 | 2 | 2 | 1 | 10 | 1 | 1 | 1 | 3 | 5
9 | 2 | 2 | 3 | 3 | 1 | 0 | 2.2360679775 | 7.2360679775 | 0 | 3 | 10.2360679775
10 | 2 | 2 | 2 | 4 | 6 | -1 | 0 | 2 | 12.2360679775 | 0 | 0 | 12.2360679775
11 | -2 | 0 | 0 | -1 | -1 | -1 | 11.4787086646 | -1 | 2 | 17 | 30.4787086646
(11 rows)
```

Parameters

Column	Type	Description
--------	------	-------------

[Orders SQL](#) TEXT [Orders SQL](#) as described below.

[Vehicles SQL](#) TEXT [Vehicles SQL](#) as described below.

Pick-Deliver optional parameters

Column	Type	Default	Description
factor	NUMERIC	1	Travel time multiplier. See <a href="#">Factor handling</a>
max_cycles	INTEGER	10	Maximum number of cycles to perform on the optimization.
initial_sol	INTEGER	4	Initial solution to be used. <ul style="list-style-type: none"><li>1 One order per truck</li><li>2 Push front order.</li><li>3 Push back order.</li><li>4 Optimize insert.</li><li>5 Push back order that allows more orders to be inserted at the back</li><li>6 Push front order that allows more orders to be inserted at the front</li></ul>

Orders SQL

A *SELECT* statement that returns the following columns:

```
id, demand
p_x, p_y, p_open, p_close, [p_service]
d_x, d_y, d_open, d_close, [d_service]
```

Where:

Column	Type	Description
id	ANY-INTEGER	Identifier of the pick-delivery order pair.
demand	ANY-NUMERICAL	Number of units in the order
p_open	ANY-NUMERICAL	The time, relative to 0, when the pickup location opens.
p_close	ANY-NUMERICAL	The time, relative to 0, when the pickup location closes.
[p_service]	ANY-NUMERICAL	The duration of the loading at the pickup location. <ul style="list-style-type: none"><li>When missing: 0 time units are used</li></ul>
d_open	ANY-NUMERICAL	The time, relative to 0, when the delivery location opens.
d_close	ANY-NUMERICAL	The time, relative to 0, when the delivery location closes.
[d_service]	ANY-NUMERICAL	The duration of the unloading at the delivery location. <ul style="list-style-type: none"><li>When missing: 0 time units are used</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Column	Type	Description
p_x	ANY-NUMERICAL	\(x\) value of the pick up location
p_y	ANY-NUMERICAL	\(y\) value of the pick up location
d_x	ANY-NUMERICAL	\(x\) value of the delivery location
d_y	ANY-NUMERICAL	\(y\) value of the delivery location

Where:

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Vehicles SQL

A *SELECT* statement that returns the following columns:

id, capacity  
start\_x, start\_y, start\_open, start\_close [, start\_service, ]  
[ end\_x, end\_y, end\_open, end\_close, end\_service ]

where:

Column	Type	Description
id	ANY-NUMERICAL	Identifier of the vehicle.
capacity	ANY-NUMERICAL	Maiximum capacity units
start_open	ANY-NUMERICAL	The time, relative to 0, when the starting location opens.
start_close	ANY-NUMERICAL	The time, relative to 0, when the starting location closes.
[start_service]	ANY-NUMERICAL	The duration of the loading at the starting location. <ul style="list-style-type: none"><li>When missing: A duration of \(\) time units is used.</li></ul>
[end_open]	ANY-NUMERICAL	The time, relative to 0, when the ending location opens. <ul style="list-style-type: none"><li>When missing: The value of start_open is used</li></ul>
[end_close]	ANY-NUMERICAL	The time, relative to 0, when the ending location closes. <ul style="list-style-type: none"><li>When missing: The value of start_close is used</li></ul>
[end_service]	ANY-NUMERICAL	The duration of the loading at the ending location. <ul style="list-style-type: none"><li>When missing: A duration in start_service is used.</li></ul>

Column	Type	Description
start_x	ANY-NUMERICAL	\(x\) value of the starting location
start_y	ANY-NUMERICAL	\(y\) value of the starting location
[end_x]	ANY-NUMERICAL	\(x\) value of the ending location <ul style="list-style-type: none"><li>When missing: start_x is used.</li></ul>
[end_y]	ANY-NUMERICAL	\(y\) value of the ending location <ul style="list-style-type: none"><li>When missing: start_y is used.</li></ul>

Where:

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of  
(seq, vehicle\_seq, vehicle\_id, stop\_seq, stop\_type,  
travel\_time, arrival\_time, wait\_time, service\_time, departure\_time)  
UNION  
(summary row)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
vehicle_seq	INTEGER	Sequential value starting from 1 for current vehicles. The $\backslash(n_{th})$ vehicle in the solution. <ul style="list-style-type: none"> <li>Value <math>\backslash(-2)</math> indicates it is the summary row.</li> </ul> Current vehicle identifier.
vehicle_id	BIGINT	<ul style="list-style-type: none"> <li>Summary row has the <b>total capacity violations</b>. <ul style="list-style-type: none"> <li>A capacity violation happens when overloading or underloading a vehicle.</li> </ul> </li> </ul> Sequential value starting from 1 for the stops made by the current vehicle. The $\backslash(m_{th})$ stop of the current vehicle.
stop_seq	INTEGER	<ul style="list-style-type: none"> <li>Summary row has the <b>total time windows violations</b>. <ul style="list-style-type: none"> <li>A time window violation happens when arriving after the location has closed.</li> </ul> </li> <li>Kind of stop location the vehicle is at <ul style="list-style-type: none"> <li><math>\backslash(-1)</math>: at the solution summary row</li> </ul> </li> </ul>
stop_type	INTEGER	<ul style="list-style-type: none"> <li><math>\backslash(1)</math>: Starting location</li> <li><math>\backslash(2)</math>: Pickup location</li> <li><math>\backslash(3)</math>: Delivery location</li> <li><math>\backslash(6)</math>: Ending location and indicates the vehicle's summary row</li> </ul>
order_id	BIGINT	Pickup-Delivery order pair identifier. <ul style="list-style-type: none"> <li>Value <math>\backslash(-1)</math>: When no order is involved on the current stop location.</li> </ul>
cargo	FLOAT	Cargo units of the vehicle when leaving the stop. <ul style="list-style-type: none"> <li>Value <math>\backslash(-1)</math> on solution summary row.</li> </ul>
travel_time	FLOAT	Travel time from previous stop_seq to current stop_seq. <ul style="list-style-type: none"> <li>Summary has the <b>total traveling time</b>: <ul style="list-style-type: none"> <li>The sum of all the travel_time.</li> </ul> </li> </ul>
arrival_time	FLOAT	Time spent waiting for current location to open. <ul style="list-style-type: none"> <li><math>\backslash(-1)</math>: at the solution summary row.</li> <li><math>\backslash(0)</math>: at the starting location.</li> </ul>
wait_time	FLOAT	Time spent waiting for current location to open. <ul style="list-style-type: none"> <li>Summary row has the <b>total waiting time</b>: <ul style="list-style-type: none"> <li>The sum of all the wait_time.</li> </ul> </li> </ul>
service_time	FLOAT	Service duration at current location. <ul style="list-style-type: none"> <li>Summary row has the <b>total service time</b>: <ul style="list-style-type: none"> <li>The sum of all the service_time.</li> </ul> </li> <li>The time at which the vehicle departs from the stop. <ul style="list-style-type: none"> <li><math>\backslash(arrival\_time + wait\_time + service\_time)</math></li> </ul> </li> </ul>
departure_time	FLOAT	<ul style="list-style-type: none"> <li>The ending location has the <b>total time</b> used by the current vehicle.</li> <li>Summary row has the <b>total solution time</b>: <ul style="list-style-type: none"> <li><math>\backslash(total\ traveling\ time + total\ waiting\ time + total\ service\ time)</math></li> </ul> </li> </ul>

#### Example¶

- [The vehicles](#)
- [The original orders](#)
- [The orders](#)
- [The query](#)

This data example **lc101** is from data published at <https://www.sintef.no/projectweb/top/pdptw/li-lim-benchmark/>

#### The vehicles¶

There are 25 vehicles in the problem all with the same characteristics.

```
CREATE TABLE v_lc101(
  id BIGINT NOT NULL primary key,
  capacity BIGINT DEFAULT 200,
  start_x FLOAT DEFAULT 30,
  start_y FLOAT DEFAULT 50,
  start_open INTEGER DEFAULT 0,
  start_close INTEGER DEFAULT 1236);
CREATE TABLE
/* create 25 vehicles */
INSERT INTO v_lc101 (id)
(SELECT * FROM generate_series(1, 25));
INSERT 0 25
```

[The original orders¶](#)

The data comes in different rows for the pickup and the delivery of the same order.

```
CREATE table lc101_c(
  id BIGINT not null primary key,
  x DOUBLE PRECISION,
  y DOUBLE PRECISION,
  demand INTEGER,
  open INTEGER,
  close INTEGER,
  service INTEGER,
  pindex BIGINT,
  dindex BIGINT
);
CREATE TABLE
/* the original data */
INSERT INTO lc101_c(
  id, x, y, demand, open, close, service, pindex, dindex) VALUES
( 1, 45, 68, -10, 912, 967, 90, 11, 0),
( 2, 45, 70, -20, 825, 870, 90, 6, 0),
( 3, 42, 66, 10, 65, 146, 90, 0, 75),
( 4, 42, 68, -10, 727, 782, 90, 9, 0),
( 5, 42, 65, 10, 15, 67, 90, 0, 7),
( 6, 40, 69, 20, 621, 702, 90, 0, 2),
( 7, 40, 66, -10, 170, 225, 90, 5, 0),
( 8, 38, 68, 20, 255, 324, 90, 0, 10),
( 9, 38, 70, 10, 534, 605, 90, 0, 4),
(10, 35, 66, -20, 357, 410, 90, 8, 0),
(11, 35, 69, 10, 448, 505, 90, 0, 1),
(12, 25, 85, -20, 652, 721, 90, 18, 0),
(13, 22, 75, 30, 30, 92, 90, 0, 17),
(14, 22, 85, -40, 567, 620, 90, 16, 0),
(15, 20, 80, -10, 384, 429, 90, 19, 0),
(16, 20, 85, 40, 475, 528, 90, 0, 14),
(17, 18, 75, -30, 99, 148, 90, 13, 0),
(18, 15, 75, 20, 179, 254, 90, 0, 12),
(19, 15, 80, 10, 278, 345, 90, 0, 15),
(20, 30, 50, 10, 10, 73, 90, 0, 24),
(21, 30, 52, -10, 914, 965, 90, 30, 0),
(22, 28, 52, -20, 812, 883, 90, 28, 0),
(23, 28, 55, 10, 732, 777, 0, 0, 103),
(24, 25, 50, -10, 65, 144, 90, 20, 0),
(25, 25, 52, 40, 169, 224, 90, 0, 27),
(26, 25, 55, -10, 622, 701, 90, 29, 0),
(27, 23, 52, -40, 261, 316, 90, 25, 0),
(28, 23, 55, 20, 546, 593, 90, 0, 22),
(29, 20, 50, 10, 358, 405, 90, 0, 26),
(30, 20, 55, 10, 449, 504, 90, 0, 21),
(31, 10, 35, -30, 200, 237, 90, 32, 0),
(32, 10, 40, 30, 31, 100, 90, 0, 31),
(33, 8, 40, 40, 87, 158, 90, 0, 37),
(34, 8, 45, -30, 751, 816, 90, 38, 0),
(35, 5, 35, 10, 283, 344, 90, 0, 39),
(36, 5, 45, 10, 665, 716, 0, 0, 105),
(37, 2, 40, -40, 383, 434, 90, 33, 0),
(38, 0, 40, 30, 479, 522, 90, 0, 34),
(39, 0, 45, -10, 567, 624, 90, 35, 0),
(40, 35, 30, -20, 264, 321, 90, 42, 0),
(41, 35, 32, -10, 166, 235, 90, 43, 0),
(42, 33, 32, 20, 68, 149, 90, 0, 40),
(43, 33, 35, 10, 16, 80, 90, 0, 41),
(44, 32, 30, 10, 359, 412, 90, 0, 46),
(45, 30, 30, 10, 541, 600, 90, 0, 48),
(46, 30, 32, -10, 448, 509, 90, 44, 0),
(47, 30, 35, -10, 1054, 1127, 90, 49, 0),
(48, 28, 30, -10, 632, 693, 90, 45, 0),
(49, 28, 35, 10, 1001, 1066, 90, 0, 47),
(50, 26, 32, 10, 815, 880, 90, 0, 52),
(51, 25, 30, 10, 725, 786, 0, 0, 101),
(52, 25, 35, -10, 912, 969, 90, 50, 0),
(53, 44, 5, 20, 286, 347, 90, 0, 58),
(54, 42, 10, 40, 186, 257, 90, 0, 60),
(55, 42, 15, -40, 95, 158, 90, 57, 0),
(56, 40, 5, 30, 385, 436, 90, 0, 59),
(57, 40, 15, 40, 35, 87, 90, 0, 55),
(58, 38, 5, -20, 471, 534, 90, 53, 0),
(59, 38, 15, -30, 651, 740, 90, 56, 0),
(60, 35, 5, -40, 562, 629, 90, 54, 0),
(61, 50, 30, -10, 531, 610, 90, 67, 0),
(62, 50, 35, 20, 262, 317, 90, 0, 68),
(63, 50, 40, 50, 171, 218, 90, 0, 74),
(64, 48, 30, 10, 632, 693, 0, 0, 102),
(65, 48, 40, 10, 76, 129, 90, 0, 72),
(66, 47, 35, 10, 826, 875, 90, 0, 69),
(67, 47, 40, 10, 12, 77, 90, 0, 61),
(68, 45, 30, -20, 734, 777, 90, 62, 0),
(69, 45, 35, -10, 916, 969, 90, 66, 0),
(70, 95, 30, -30, 387, 456, 90, 81, 0),
(71, 95, 35, 20, 293, 360, 90, 0, 77),
(72, 53, 30, -10, 450, 505, 90, 65, 0),
(73, 92, 30, -10, 478, 551, 90, 76, 0),
(74, 53, 35, -50, 353, 412, 90, 63, 0),
(75, 45, 65, -10, 997, 1068, 90, 3, 0),
(76, 90, 35, 10, 203, 260, 90, 0, 73),
(77, 88, 30, -20, 574, 643, 90, 71, 0),
(78, 88, 35, 20, 109, 170, 0, 0, 104),
(79, 87, 30, 10, 668, 731, 90, 0, 80),
(80, 85, 25, -10, 769, 820, 90, 79, 0),
(81, 85, 35, 30, 47, 124, 90, 0, 70),
(82, 75, 55, 20, 369, 420, 90, 0, 85),
(83, 72, 55, -20, 265, 338, 90, 87, 0),
(84, 70, 58, 20, 458, 523, 90, 0, 89),
(85, 68, 60, -20, 555, 612, 90, 82, 0),
(86, 66, 55, 10, 173, 238, 90, 0, 91),
(87, 65, 55, 20, 85, 144, 90, 0, 83),
(88, 65, 60, -10, 645, 708, 90, 90, 0),
(89, 63, 58, -20, 737, 802, 90, 84, 0),
(90, 60, 55, 10, 20, 84, 90, 0, 88),
(91, 60, 60, -10, 836, 889, 90, 86, 0),
(92, 67, 85, 20, 368, 441, 90, 0, 93),
(93, 65, 85, -20, 475, 518, 90, 92, 0),
(94, 65, 82, -10, 285, 336, 90, 96, 0),
(95, 62, 80, -20, 196, 239, 90, 98, 0),
(96, 60, 80, 10, 95, 156, 90, 0, 94),
(97, 60, 85, 30, 561, 622, 0, 0, 106),
(98, 58, 75, 20, 30, 84, 90, 0, 95),
(99, 55, 80, -20, 743, 820, 90, 100, 0),
(100, 55, 85, 20, 647, 726, 90, 0, 99),
(101, 25, 30, -10, 725, 786, 90, 51, 0),
(102, 48, 30, -10, 632, 693, 90, 64, 0),
(103, 28, 55, -10, 732, 777, 90, 23, 0),
(104, 88, 35, -20, 109, 170, 90, 78, 0),
(105, 5, 45, -10, 665, 716, 90, 36, 0),
(106, 60, 85, -30, 561, 622, 90, 97, 0);
INSERT 0 106
```

[The orders¶](#)

The original data needs to be converted to an appropriate table:

```
WITH deliveries AS (SELECT * FROM lc101_c WHERE dindex = 0)
SELECT
  row_number() over() AS id, p.demand,
  p.id as p_node_id, p.x AS p_x, p.y AS p_y, p.open AS p_open, p.close as p_close, p.service as p_service,
  d.id as d_node_id, d.x AS d_x, d.y AS d_y, d.open AS d_open, d.close as d_close, d.service as d_service
INTO c_lc101
FROM deliveries as d JOIN lc101_c as p ON (d.pindex = p.id);
SELECT 53
SELECT * FROM c_lc101 LIMIT 1;
id | demand | p_node_id | p_x | p_y | p_open | p_close | p_service | d_node_id | d_x | d_y | d_open | d_close | d_service
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 10 | 3 | 42 | 66 | 65 | 146 | 90 | 75 | 45 | 65 | 997 | 1068 | 90
(1 row)
```

[The query¶](#)

Showing only the relevant information to compare with the best solution information published on<https://www.sintef.no/projectweb/top/pdptw/100-customers/>

- The best solution found for **lc101** is a travel time: 828.94
- This implementation's travel time: 854.54

```
SELECT travel_time, 828.94 AS best
FROM pgr_pickDeliverEuclidean(
  $$SELECT * FROM c_lc101 $$,
  $$SELECT * FROM v_lc101 $$,
  max_cycles => 2, initial_sol => 4) WHERE vehicle_seq = -2;
travel_time | best
-----+-----
854.5412705652799 | 828.94
(1 row)
```

[See Also¶](#)

- [Vehicle Routing Functions - Category](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

**pgr\_vrpOneDepot - Experimental¶**

☐ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

**No documentation available**

Availability

- Version 2.1.0
  - New experimental function.
- TBD

[Description¶](#)

- TBD

[Signatures¶](#)

- TBD

[Parameters¶](#)

- TBD

[Inner Queries¶](#)

- TBD



Result columns1

- TBD

Additional Example:1

```
BEGIN;
BEGIN
SET client_min_messages TO NOTICE;
SET
SELECT * FROM pgr_vrpOneDepot(
'SELECT * FROM solomon_100_RC_101',
'SELECT * FROM vrp_vehicles',
'SELECT * FROM vrp_distance',
1);
oid | opos | vid | tarrival | tdepart
-----+-----+-----+-----+-----
-1 | 1 | 1 | 0 | 0
7 | 2 | 1 | 0 | 0
9 | 3 | 1 | 0 | 0
8 | 4 | 1 | 0 | 0
6 | 5 | 1 | 0 | 0
5 | 6 | 1 | 0 | 0
4 | 7 | 1 | 0 | 0
2 | 8 | 1 | 0 | 0
6 | 9 | 1 | 40 | 51
8 | 10 | 1 | 62 | 89
9 | 11 | 1 | 94 | 104
7 | 12 | 1 | 110 | 120
4 | 13 | 1 | 131 | 141
2 | 14 | 1 | 144 | 155
5 | 15 | 1 | 162 | 172
-1 | 16 | 1 | 208 | 208
-1 | 1 | 2 | 0 | 0
10 | 2 | 2 | 0 | 0
11 | 3 | 2 | 0 | 0
10 | 4 | 2 | 34 | 101
11 | 5 | 2 | 106 | 129
-1 | 6 | 2 | 161 | 161
-1 | 1 | 3 | 0 | 0
3 | 2 | 3 | 0 | 0
3 | 3 | 3 | 31 | 60
-1 | 4 | 3 | 91 | 91
-1 | 0 | 0 | -1 | 460
(27 rows)

ROLLBACK;
ROLLBACK
```

Data

```
DROP TABLE IF EXISTS solomon_100_RC_101 cascade;
CREATE TABLE solomon_100_RC_101 (
id integer NOT NULL PRIMARY KEY,
order_unit integer,
open_time integer,
close_time integer,
service_time integer,
x float8,
y float8
);

INSERT INTO solomon_100_RC_101 (id, x, y, order_unit, open_time, close_time, service_time) VALUES
(1, 40.000000, 50.000000, 0, 0, 240, 0),
(2, 25.000000, 85.000000, 20, 145, 175, 10),
(3, 22.000000, 75.000000, 30, 50, 80, 10),
(4, 22.000000, 85.000000, 10, 109, 139, 10),
(5, 20.000000, 80.000000, 40, 141, 171, 10),
(6, 20.000000, 85.000000, 20, 41, 71, 10),
(7, 18.000000, 75.000000, 20, 95, 125, 10),
(8, 15.000000, 75.000000, 20, 79, 109, 10),
(9, 15.000000, 80.000000, 10, 91, 121, 10),
(10, 10.000000, 35.000000, 20, 91, 121, 10),
(11, 10.000000, 40.000000, 30, 119, 149, 10);

DROP TABLE IF EXISTS vrp_vehicles cascade;
CREATE TABLE vrp_vehicles (
vehicle_id integer not null primary key,
capacity integer,
case_no integer
);

INSERT INTO vrp_vehicles (vehicle_id, capacity, case_no) VALUES
(1, 200, 5),
(2, 200, 5),
(3, 200, 5);

DROP TABLE IF EXISTS vrp_distance cascade;
WITH
the_matrix_info AS (
SELECT A.id AS src_id, B.id AS dest_id, sqrt((a.x - b.x) * (a.x - b.x) + (a.y - b.y) * (a.y - b.y)) AS cost
FROM solomon_100_rc_101 AS A, solomon_100_rc_101 AS B WHERE A.id != B.id
)
SELECT src_id, dest_id, cost, cost AS distance, cost AS travelltime
INTO vrp_distance
FROM the_matrix_info;
```

See Also1

- [https://en.wikipedia.org/wiki/Vehicle\\_routing\\_problem](https://en.wikipedia.org/wiki/Vehicle_routing_problem)

Indices and tables

- [Index](#)
- [Search Page](#)

Introduction1

Vehicle Routing Problems *VRP* are **NP-hard** optimization problem, it generalises the travelling salesman problem (TSP).

- The objective of the VRP is to minimize the total route cost.
- There are several variants of the VRP problem,

**pgRouting does not try to implement all variants.**

Characteristics1

- Capacitated Vehicle Routing Problem *CVRP* where The vehicles have limited carrying capacity of the goods.
- Vehicle Routing Problem with Time Windows *VRPTW* where the locations have time windows within which the vehicle's visits must be made.

- Vehicle Routing Problem with Pickup and Delivery *VRPPD* where a number of goods need to be moved from certain pickup locations to other delivery locations.

#### Limitations

- No multiple time windows for a location.
- Less vehicle used is considered better.
- Less total duration is better.
- Less wait time is better.

#### [Pick & Delivery¶](#)

Problem: *CVRPPDTW* Capacitated Pick and Delivery Vehicle Routing problem with Time Windows

- Times are relative to 0
- The vehicles
  - have start and ending service duration times.
  - have opening and closing times for the start and ending locations.
  - have a capacity.
- The orders
  - Have pick up and delivery locations.
  - Have opening and closing times for the pickup and delivery locations.
  - Have pickup and delivery duration service times.
  - have a demand request for moving goods from the pickup location to the delivery location.
- Time based calculations:
  - Travel time between customers is  $\backslash(\text{distance} / \text{speed})$
  - Pickup and delivery order pair is done by the same vehicle.
  - A pickup is done before the delivery.

#### [Parameters¶](#)

#### [Pick & deliver¶](#)

Used in [pgr\\_pickDeliverEuclidean - Experimental](#)

Column	Type	Description
<a href="#">Orders SQL</a>	TEXT	<a href="#">Orders SQL</a> as described below.

<a href="#">Vehicles SQL</a>	TEXT	<a href="#">Vehicles SQL</a> as described below.
------------------------------	------	--

Used in [pgr\\_pickDeliver - Experimental](#)

Column	Type	Description
<a href="#">Orders SQL</a>	TEXT	<a href="#">Orders SQL</a> as described below.

<a href="#">Vehicles SQL</a>	TEXT	<a href="#">Vehicles SQL</a> as described below.
------------------------------	------	--

<a href="#">Matrix SQL</a>	TEXT	<a href="#">Matrix SQL</a> as described below.
----------------------------	------	--

#### [Pick-Deliver optional parameters¶](#)

Column	Type	Default	Description
factor	NUMERIC	1	Travel time multiplier. See <a href="#">Factor handling</a>
max_cycles	INTEGER	10	Maximum number of cycles to perform on the optimization.
			Initial solution to be used.
			<ul style="list-style-type: none"> <li>• 1 One order per truck</li> <li>• 2 Push front order.</li> <li>• 3 Push back order.</li> </ul>
initial_sol	INTEGER	4	<ul style="list-style-type: none"> <li>• 4 Optimize insert.</li> <li>• 5 Push back order that allows more orders to be inserted at the back</li> <li>• 6 Push front order that allows more orders to be inserted at the front</li> </ul>

#### [Inner Queries¶](#)

#### [Orders SQL¶](#)

Common columns for the orders SQL in both implementations:

Column	Type	Description
id	ANY-INTEGER	Identifier of the pick-delivery order pair.
demand	ANY-NUMERICAL	Number of units in the order
p_open	ANY-NUMERICAL	The time, relative to 0, when the pickup location opens.
p_close	ANY-NUMERICAL	The time, relative to 0, when the pickup location closes.
[p_service]	ANY-NUMERICAL	The duration of the loading at the pickup location. <ul style="list-style-type: none"> <li>When missing: 0 time units are used</li> </ul>
d_open	ANY-NUMERICAL	The time, relative to 0, when the delivery location opens.
d_close	ANY-NUMERICAL	The time, relative to 0, when the delivery location closes.
[d_service]	ANY-NUMERICAL	The duration of the unloading at the delivery location. <ul style="list-style-type: none"> <li>When missing: 0 time units are used</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

For [pgr\\_pickDeliver - Experimental](#) the pickup and delivery identifiers of the locations are needed:

Column	Type	Description
p_node_id	ANY-INTEGER	The node identifier of the pickup, must match a vertex identifier in the <a href="#">Matrix SQL</a> .
d_node_id	ANY-INTEGER	The node identifier of the delivery, must match a vertex identifier in the <a href="#">Matrix SQL</a> .

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

For [pgr\\_pickDeliverEuclidean - Experimental](#) the  $\backslash((x, y)\backslash)$  values of the locations are needed:

Column	Type	Description
p_x	ANY-NUMERICAL	$\backslash(x\backslash)$ value of the pick up location
p_y	ANY-NUMERICAL	$\backslash(y\backslash)$ value of the pick up location
d_x	ANY-NUMERICAL	$\backslash(x\backslash)$ value of the delivery location
d_y	ANY-NUMERICAL	$\backslash(y\backslash)$ value of the delivery location

Where:

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

[Vehicles SQL ¶](#)

Common columns for the vehicles SQL in both implementations:

Column	Type	Description
id	ANY-NUMERICAL	Identifier of the vehicle.
capacity	ANY-NUMERICAL	Maiximum capacity units
start_open	ANY-NUMERICAL	The time, relative to 0, when the starting location opens.
start_close	ANY-NUMERICAL	The time, relative to 0, when the starting location closes.
[start_service]	ANY-NUMERICAL	The duration of the loading at the starting location. <ul style="list-style-type: none"> <li>When missing: A duration of <math>\backslash(0\backslash)</math> time units is used.</li> </ul>

Column	Type	Description
[end_open]	ANY-NUMERICAL	The time, relative to 0, when the ending location opens. <ul style="list-style-type: none"> <li>When missing: The value of start_open is used</li> </ul>
[end_close]	ANY-NUMERICAL	The time, relative to 0, when the ending location closes. <ul style="list-style-type: none"> <li>When missing: The value of start_close is used</li> </ul>
[end_service]	ANY-NUMERICAL	The duration of the loading at the ending location. <ul style="list-style-type: none"> <li>When missing: A duration in start_service is used.</li> </ul>

For [pgr\\_pickDeliver - Experimental](#) the starting and ending identifiers of the locations are needed:

Column	Type	Description
start_node_id	ANY-INTEGER	The node identifier of the start location, must match a vertex identifier in the <a href="#">Matrix SQL</a> .
[end_node_id]	ANY-INTEGER	The node identifier of the end location, must match a vertex identifier in the <a href="#">Matrix SQL</a> . <ul style="list-style-type: none"> <li>When missing: end_node_id is used.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

For [pgr\\_pickDeliverEuclidean - Experimental](#) the  $\backslash((x, y)\backslash)$  values of the locations are needed:

Column	Type	Description
start_x	ANY-NUMERICAL	$\backslash(x)\backslash$ value of the starting location
start_y	ANY-NUMERICAL	$\backslash(y)\backslash$ value of the starting location
[end_x]	ANY-NUMERICAL	$\backslash(x)\backslash$ value of the ending location <ul style="list-style-type: none"> <li>When missing: start_x is used.</li> </ul>
[end_y]	ANY-NUMERICAL	$\backslash(y)\backslash$ value of the ending location <ul style="list-style-type: none"> <li>When missing: start_y is used.</li> </ul>

Where:

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

[Matrix SQL¶](#)

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

[Result columns¶](#)

Returns set of  
(seq, vehicle\_seq, vehicle\_id, stop\_seq, stop\_type,  
travel\_time, arrival\_time, wait\_time, service\_time, departure\_time)  
UNION  
(summary row)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
vehicle_seq	INTEGER	Sequential value starting from 1 for current vehicles. The $\backslash(n_{th})\backslash$ vehicle in the solution. <ul style="list-style-type: none"> <li>Value <math>\backslash(-2)\backslash</math> indicates it is the summary row.</li> </ul>
vehicle_id	BIGINT	Current vehicle identifier. <ul style="list-style-type: none"> <li>Summary row has the <b>total capacity violations</b>. <ul style="list-style-type: none"> <li>A capacity violation happens when overloading or underloading a vehicle.</li> </ul> </li> </ul>

Column	Type	Description
		Sequential value starting from 1 for the stops made by the current vehicle. The $m_{th}$ stop of the current vehicle.
stop_seq	INTEGER	<ul style="list-style-type: none"> <li>Summary row has the <b>total time windows violations</b>. <ul style="list-style-type: none"> <li>A time window violation happens when arriving after the location has closed.</li> </ul> </li> <li>Kind of stop location the vehicle is at <ul style="list-style-type: none"> <li><math>-1</math>: at the solution summary row</li> </ul> </li> </ul>
stop_type	INTEGER	<ul style="list-style-type: none"> <li><math>1</math>: Starting location</li> <li><math>2</math>: Pickup location</li> <li><math>3</math>: Delivery location</li> <li><math>6</math>: Ending location and indicates the vehicle's summary row</li> </ul>
order_id	BIGINT	Pickup-Delivery order pair identifier. <ul style="list-style-type: none"> <li>Value <math>-1</math>: When no order is involved on the current stop location.</li> </ul>
cargo	FLOAT	Cargo units of the vehicle when leaving the stop. <ul style="list-style-type: none"> <li>Value <math>-1</math> on solution summary row.</li> </ul>
travel_time	FLOAT	Travel time from previous stop_seq to current stop_seq. <ul style="list-style-type: none"> <li>Summary has the <b>total traveling time</b>: <ul style="list-style-type: none"> <li>The sum of all the travel_time.</li> </ul> </li> </ul>
arrival_time	FLOAT	Time spent waiting for current location to open. <ul style="list-style-type: none"> <li><math>-1</math>: at the solution summary row.</li> <li><math>0</math>: at the starting location.</li> </ul>
wait_time	FLOAT	Time spent waiting for current location to open. <ul style="list-style-type: none"> <li>Summary row has the <b>total waiting time</b>: <ul style="list-style-type: none"> <li>The sum of all the wait_time.</li> </ul> </li> </ul>
service_time	FLOAT	Service duration at current location. <ul style="list-style-type: none"> <li>Summary row has the <b>total service time</b>: <ul style="list-style-type: none"> <li>The sum of all the service_time.</li> </ul> </li> </ul>
departure_time	FLOAT	<ul style="list-style-type: none"> <li>The time at which the vehicle departs from the stop. <ul style="list-style-type: none"> <li><math>(arrival\_time + wait\_time + service\_time)</math></li> </ul> </li> <li>The ending location has the <b>total time</b> used by the current vehicle.</li> <li>Summary row has the <b>total solution time</b>: <ul style="list-style-type: none"> <li><math>(total\ traveling\ time + total\ waiting\ time + total\ service\ time)</math></li> </ul> </li> </ul>

#### Summary Row

Column	Type	Description
seq	INTEGER	Continues the sequence
vehicle_seq	INTEGER	Value $-2$ indicates it is the summary row.
vehicle_id	BIGINT	<b>total capacity violations:</b> <ul style="list-style-type: none"> <li>A capacity violation happens when overloading or underloading a vehicle.</li> </ul>
stop_seq	INTEGER	<b>total time windows violations:</b> <ul style="list-style-type: none"> <li>A time window violation happens when arriving after the location has closed.</li> </ul>
stop_type	INTEGER	$-1$
order_id	BIGINT	$-1$
cargo	FLOAT	$-1$
travel_time	FLOAT	<b>total traveling time:</b> <ul style="list-style-type: none"> <li>The sum of all the travel_time.</li> </ul>
arrival_time	FLOAT	$-1$

Column	Type	Description
<b>total waiting time:</b>		
wait_time	FLOAT	<ul style="list-style-type: none"> <li>The sum of all the wait_time.</li> </ul>
<b>total service time:</b>		
service_time	FLOAT	<ul style="list-style-type: none"> <li>The sum of all the service_time.</li> </ul>
Summary row has the <b>total solution time:</b>		
departure_time	FLOAT	<ul style="list-style-type: none"> <li><math>\backslash(\text{total}\backslash \text{traveling}\backslash \text{time} + \text{total}\backslash \text{waiting}\backslash \text{time} + \text{total}\backslash \text{service}\backslash \text{time})</math></li> </ul>

#### Handling Parameters¶

To define a problem, several considerations have to be done, to get consistent results. This section gives an insight of how parameters are to be considered.

- [Capacity and Demand Units Handling](#)
- [Locations](#)
- [Time Handling](#)
- [Factor Handling](#)

#### Capacity and Demand Units Handling¶

The *capacity* of a vehicle, can be measured in:

- Volume units like  $\backslash(\text{m}^3\backslash)$ .
- Area units like  $\backslash(\text{m}^2\backslash)$  (when no stacking is allowed).
- Weight units like  $\backslash(\text{kg}\backslash)$ .
- Number of boxes that fit in the vehicle.
- Number of seats in the vehicle

The *demand* request of the pickup-deliver orders must use the same units as the units used in the vehicle's *capacity*.

To handle problems like: 10 (equal dimension) boxes of apples and 5 kg of feathers that are to be transported (not packed in boxes).

- If the vehicle's **capacity** is measured in *boxes*, a conversion of *kg of feathers to number of boxes* is needed.
- If the vehicle's **capacity** is measured in *kg*, a conversion of *box of apples to kg* is needed.

Showing how the 2 possible conversions can be done

Let: -  $\backslash(f\_boxes\backslash)$ : number of boxes needed for 1 kg of feathers. -  $\backslash(a\_weight\backslash)$ : weight of 1 box of apples.

**Capacity Units      apples      feathers**

boxes      10       $\backslash(5 * f\_boxes\backslash)$

kg       $\backslash(10 * a\_weight\backslash)$       5

#### Locations¶

- When using [pgr\\_pickDeliverEuclidean - Experimental](#):
  - The vehicles have  $\backslash((x, y)\backslash)$  pairs for start and ending locations.
  - The orders Have  $\backslash((x, y)\backslash)$  pairs for pickup and delivery locations.
- When using [pgr\\_pickDeliver - Experimental](#):
  - The vehicles have identifiers for the start and ending locations.
  - The orders have identifiers for the pickup and delivery locations.
  - All the identifiers are indices to the given matrix.

#### Time Handling¶

The times are relative to 0. All time units have to be converted to a 0 reference and the same time units.

Suppose that a vehicle's driver starts the shift at 9:00 am and ends the shift at 4:30 pm and the service time duration is 10 minutes with 30 seconds.

Meaning of 0	time units	9:00 am	4:30 pm	10 min 30 secs
0:00 am	hours	9	16.5	$\backslash(10.5 / 60 = 0.175\backslash)$
0:00 am	minutes	$\backslash(9*60 = 54\backslash)$	$\backslash(16.5*60 = 990\backslash)$	10.5
9:00 am	hours	0	7.5	$\backslash(10.5 / 60 = 0.175\backslash)$
9:00 am	minutes	0	$\backslash(7.5*60 = 540\backslash)$	10.5

#### Factor handling¶

factor acts as a multiplier to convert from distance values to time units the matrix values or the euclidean values.

- When the values are already in the desired time units
  - factor should be 1
  - When factor > 1 the travel times are faster
  - When factor < 1 the travel times are slower

For the [pgr\\_pickDeliverEuclidean - Experimental](#):

Working with time units in seconds, and x/y in lat/lon: Factor: would depend on the location of the points and on the average velocity say 25m/s is the velocity.

Latitude	Conversion	Factor
45	1 longitude degree is (78846.81m)/(25m/s)	3153 s
0	1 longitude degree is (111319.46 m)/(25m/s)	4452 s

For the [pgr\\_pickDeliver - Experimental](#):

Given  $v = d / t$  therefore  $t = d / v$  And the factor becomes  $(1 / v)$

Where:

v:  
Velocity

d:  
Distance

t:  
Time

For the following equivalences  $(10m/s \approx 600m/min \approx 36 km/hr)$

Working with time units in seconds and the matrix been in meters: For a 1000m length value on the matrix:

Units	velocity	Conversion	Factor	Result
seconds	$(10 m/s)$	$(\frac{1}{10m/s})$	$(0.1s/m)$	$(1000m * 0.1s/m = 100s)$
minutes	$(\frac{600}{m/min})$	$(\frac{1}{600m/min})$	$(0.0016min/m)$	$(1000m * 0.0016min/m = 1.6min)$
Hours	$(36 km/hr)$	$(\frac{1}{36 km/hr})$	$(0.0277hr/km)$	$(1km * 0.0277hr/km = 0.0277hr)$

[See Also](#)

- [https://en.wikipedia.org/wiki/Vehicle\\_routing\\_problem](https://en.wikipedia.org/wiki/Vehicle_routing_problem)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

**withPoints - Category**

When points are added to the graph.

Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.
- [withPoints - Family of functions](#) - Functions based on Dijkstra algorithm.
- From the [TRSP - Family of functions](#):
  - [pgr\\_trsp\\_withPoints - Proposed](#) - Vertex/Point routing with restrictions.
  - [pgr\\_trspVia\\_withPoints - Proposed](#) - Via Vertex/point routing with restrictions.

**Introduction**

The **with points** category modifies the graph on the fly by adding points on edges as required by the [Points SQL](#) query.

The functions within this category give the ability to process between arbitrary points located outside the original graph.

This category of functions was thought for routing vehicles, but might as well work for some other application not involving vehicles.

When given a point identifier `pid` that its being mapped to an edge with an identifier `edge_id`, with a fraction from the source to the target along the edge `fraction` and some additional information about which side of the edge the point is on `side`, then processing from arbitrary points can be done on fixed networks.

All this functions consider as many traits from the “real world” as possible:

- Kind of graph:
  - **directed** graph
  - **undirected** graph
- Arriving at the point:
  - Compulsory arrival on the side of the segment where the point is located.
  - On either side of the segment.
- Countries with:
  - **Right** side driving
  - **Left** side driving
- Some points are:
  - **Permanent**: for example the set of points of clients stored in a table in the data base.
    - The graph has been modified to permanently have those points as vertices.
    - There is a table on the database that describes the points
  - **Temporal**: for example points given through a web application
    - Use `pgr_findCloseEdges` in the [Points SQL](#).
- The numbering of the points are handled with negative sign.
  - This sign change is to avoid confusion when there is a vertex with the same identifier as the point identifier.
  - Original point identifiers are to be positive.
  - Transformation to negative is done internally.
  - Interpretation of the sign on the node information of the output
    - positive sign is a vertex of the original graph
    - negative sign is a point of the [Points SQL](#)

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Points SQL</a>	TEXT	<a href="#">Points SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path. Negative value is for point's identifier.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices. Negative values are for point's identifiers.
end vid	BIGINT	Identifier of the ending vertex of the path. Negative value is for point's identifier.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices. Negative values are for point's identifiers.

Optional parameters

Parameter	Type	Default	Description
Value in [r, l] indicating if the driving side is:			
driving_side	CHAR	r	<ul style="list-style-type: none"><li>• r for right driving side</li><li>• l for left driving side</li><li>• Any other value will be considered as r</li></ul>
details	BOOLEAN	false	<ul style="list-style-type: none"><li>• When true the results will include the points that are in the path.</li><li>• When false the results will not include the points that are in the path.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.



Column	Type	Default	Description
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Points [SQL¶](#)

Parameter	Type	Default	Description
			Identifier of the point. <ul style="list-style-type: none"> <li>Use with positive value, as internally will be converted to negative value</li> <li>If column is present, it can not be NULL.</li> <li>If column is not present, a sequential negative <b>value</b> will be given automatically.</li> </ul>
pid	ANY-INTEGER	value	
edge_id	ANY-INTEGER		Identifier of the "closest" edge to the point.
fraction	ANY-NUMERICAL		Value in <0,1> that indicates the relative position from the first end point of the edge.
			Value in [b, r, l, NULL] indicating if the point is: <ul style="list-style-type: none"> <li>In the right r,</li> <li>In the left l,</li> <li>In both sides b, NULL</li> </ul>
side	CHAR	b	

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations [SQL¶](#)

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Advanced documentation[¶](#)

Contents

- About points
- Driving side
  - Right driving side
  - Left driving side
  - Driving side does not matter
- Creating temporary vertices
  - On a right hand side driving network
  - On a left hand side driving network
  - When driving side does not matter

[About points¶](#)

For this section the following city (see [Sample Data](#)) some interesting points such as restaurant, supermarket, post office, etc. will be used as example.

[\\_images/Fig1-originalData.png](#)



- The graph is **directed**
- Red arrows show the (source, target) of the edge on the edge table
- Blue arrows show the (target, source) of the edge on the edge table
- Each point location shows where it is located with relation of the edge(source, target)
  - On the right for points **2** and **4**.
  - On the left for points **1**, **3** and **5**.
  - On both sides for point **6**.

The representation on the data base follows the [Points SQL](#) description, and for this example:

```
SELECT pid, edge_id, fraction, side FROM pointsOfInterest;
pid | edge_id | fraction | side
-----
1 | 1 | 0.4 | l
4 | 6 | 0.3 | r
3 | 12 | 0.6 | l
2 | 15 | 0.4 | r
5 | 5 | 0.8 | l
6 | 4 | 0.7 | b
(6 rows)
```

[Driving side¶](#)

In the following images:

- The squared vertices are the temporary vertices,
- The temporary vertices are added according to the driving side,
- visually showing the differences on how depending on the driving side the data is interpreted.

[Right driving side¶](#)

[\\_images/rightDrivingSide.png](#)



- Point **1** located on edge (6, 5)
- Point **2** located on edge (16, 17)
- Point **3** located on edge (8, 12)
- Point **4** located on edge (1, 3)
- Point **5** located on edge (10, 11)
- Point **6** located on edges (6, 7) and (7, 6)

[Left driving side¶](#)



- Point 1 located on edge (5, 6)
- Point 2 located on edge (17, 16)
- Point 3 located on edge (8, 12)
- Point 4 located on edge (3, 1)
- Point 5 located on edge (10, 11)
- Point 6 located on edges (6, 7) and (7, 6)

Driving side does not matter¶

- Like having all points to be considered in both sides
- Preferred usage on **undirected** graphs
- On the [TRSP - Family of functions](#) this option is not valid



- Point 1 located on edge (5, 6) and (6, 5)
- Point 2 located on edge (17, 16) and 16, 17
- Point 3 located on edge (8, 12)
- Point 4 located on edge (3, 1) and (1, 3)
- Point 5 located on edge (10, 11)
- Point 6 located on edges (6, 7) and (7, 6)

Creating temporary vertices¶

This section will demonstrate how a temporary vertex is created internally on the graph.

Problem

For edge:

```
SELECT id, source, target, cost, reverse_cost
FROM edges WHERE id = 15;
id | source | target | cost | reverse_cost
-----+-----+-----+-----+-----
15 | 16 | 17 | 1 | 1
(1 row)
```

insert point:

```
SELECT pid, edge_id, fraction, side
FROM pointsOfInterest WHERE pid = 2;
pid | edge_id | fraction | side
-----+-----+-----+-----
2 | 15 | 0.4 | r
(1 row)
```

On a right hand side driving network¶

Right driving side

[\\_images/rightDrivingSide.png](#)



- Arrival to point -2 can be achieved only via vertex **16**.
- Does not affects edge (17, 16), therefore the edge is kept.
- It only affects the edge (16, 17), therefore the edge is removed.
- Create two new edges:
  - Edge (16, -2) with cost 0.4 (original cost \* fraction ==  $\backslash(1 * 0.4\backslash)$ )
  - Edge (-2, 17) with cost 0.6 (the remaining cost)
- The total cost of the additional edges is equal to the original cost.
- If more points are on the same edge, the process is repeated recursively.

[On a left hand side driving network](#)

Left driving side

[\\_images/leftDrivingSide.png](#)



- Arrival to point -2 can be achieved only via vertex **17**.
- Does not affects edge (16, 17), therefore the edge is kept.
- It only affects the edge (17, 16), therefore the edge is removed.
- Create two new edges:
  - Work with the original edge (16, 17) as the fraction is a fraction of the original:
    - Edge (16, -2) with cost 0.4 (original cost \* fraction ==  $\backslash(1 * 0.4\backslash)$ )
    - Edge (-2, 17) with cost 0.6 (the remaining cost)
    - If more points are on the same edge, the process is repeated recursively.
  - Flip the Edges and add them to the graph:
    - Edge (17, -2) becomes (-2, 16) with cost 0.4 and is added to the graph.
    - Edge (-2, 16) becomes (17, -2) with cost 0.6 and is added to the graph.
- The total cost of the additional edges is equal to the original cost.

[When driving side does not matter](#)

[\\_images/noMatterDrivingSide.png](#)



- Arrival to point -2 can be achieved via vertices **16** or **17**.
- Affects the edges (16, 17) and (17, 16), therefore the edges are removed.
- Create four new edges:
  - Work with the original edge (16, 17) as the fraction is a fraction of the original:
    - Edge (16, -2) with cost 0.4 (original cost \* fraction ==  $(1 * 0.4)$ )
    - Edge (-2, 17) with cost 0.6 (the remaining cost)
    - If more points are on the same edge, the process is repeated recursively.
  - Flip the Edges and add all the edges to the graph:
    - Edge (16, -2) is added to the graph.
    - Edge (-2, 17) is added to the graph.
    - Edge (16, -2) becomes (-2, 16) with cost 0.4 and is added to the graph.
    - Edge (-2, 17) becomes (17, -2) with cost 0.6 and is added to the graph.

See Also

- [withPoints - Family of functions](#)

Indices and tables

- [Index](#)
- [Search Page](#)

See Also

Indices and tables

- [Index](#)
- [Search Page](#)

[All Pairs - Family of Functions](#)

- [pgr\\_floydWarshall](#) - Floyd-Warshall's algorithm.
- [pgr\\_johnson](#) - Johnson's algorithm

[A\\* - Family of functions](#)

- [pgr\\_aStar](#) - A\* algorithm for the shortest path.
- [pgr\\_aStarCost](#) - Get the aggregate cost of the shortest paths.
- [pgr\\_aStarCostMatrix](#) - Get the cost matrix of the shortest paths.

[Bidirectional A\\* - Family of functions](#)

- [pgr\\_bdAstar](#) - Bidirectional A\* algorithm for obtaining paths.
- [pgr\\_bdAstarCost](#) - Bidirectional A\* algorithm to calculate the cost of the paths.
- [pgr\\_bdAstarCostMatrix](#) - Bidirectional A\* algorithm to calculate a cost matrix of paths.

[Bidirectional Dijkstra - Family of functions](#)

- [pgr\\_bdDijkstra](#) - Bidirectional Dijkstra algorithm for the shortest paths.
- [pgr\\_bdDijkstraCost](#) - Bidirectional Dijkstra to calculate the cost of the shortest paths
- [pgr\\_bdDijkstraCostMatrix](#) - Bidirectional Dijkstra algorithm to create a matrix of costs of the shortest paths.

[Components - Family of functions](#)

- [pgr\\_connectedComponents](#) - Connected components of an undirected graph.
- [pgr\\_strongComponents](#) - Strongly connected components of a directed graph.
- [pgr\\_biconnectedComponents](#) - Biconnected components of an undirected graph.
- [pgr\\_articulationPoints](#) - Articulation points of an undirected graph.
- [pgr\\_bridges](#) - Bridges of an undirected graph.

[Contraction - Family of functions](#)

- [pgr\\_contraction](#)

[Dijkstra - Family of functions](#)

- [pgr\\_dijkstra](#) - Dijkstra's algorithm for the shortest paths.

- [pgr\\_dijkstraCost](#) - Get the aggregate cost of the shortest paths.
- [pgr\\_dijkstraCostMatrix](#) - Use pgr\_dijkstra to create a costs matrix.
- [pgr\\_drivingDistance](#) - Use pgr\_dijkstra to calculate catchment information.
- [pgr\\_KSP](#) - Use Yen algorithm with pgr\_dijkstra to get the K shortest paths.

#### Flow - Family of functions

- [pgr\\_maxFlow](#) - Only the Max flow calculation using Push and Relabel algorithm.
- [pgr\\_boykovKolmogorov](#) - Boykov and Kolmogorov with details of flow on edges.
- [pgr\\_edmondsKarp](#) - Edmonds and Karp algorithm with details of flow on edges.
- [pgr\\_pushRelabel](#) - Push and relabel algorithm with details of flow on edges.
- Applications
  - [pgr\\_edgeDisjointPaths](#) - Calculates edge disjoint paths between two groups of vertices.
  - [pgr\\_maxCardinalityMatch](#) - Calculates a maximum cardinality matching in a graph.

#### Kruskal - Family of functions

- [pgr\\_kruskal](#)
- [pgr\\_kruskalBFS](#)
- [pgr\\_kruskalDD](#)
- [pgr\\_kruskalDFS](#)

#### Metrics - Family of functions

- [pgr\\_degree](#) - Returns a set of vertices and corresponding count of incident edges to the vertex.

#### Prim - Family of functions

- [pgr\\_prim](#)
- [pgr\\_primBFS](#)
- [pgr\\_primDD](#)
- [pgr\\_primDFS](#)

#### Reference

- [pgr\\_version](#)
- [pgr\\_full\\_version](#)

#### Topology - Family of Functions

The following functions modify the database directly therefore the user must have special permissions given by the administrators to use them.

- [pgr\\_createTopology](#) - [Deprecated since v3.8.0](#) - create a topology based on the geometry.
- [pgr\\_createVerticesTable](#) - [Deprecated since 3.8.0](#) - reconstruct the vertices table based on the source and target information.
- [pgr\\_analyzeGraph](#) - [Deprecated since 3.8.0](#) - to analyze the edges and vertices of the edge table.
- [pgr\\_analyzeOneWay](#) - [Deprecated since 3.8.0](#) - to analyze directionality of the edges.
- [pgr\\_nodeNetwork](#) - [Deprecated since 3.8.0](#) - to create nodes to a not noded edge table.

#### Traveling Sales Person - Family of functions

- [pgr\\_TSP](#) - When input is given as matrix cell information.
- [pgr\\_TSPeuclidean](#) - When input are coordinates.

#### [pgr\\_trsp](#) - [Proposed](#) - Turn Restriction Shortest Path (TRSP)

#### Utilities

- [pgr\\_extractVertices](#) - Extracts vertex information based on the edge table information.
- [pgr\\_findCloseEdges](#) - Finds close edges of points on the fly
- [pgr\\_separateCrossing](#) - Breaks geometries that cross each other.
- [pgr\\_separateTouching](#) - Breaks geometries that (almost) touch each other.

## Functions by categories¶

#### [Cost - Category](#)

- [pgr\\_aStarCost](#)
- [pgr\\_bdAStarCost](#)
- [pgr\\_dijkstraCost](#)
- [pgr\\_bdDijkstraCost](#)
- [pgr\\_dijkstraNearCost](#) - [Proposed](#)

#### [Cost Matrix - Category](#)

- [pgr\\_aStarCostMatrix](#)
- [pgr\\_dijkstraCostMatrix](#)
- [pgr\\_bdAStarCostMatrix](#)
- [pgr\\_bdDijkstraCostMatrix](#)

#### [Driving Distance - Category](#)

- [pgr\\_drivingDistance](#) - Driving Distance based on Dijkstra's algorithm
- [pgr\\_primDD](#) - Driving Distance based on Prim's algorithm
- [pgr\\_kruskalDD](#) - Driving Distance based on Kruskal's algorithm

- Post processing
  - [pgr\\_alphaShape](#) - Alpha shape computation

#### [K shortest paths - Category](#)

- [pgr\\_KSP](#) - Yen's algorithm based on [pgr\\_dijkstra](#)

#### [Spanning Tree - Category](#)

- [Kruskal - Family of functions](#)
- [Prim - Family of functions](#)

#### [BFS - Category](#)

- [pgr\\_kruskalBFS](#)
- [pgr\\_primBFS](#)

#### [DFS - Category](#)

- [pgr\\_kruskalDFS](#)
- [pgr\\_primDFS](#)

## Available Functions but not official pgRouting functions¶

- [Proposed Functions](#)
- [Experimental Functions](#)

## Proposed Functions¶

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Families

#### [Dijkstra - Family of functions](#)

- [pgr\\_dijkstraVia - Proposed](#) - Get a route of a sequence of vertices.
- [pgr\\_dijkstraNear - Proposed](#) - Get the route to the nearest vertex.
- [pgr\\_dijkstraNearCost - Proposed](#) - Get the cost to the nearest vertex.

#### [withPoints - Family of functions](#)

- [pgr\\_withPoints - Proposed](#) - Route from/to points anywhere on the graph.
- [pgr\\_withPointsCost - Proposed](#) - Costs of the shortest paths.
- [pgr\\_withPointsCostMatrix - proposed](#) - Costs of the shortest paths.
- [pgr\\_withPointsKSP - Proposed](#) - K shortest paths.
- [pgr\\_withPointsDD - Proposed](#) - Driving distance.
- [pgr\\_withPointsVia - Proposed](#) - Via routing

#### [TRSP - Family of functions](#)

- [pgr\\_trsp - Proposed](#) - Vertex - Vertex routing with restrictions.
- [pgr\\_trspVia - Proposed](#) - Via Vertices routing with restrictions.
- [pgr\\_trsp\\_withPoints - Proposed](#) - Vertex/Point routing with restrictions.
- [pgr\\_trspVia\\_withPoints - Proposed](#) - Via Vertex/point routing with restrictions.

## TRSP - Family of functions¶

When points are also given as input:

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.

- Documentation might need refinement.
- [pgr\\_trsp - Proposed](#) - Vertex - Vertex routing with restrictions.
- [pgr\\_trspVia - Proposed](#) - Via Vertices routing with restrictions.
- [pgr\\_trsp\\_withPoints - Proposed](#) - Vertex/Point routing with restrictions.
- [pgr\\_trspVia\\_withPoints - Proposed](#) - Via Vertex/point routing with restrictions.

Warning

Read the [Migration guide](#) about how to migrate from the deprecated TRSP functionality to the new signatures or replacement functions.

☐ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting
- [pgr\\_turnRestrictedPath - Experimental](#) - Routing with restrictions.

**pgr\_trsp - Proposed**

pgr\_trsp - routing vertices with restrictions.

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Availability

- Version 3.4.0
  - New proposed signatures:
    - pgr\_trsp(One to One)
    - pgr\_trsp(One to Many)
    - pgr\_trsp(Many to One)
    - pgr\_trsp(Many to Many)
    - pgr\_trsp(Combinations)
  - Deprecated signatures
    - pgr\_trsp(text,integer,integer,boolean,boolean,text)
    - pgr\_trsp(text,integer,float,integer,float,boolean,boolean,text)
    - pgr\_trspViaVertices(text,array,boolean,boolean,text)
    - pgr\_trspviaedges(text,integer[],double precision[],boolean,boolean,text)
- Version 2.1.0
  - New prototypes
    - pgr\_trspViaVertices
    - pgr\_trspViaEdges



- Version 2.0.0
  - Official function.

Description

Turn restricted shortest path (TRSP) is an algorithm that receives turn restrictions in form of a query like those found in real world navigable road networks.

The main characteristics are:

- It does no guarantee the shortest path as it might contain restriction paths.

The general algorithm is as follows:

- Execute a Dijkstra.
- If the solution passes thru a restriction then.
  - Execute the **TRSP** algorithm with restrictions.

Boost Graph Inside

Signatures

Summary

`pgr_trsp(Edges SQL, Restrictions SQL, start vid, end vid, [directed])`  
`pgr_trsp(Edges SQL, Restrictions SQL, start vid, end vids, [directed])`  
`pgr_trsp(Edges SQL, Restrictions SQL, start vids, end vid, [directed])`  
`pgr_trsp(Edges SQL, Restrictions SQL, start vids, end vids, [directed])`  
`pgr_trsp(Edges SQL, Restrictions SQL, Combinations SQL, [directed])`  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

One to One

`pgr_trsp(Edges SQL, Restrictions SQL, start vid, end vid, [directed])`

Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertex \10\ on an undirected graph.

```
SELECT * FROM pgr_trsp(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  6, 10,
  false);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 2 | 1 | 0
2 | 2 | 6 | 10 | 10 | -1 | 0 | 1
(2 rows)
```

One to Many

`pgr_trsp(Edges SQL, Restrictions SQL, start vid, end vids, [directed])`

Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertices \10, 1\ on an undirected graph.

```
SELECT * FROM pgr_trsp(
  $$SELECT id, source, target, cost FROM edges$$,
  $$SELECT * FROM restrictions$$,
  6, ARRAY[10, 1],
  false);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 10 | 1 | 1
3 | 3 | 6 | 10 | 8 | 12 | 1 | 2
4 | 4 | 6 | 10 | 12 | 11 | 1 | 3
5 | 5 | 6 | 10 | 11 | 8 | 1 | 4
6 | 6 | 6 | 10 | 7 | 7 | 1 | 5
7 | 7 | 6 | 10 | 3 | 6 | 1 | 6
8 | 8 | 6 | 10 | 1 | -1 | 0 | 7
9 | 1 | 6 | 10 | 6 | 4 | 1 | 0
10 | 2 | 6 | 10 | 7 | 8 | 1 | 1
11 | 3 | 6 | 10 | 11 | 5 | 1 | 2
12 | 4 | 6 | 10 | 10 | -1 | 0 | 3
(12 rows)
```

Many to One

`pgr_trsp(Edges SQL, Restrictions SQL, start vids, end vid, [directed])`

Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \6, 1\ to vertex \8\ on a directed graph.

```
SELECT * FROM pgr_trsp(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  ARRAY[6, 1], 8);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 8 | 1 | 6 | 1 | 0
2 | 2 | 1 | 8 | 3 | 7 | 1 | 1
3 | 3 | 1 | 8 | 7 | 10 | 101 | 2
4 | 4 | 1 | 8 | 8 | -1 | 0 | 103
5 | 1 | 6 | 8 | 6 | 4 | 1 | 0
6 | 2 | 6 | 8 | 7 | 10 | 1 | 1
7 | 3 | 6 | 8 | 8 | -1 | 0 | 2
(7 rows)
```

Many to Many

pgr\_trsp([Edges SQL](#), [Restrictions SQL](#), start vids, end vids, [directed])

Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \{6, 1\} to vertices \{10, 8\} on an undirected graph.

```
SELECT * FROM pgr_trsp(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  ARRAY[6, 1], ARRAY[10, 8],
  false);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 8 | 6 | 1 | 0 |
2 | 2 | 1 | 8 | 3 | 7 | 1 | 1
3 | 3 | 1 | 8 | 7 | 4 | 1 | 2
4 | 4 | 1 | 8 | 6 | 2 | 1 | 3
5 | 5 | 1 | 8 | 10 | 5 | 1 | 4
6 | 6 | 1 | 8 | 11 | 11 | 1 | 5
7 | 7 | 1 | 8 | 12 | 12 | 1 | 6
8 | 8 | 1 | 8 | 8 | -1 | 0 | 7
9 | 1 | 1 | 10 | 1 | 6 | 1 | 0
10 | 2 | 1 | 10 | 3 | 7 | 1 | 1
11 | 3 | 1 | 10 | 7 | 4 | 1 | 2
12 | 4 | 1 | 10 | 6 | 2 | 1 | 3
13 | 5 | 1 | 10 | 10 | -1 | 0 | 4
14 | 1 | 6 | 8 | 6 | 4 | 1 | 0
15 | 2 | 6 | 8 | 7 | 10 | 1 | 1
16 | 3 | 6 | 8 | 8 | -1 | 0 | 2
17 | 1 | 6 | 10 | 6 | 2 | 1 | 0
18 | 2 | 6 | 10 | 10 | -1 | 0 | 1
(18 rows)
```

Combinations

pgr\_trsp([Edges SQL](#), [Restrictions SQL](#), [Combinations SQL](#), [directed])

Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table on an undirected graph.

```
SELECT * FROM pgr_trsp(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  $$SELECT * FROM (VALUES (6, 10), (6, 1), (6, 8), (1, 8)) AS combinations (source, target)$$);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 8 | 6 | 1 | 0 |
2 | 2 | 1 | 8 | 3 | 7 | 1 | 1
3 | 3 | 1 | 8 | 7 | 10 | 101 | 2
4 | 4 | 1 | 8 | 8 | -1 | 0 | 103
5 | 1 | 6 | 1 | 6 | 4 | 1 | 0
6 | 2 | 6 | 1 | 7 | 10 | 1 | 1
7 | 3 | 6 | 1 | 8 | 12 | 1 | 2
8 | 4 | 6 | 1 | 12 | 13 | 1 | 3
9 | 5 | 6 | 1 | 17 | 15 | 1 | 4
10 | 6 | 6 | 1 | 16 | 9 | 1 | 5
11 | 7 | 6 | 1 | 11 | 8 | 1 | 6
12 | 8 | 6 | 1 | 7 | 7 | 1 | 7
13 | 9 | 6 | 1 | 3 | 6 | 1 | 8
14 | 10 | 6 | 1 | 1 | -1 | 0 | 9
15 | 1 | 6 | 8 | 6 | 4 | 1 | 0
16 | 2 | 6 | 8 | 7 | 10 | 1 | 1
17 | 3 | 6 | 8 | 8 | -1 | 0 | 2
18 | 1 | 6 | 10 | 6 | 4 | 1 | 0
19 | 2 | 6 | 10 | 7 | 10 | 1 | 1
20 | 3 | 6 | 10 | 8 | 12 | 1 | 2
21 | 4 | 6 | 10 | 12 | 13 | 1 | 3
22 | 5 | 6 | 10 | 17 | 15 | 1 | 4
23 | 6 | 6 | 10 | 16 | 16 | 1 | 5
24 | 7 | 6 | 10 | 15 | 3 | 1 | 6
25 | 8 | 6 | 10 | 10 | -1 | 0 | 7
(25 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	SQL query as described.
<a href="#">Restrictions SQL</a>	TEXT	SQL query as described.
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	ANY-INTEGER	Identifier of the departure vertex.
start vids	ARRAY [ANY-INTEGER]	Array of identifiers of destination vertices.
end vid	ANY-INTEGER	Identifier of the departure vertex.
end vids	ARRAY [ANY-INTEGER]	Array of identifiers of destination vertices.

Where:  
ANY-INTEGER:  
SMALLINT, INTEGER, BIGINT

Optional parameters

Column	Type	Default	Description
--------	------	---------	-------------

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true the graph is considered <i>Directed</i></li> <li>When false the graph is considered as <i>Undirected</i>.</li> </ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Restrictions SQL

Column	Type	Description
path	ARRAY [ANY-INTEGER]	Sequence of edge identifiers that form a path that is not allowed to be taken. - Empty arrays onNULL arrays are ignored. - Arrays that have a NULL element will raise an exception.
Cost	ANY-NUMERICAL	Cost of taking the forbidden path.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Path identifier. <ul style="list-style-type: none"> <li>Has value 1 for the first of a path fromstart_vid to end_vid.</li> </ul>
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex.

Column	Type	Description
end_vid	BIGINT	Identifier of the ending vertex.
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

See Also

- [TRSP - Family of functions](#)
- [Deprecated documentation](#)
- [Migration guide](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_trspVia - Proposed

pgr\_trspVia Route that goes through a list of vertices with restrictions.

Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Availability

- Version 3.4.0
  - New proposed function.

Description

Given a list of vertices and a graph, this function is equivalent to finding the shortest path between(vertex\_i) and (vertex\_{i+1}) for all (i < size\_of(via;vertices)) trying not to use restricted paths.

The paths represents the sections of the route.

The general algorithm is as follows:

- Execute a [pgr\\_dijkstraVia - Proposed](#).
- For the set of sub paths of the solution that pass through a restriction then
  - Execute the **TRSP** algorithm with restrictions for the paths.
  - NOTE** when this is done, U\_turn\_on\_edge flag is ignored.

Boost Graph Inside

Signatures

One Via

pgr\_trspVia(Edges SQL, Restrictions SQL, via vertices, [options])

options: [directed, strict, U\_turn\_on\_edge]

Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost, route\_agg\_cost)  
OR EMPTY SET

Example:

Find the route that visits the vertices\({5, 1, 8}\) in that order on an directed graph.

```
SELECT * FROM pgr_trspVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  ARRAY[5, 1, 8]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
```

1	1	1	5	1	5	1	1	0	0
2	1	2	5	1	6	4	1	1	1
3	1	3	5	1	7	10	1	2	2
4	1	4	5	1	8	12	1	3	3
5	1	5	5	1	12	13	1	4	4
6	1	6	5	1	17	15	1	5	5
7	1	7	5	1	16	9	1	6	6
8	1	8	5	1	11	8	1	7	7
9	1	9	5	1	7	7	1	8	8
10	1	10	5	1	3	6	1	9	9

11		1		11		5		1		1		-1		0		10		10
12		2		1		1		8		1		6		1		0		10
13		2		2		1		8		3		7		1		1		11
14		2		3		1		8		7		10		101		2		12
15		2		4		1		8		8		-2		0		103		113
(15 rows)																		

Parameters1

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> query as described.
<a href="#">Restrictions SQL</a>	TEXT	<a href="#">Restrictions SQL</a> query as described.
via vertices	ARRAY[ <b>ANY-INTEGER</b> ]	Array of ordered vertices identifiers that are going to be visited.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Optional parameters1

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Via optional parameters1

Parameter	Type	Default	Description
strict	BOOLEAN	false	<ul style="list-style-type: none"><li>When true if a path is missing stops and returns <b>EMPTY SET</b></li><li>When false ignores missing paths returning all paths found</li></ul>
U_turn_on_edge	BOOLEAN	true	<ul style="list-style-type: none"><li>When true departing from a visited vertex will not try to avoid</li></ul>

Inner Queries1

Edges SQL1

Column	Type	Default	Description
id	<b>ANY-INTEGER</b>		Identifier of the edge.
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
target	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
cost	<b>ANY-NUMERICAL</b>		Weight of the edge (source, target)
reverse_cost	<b>ANY-NUMERICAL</b>	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Restrictions SQL1

Column	Type	Description
path	ARRAY [ <b>ANY-INTEGER</b> ]	Sequence of edge identifiers that form a path that is not allowed to be taken. - Empty arrays on NULL arrays are ignored. - Arrays that have a NULL element will raise an exception.
Cost	<b>ANY-NUMERICAL</b>	Cost of taking the forbidden path.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns¶

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Identifier of a path. Has value1 for the first path.
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex of the path.
end_vid	BIGINT	Identifier of the ending vertex of the path.
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence.
		• -1 for the last node of the path.
		• -2 for the last node of the route.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.
route_agg_cost	FLOAT	Total cost from start_vid of seq = 1 to end_vid of the currentseq.

Additional Examples¶

- [The main query](#)
  - [Aggregate cost of the third path.](#)
  - [Route's aggregate cost of the route at the end of the third path.](#)
  - [Nodes visited in the route.](#)
  - [The aggregate costs of the route when the visited vertices are reached.](#)
  - [Status of "passes in front" or "visits" of the nodes.](#)
- [Simulation of how algorithm works.](#)

All this examples are about the route that visits the vertices(\{5, 7, 1, 8, 15\})in that order on a directed graph.

[The main query¶](#)

```
SELECT * FROM pgr_trspVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  ARRAY[5, 7, 1, 8, 15]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 5 | 7 | 5 | 1 | 1 | 0 | 0
2 | 1 | 2 | 5 | 7 | 6 | 4 | 1 | 1 | 1
3 | 1 | 3 | 5 | 7 | 7 | -1 | 0 | 2 | 2
4 | 2 | 1 | 7 | 1 | 7 | 7 | 1 | 0 | 2
5 | 2 | 2 | 7 | 1 | 3 | 6 | 1 | 1 | 3
6 | 2 | 3 | 7 | 1 | 1 | -1 | 0 | 2 | 4
7 | 3 | 1 | 1 | 8 | 1 | 6 | 1 | 0 | 4
8 | 3 | 2 | 1 | 8 | 3 | 7 | 1 | 1 | 5
9 | 3 | 3 | 1 | 8 | 7 | 10 | 101 | 2 | 6
10 | 3 | 4 | 1 | 8 | 8 | -1 | 0 | 103 | 107
11 | 4 | 1 | 8 | 15 | 8 | 12 | 1 | 0 | 107
12 | 4 | 2 | 8 | 15 | 12 | 13 | 1 | 1 | 108
13 | 4 | 3 | 8 | 15 | 17 | 15 | 1 | 2 | 109
14 | 4 | 4 | 8 | 15 | 16 | 16 | 1 | 3 | 110
15 | 4 | 5 | 8 | 15 | 15 | -2 | 0 | 4 | 111
(15 rows)
```

[Aggregate cost of the third path¶](#)

```
SELECT agg_cost FROM pgr_trspVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  ARRAY[5, 7, 1, 8, 15])
WHERE path_id = 3 AND edge < 0;
agg_cost
-----
103
(1 row)
```

[Route's aggregate cost of the route at the end of the third path¶](#)

```
SELECT route_agg_cost FROM pgr_trspVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  ARRAY[5, 7, 1, 8, 15])
WHERE path_id = 3 AND edge < 0;
route_agg_cost
-----
107
(1 row)
```

[Nodes visited in the route.¶](#)

```
SELECT row_number() over () as node_seq, node
FROM pgr_trspVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
```

```

$$SELECT path, cost FROM restrictions$$,
ARRAY[5, 7, 1, 8, 15])
WHERE edge <> -1 ORDER BY seq;
node_seq | node
-----+-----
1 | 5
2 | 6
3 | 7
4 | 3
5 | 1
6 | 3
7 | 7
8 | 8
9 | 12
10 | 17
11 | 16
12 | 15
(12 rows)

```

The aggregate costs of the route when the visited vertices are reached.

```

SELECT path_id, route_agg_cost FROM pgr_trspVia(
$$SELECT id, source, target, cost, reverse_cost FROM edges$$,
$$SELECT path, cost FROM restrictions$$,
ARRAY[5, 7, 1, 8, 15])
WHERE edge < 0;
path_id | route_agg_cost
-----+-----
1 | 2
2 | 4
3 | 107
4 | 111
(4 rows)

```

Status of "passes in front" or "visits" of the nodes.

```

SELECT seq, route_agg_cost, node, agg_cost,
CASE WHEN edge = -1 THEN $$visits$$
ELSE $$passes in front$$
END as status
FROM pgr_trspVia(
$$SELECT id, source, target, cost, reverse_cost FROM edges$$,
$$SELECT path, cost FROM restrictions$$,
ARRAY[5, 7, 1, 8, 15])
WHERE agg_cost <= 0 or seq = 1;
seq | route_agg_cost | node | agg_cost | status
-----+-----+-----+-----+-----
1 | 0 | 5 | 0 | passes in front
2 | 1 | 6 | 1 | passes in front
3 | 2 | 7 | 2 | visits
5 | 3 | 3 | 1 | passes in front
6 | 4 | 1 | 2 | visits
8 | 5 | 3 | 1 | passes in front
9 | 6 | 7 | 2 | passes in front
10 | 107 | 8 | 103 | visits
12 | 108 | 12 | 1 | passes in front
13 | 109 | 17 | 2 | passes in front
14 | 110 | 16 | 3 | passes in front
15 | 111 | 15 | 4 | passes in front
(12 rows)

```

Simulation of how algorithm works.

The algorithm performs a [pgr\\_dijkstraVia - Proposed](#)

```

SELECT * FROM pgr_dijkstraVia(
$$SELECT id, source, target, cost, reverse_cost FROM edges$$,
ARRAY[6, 3, 6]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 6 | 3 | 6 | 4 | 1 | 0 | 0
2 | 1 | 2 | 6 | 3 | 7 | 7 | 1 | 1 | 1
3 | 1 | 3 | 6 | 3 | 3 | -1 | 0 | 2 | 2
4 | 2 | 1 | 3 | 6 | 3 | 7 | 1 | 0 | 2
5 | 2 | 2 | 3 | 6 | 7 | 4 | 1 | 1 | 3
6 | 2 | 3 | 3 | 6 | 6 | -2 | 0 | 2 | 4
(6 rows)

```

Detects which of the sub paths pass through a restriction in this case is for the path\_id = 5 from 6 to 3 because the path (15 → 1) is restricted.

Executes the [pgr\\_trsp - Proposed](#) algorithm for the conflicting paths.

```

SELECT 1 AS path_id, path_seq, start_vid, end_vid, node, edge, cost, agg_cost FROM pgr_trsp(
$$SELECT id, source, target, cost, reverse_cost FROM edges$$,
$$SELECT path, cost FROM restrictions$$,
6, 3);
path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 3 | 6 | 4 | 1 | 0
1 | 2 | 6 | 3 | 7 | 10 | 1 | 1
1 | 3 | 6 | 3 | 8 | 12 | 1 | 2
1 | 4 | 6 | 3 | 12 | 13 | 1 | 3
1 | 5 | 6 | 3 | 17 | 15 | 1 | 4
1 | 6 | 6 | 3 | 16 | 9 | 1 | 5
1 | 7 | 6 | 3 | 11 | 8 | 1 | 6
1 | 8 | 6 | 3 | 7 | 7 | 1 | 7
1 | 9 | 6 | 3 | 3 | -1 | 0 | 8
(9 rows)

```

From the [pgr\\_dijkstraVia - Proposed](#) result it removes the conflicting paths and builds the solution with the results of the [pgr\\_trsp - Proposed](#) algorithm:

```

WITH
solutions AS (
SELECT path_id, path_seq, start_vid, end_vid, node, edge, cost, agg_cost FROM pgr_dijkstraVia(
$$SELECT id, source, target, cost, reverse_cost FROM edges$$,
ARRAY[6, 3, 6]) WHERE path_id != 1
UNION
SELECT 1 AS path_id, path_seq, start_vid, end_vid, node, edge, cost, agg_cost FROM pgr_trsp(
$$SELECT id, source, target, cost, reverse_cost FROM edges$$,
$$SELECT path, cost FROM restrictions$$,
6, 3)),
with_seq AS (
SELECT row_number() over(ORDER BY path_id, path_seq) AS seq, *
FROM solutions),
aggregation AS (SELECT seq, SUM(cost) OVER(ORDER BY seq) AS route_agg_cost FROM with_seq)
SELECT with_seq.*, COALESCE(route_agg_cost, 0) AS route_agg_cost
FROM with_seq LEFT JOIN aggregation ON (with_seq.seq = aggregation.seq + 1);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 6 | 3 | 6 | 4 | 1 | 0 | 0
2 | 1 | 2 | 6 | 3 | 7 | 10 | 1 | 1 | 1
3 | 1 | 3 | 6 | 3 | 8 | 12 | 1 | 2 | 2

```

4	1	4	6	3	12	13	1	3	3
5	1	5	6	3	17	15	1	4	4
6	1	6	6	3	16	9	1	5	5
7	1	7	6	3	11	8	1	6	6
8	1	8	6	3	7	7	1	7	7
9	1	9	6	3	3	-1	0	8	8
10	2	1	3	6	3	7	1	0	8
11	2	2	3	6	7	4	1	1	9
12	2	3	3	6	6	-2	0	2	10

(12 rows)

Getting the same result as pgr\_trspVia:

```
SELECT * FROM pgr_trspVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  ARRAY[6, 3, 6]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | 6 | 3 | 6 | 4 | 1 | 0 | 0
2 | 1 | 2 | 2 | 6 | 3 | 7 | 10 | 1 | 1 | 1
3 | 1 | 3 | 3 | 6 | 3 | 8 | 12 | 1 | 2 | 2
4 | 1 | 4 | 4 | 6 | 3 | 12 | 13 | 1 | 3 | 3
5 | 1 | 5 | 5 | 6 | 3 | 17 | 15 | 1 | 4 | 4
6 | 1 | 6 | 6 | 6 | 3 | 16 | 9 | 1 | 5 | 5
7 | 1 | 7 | 6 | 3 | 11 | 8 | 1 | 6 | 6
8 | 1 | 8 | 6 | 3 | 7 | 7 | 1 | 7 | 7
9 | 1 | 9 | 6 | 3 | 3 | -1 | 0 | 8 | 8
10 | 2 | 1 | 3 | 6 | 3 | 7 | 1 | 0 | 8
11 | 2 | 2 | 3 | 6 | 7 | 4 | 1 | 1 | 9
12 | 2 | 3 | 3 | 6 | 6 | -2 | 0 | 2 | 10
(12 rows)
```

Example 8:

Sometimes U\_turn\_on\_edge flag is ignored when is set to false.  
The first step, doing a [pgr\\_dijkstraVia - Proposed](#) does consider not making a U turn on the same edge. But the path(16 \rightarrow 13\ (Rows 4 and 5) is restricted and the result is using it.

```
SELECT * FROM pgr_dijkstraVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  ARRAY[6, 7, 6], U_turn_on_edge => false);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | 6 | 7 | 6 | 4 | 1 | 0 | 0
2 | 1 | 2 | 6 | 7 | 7 | -1 | 0 | 1 | 1
3 | 2 | 1 | 7 | 6 | 7 | 8 | 1 | 0 | 1
4 | 2 | 2 | 7 | 6 | 11 | 9 | 1 | 1 | 2
5 | 2 | 3 | 7 | 6 | 16 | 16 | 1 | 2 | 3
6 | 2 | 4 | 7 | 6 | 15 | 3 | 1 | 3 | 4
7 | 2 | 5 | 7 | 6 | 10 | 2 | 1 | 4 | 5
8 | 2 | 6 | 7 | 6 | 6 | -2 | 0 | 5 | 6
(8 rows)
```

When executing the [pgr\\_trsp - Proposed](#) algorithm for the conflicting path, there is no U\_turn\_on\_edge flag.

```
SELECT 1 AS path_id, path_seq, start_vid, end_vid, node, edge, cost, agg_cost FROM pgr_trsp(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  7, 6);
path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 7 | 6 | 7 | 4 | 1 | 0
1 | 2 | 2 | 7 | 6 | 6 | -1 | 0 | 1
(2 rows)
```

Therefore the result ignores the U\_turn\_on\_edge flag when set to false.

```
SELECT * FROM pgr_trspVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  ARRAY[6, 7, 6], U_turn_on_edge => false);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | 6 | 7 | 6 | 4 | 1 | 0 | 0
2 | 1 | 2 | 6 | 7 | 7 | -1 | 0 | 1 | 1
3 | 2 | 1 | 7 | 6 | 7 | 4 | 1 | 0 | 1
4 | 2 | 2 | 7 | 6 | 6 | -2 | 0 | 1 | 2
(4 rows)
```

See Also

- [Via - Category](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

**pgr\_trsp\_withPoints - Proposed**

pgr\_trsp\_withPoints Routing Vertex/Point with restrictions.

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.



Availability

- Version 3.4.0
  - New proposed function.

Description

Modify the graph to include points defined by points\_sql. Using Dijkstra algorithm, find the shortest path

Characteristics:

- Vertices of the graph are:
  - **positive** when it belongs to the [Edges SQL](#)
  - **negative** when it belongs to the [Points SQL](#)
- Driving side can not be b
- Values are returned when there is a path.
  - When the starting vertex and ending vertex are the same, there is no path.
    - The agg\_cost the non included values (v, v) is 0
  - When the starting vertex and ending vertex are the different and there is no path:
    - The agg\_cost the non included values (u, v) is ∞
- For optimization purposes, any duplicated value in the start\_vids or end\_vids are ignored.
- The returned values are ordered: - start\_vid ascending - end\_vid ascending
- Running time:  $\backslash(O((start\_vids \backslash times (V \log V + E)) \backslash )$

Boost Graph Inside

Signatures

Summary

pgr\_trsp\_withPoints([Edges SQL](#), [Restrictions SQL](#), [Points SQL](#), start\_vid, end\_vid, [options])  
pgr\_trsp\_withPoints([Edges SQL](#), [Restrictions SQL](#), [Points SQL](#), start\_vid, end\_vids, [options])  
pgr\_trsp\_withPoints([Edges SQL](#), [Restrictions SQL](#), [Points SQL](#), start\_vids, end\_vid, [options])  
pgr\_trsp\_withPoints([Edges SQL](#), [Restrictions SQL](#), [Points SQL](#), start\_vids, end\_vids, [options])  
pgr\_trsp\_withPoints([Edges SQL](#), [Restrictions SQL](#), [Combinations SQL](#), [Points SQL](#), [options])  
**options:** [directed, driving\_side, details]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

One to One

pgr\_trsp\_withPoints([Edges SQL](#), [Restrictions SQL](#), [Points SQL](#), start\_vid, end\_vid, [options])  
**options:** [directed, driving\_side, details]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From point  $\backslash(1 \backslash)$  to vertex  $\backslash(10 \backslash)$  with details on a left driving side configuration on a directed graph with details.

```
SELECT * FROM pgr_trsp_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
  $$SELECT id, path, cost FROM restrictions$$,
  $$SELECT pid, edge_id, fraction, side FROM pointsOfInterest$$,
  -1, 10,
  details => true);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -1 | 10 | -1 | 1 | 0.4 | 0
2 | 2 | 1 | -1 | 10 | 5 | 1 | 1 | 0.4
3 | 3 | 1 | -1 | 10 | 6 | 4 | 0.7 | 1.4
4 | 4 | 1 | -1 | 10 | -6 | 4 | 0.3 | 2.1
5 | 5 | 1 | -1 | 10 | 7 | 10 | 1 | 2.4
6 | 6 | 1 | -1 | 10 | 8 | 12 | 0.6 | 3.4
7 | 7 | 1 | -1 | 10 | -3 | 12 | 0.4 | 4
8 | 8 | 1 | -1 | 10 | 12 | 13 | 1 | 4.4
9 | 9 | 1 | -1 | 10 | 17 | 15 | 1 | 5.4
10 | 10 | 1 | -1 | 10 | 16 | 16 | 1 | 6.4
11 | 11 | 1 | -1 | 10 | 15 | 3 | 1 | 7.4
12 | 12 | 1 | -1 | 10 | 10 | -1 | 0 | 8.4
(12 rows)
```

One to Many

pgr\_trsp\_withPoints([Edges SQL](#), [Restrictions SQL](#), [Points SQL](#), start\_vid, end\_vids, [options])  
**options:** [directed, driving\_side, details]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From point  $\backslash(1 \backslash)$  to point  $\backslash(3 \backslash)$  and vertex  $\backslash(7 \backslash)$ .

```
SELECT * FROM pgr_trsp_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
  $$SELECT id, path, cost FROM restrictions$$,
  $$SELECT pid, edge_id, fraction, side FROM pointsOfInterest$$,
  -1, ARRAY{-3, 7});
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -1 | -3 | -1 | 1 | 1.4 | 0
2 | 2 | 1 | -1 | -3 | 6 | 4 | 1 | 1.4
3 | 3 | 1 | -1 | -3 | 7 | 10 | 1 | 2.4
4 | 4 | 1 | -1 | -3 | 8 | 12 | 0.6 | 3.4
5 | 5 | 1 | -1 | -3 | -3 | -1 | 0 | 4
6 | 1 | 1 | -1 | 7 | -1 | 1 | 1.4 | 0
7 | 2 | 1 | -1 | 7 | 6 | 4 | 1 | 1.4
8 | 3 | 1 | -1 | 7 | 7 | -1 | 0 | 2.4
(8 rows)
```

Many to One

pgr\_trsp\_withPoints([Edges SQL](#), [Restrictions SQL](#), [Points SQL](#), start\_vids, end\_vid, [options])  
**options:** [directed, driving\_side, details]

Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From point \1\ and vertex \6\ to point \3\.

```
SELECT * FROM pgr_trsp_withPoints(
$$$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
$$$SELECT id, path, cost FROM restrictions$$,
$$$SELECT pid, edge_id, fraction, side FROM pointsOfInterest$$,
ARRAY[-1, 6, -3]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | -1 | -3 | -1 | 1 | 1.4 | 0
2 | 2 | -1 | -3 | 6 | 4 | 1 | 1.4
3 | 3 | -1 | -3 | 7 | 10 | 1 | 2.4
4 | 4 | -1 | -3 | 8 | 12 | 0.6 | 3.4
5 | 5 | -1 | -3 | -3 | -1 | 0 | 4
6 | 1 | 6 | -3 | 6 | 4 | 1 | 0
7 | 2 | 6 | -3 | 7 | 10 | 1 | 1
8 | 3 | 6 | -3 | 8 | 12 | 0.6 | 2
9 | 4 | 6 | -3 | -3 | -1 | 0 | 2.6
(9 rows)
```

Many to Many

pgr\_trsp\_withPoints([Edges SQL](#), [Restrictions SQL](#), [Points SQL](#), start vids, end vids, [options])  
**options:** [directed, driving\_side, details]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From point \1\ and vertex \6\ to point \3\ and vertex \1\.

```
SELECT * FROM pgr_trsp_withPoints(
$$$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
$$$SELECT id, path, cost FROM restrictions$$,
$$$SELECT pid, edge_id, fraction, side FROM pointsOfInterest$$,
ARRAY[-1, 6], ARRAY[-3, 1]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | -1 | -3 | -1 | 1 | 1.4 | 0
2 | 2 | -1 | -3 | 6 | 4 | 1 | 1.4
3 | 3 | -1 | -3 | 7 | 10 | 1 | 2.4
4 | 4 | -1 | -3 | 8 | 12 | 0.6 | 3.4
5 | 5 | -1 | -3 | -3 | -1 | 0 | 4
6 | 1 | -1 | 1 | -1 | 1 | 1.4 | 0
7 | 2 | -1 | 1 | 6 | 4 | 1 | 1.4
8 | 3 | -1 | 1 | 7 | 10 | 1 | 2.4
9 | 4 | -1 | 1 | 8 | 12 | 1 | 3.4
10 | 5 | -1 | 1 | 12 | 13 | 1 | 4.4
11 | 6 | -1 | 1 | 17 | 15 | 1 | 5.4
12 | 7 | -1 | 1 | 16 | 9 | 1 | 6.4
13 | 8 | -1 | 1 | 11 | 8 | 1 | 7.4
14 | 9 | -1 | 1 | 7 | 7 | 1 | 8.4
15 | 10 | -1 | 1 | 3 | 6 | 1 | 9.4
16 | 11 | -1 | 1 | 1 | -1 | 0 | 10.4
17 | 1 | 6 | -3 | 6 | 4 | 1 | 0
18 | 2 | 6 | -3 | 7 | 10 | 1 | 1
19 | 3 | 6 | -3 | 8 | 12 | 0.6 | 2
20 | 4 | 6 | -3 | -3 | -1 | 0 | 2.6
21 | 1 | 6 | 1 | 6 | 4 | 1 | 0
22 | 2 | 6 | 1 | 7 | 10 | 1 | 1
23 | 3 | 6 | 1 | 8 | 12 | 1 | 2
24 | 4 | 6 | 1 | 12 | 13 | 1 | 3
25 | 5 | 6 | 1 | 17 | 15 | 1 | 4
26 | 6 | 6 | 1 | 16 | 9 | 1 | 5
27 | 7 | 6 | 1 | 11 | 8 | 1 | 6
28 | 8 | 6 | 1 | 7 | 7 | 1 | 7
29 | 9 | 6 | 1 | 3 | 6 | 1 | 8
30 | 10 | 6 | 1 | 1 | -1 | 0 | 9
(30 rows)
```

Combinations

pgr\_trsp\_withPoints([Edges SQL](#), [Restrictions SQL](#), [Combinations SQL](#), [Points SQL](#), [options])  
**options:** [directed, driving\_side, details]  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From point \1\ to vertex \10\ and from vertex \6\ to point \3\ with right side driving configuration.

```
SELECT * FROM pgr_trsp_withPoints(
$$$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
$$$SELECT id, path, cost FROM restrictions$$,
$$$SELECT pid, edge_id, fraction, side FROM pointsOfInterest$$,
$$$SELECT * FROM (VALUES (-1, 10), (6, -3)) AS t(source, target)$$,
driving_side => 'r',
details => true);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | -1 | 10 | -1 | 1 | 0.4 | 0
2 | 2 | -1 | 10 | 5 | 1 | 1 | 0.4
3 | 3 | -1 | 10 | 6 | 4 | 0.7 | 1.4
4 | 4 | -1 | 10 | -6 | 4 | 0.3 | 2.1
5 | 5 | -1 | 10 | 7 | 10 | 1 | 2.4
6 | 6 | -1 | 10 | 8 | 12 | 0.6 | 3.4
7 | 7 | -1 | 10 | -3 | 12 | 0.4 | 4
8 | 8 | -1 | 10 | 12 | 13 | 1 | 4.4
9 | 9 | -1 | 10 | 17 | 15 | 1 | 5.4
10 | 10 | -1 | 10 | 16 | 16 | 1 | 6.4
11 | 11 | -1 | 10 | 15 | 3 | 1 | 7.4
12 | 12 | -1 | 10 | 10 | -1 | 0 | 8.4
13 | 1 | 6 | -3 | 6 | 4 | 0.7 | 0
14 | 2 | 6 | -3 | -6 | 4 | 0.3 | 0.7
15 | 3 | 6 | -3 | 7 | 10 | 1 | 1
16 | 4 | 6 | -3 | 8 | 12 | 0.6 | 2
17 | 5 | 6 | -3 | -3 | -1 | 0 | 2.6
(17 rows)
```

Parameters

Column	Type	Description
--------	------	-------------

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	SQL query as described.
<a href="#">Restrictions SQL</a>	TEXT	SQL query as described.
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	ANY-INTEGER	Identifier of the departure vertex.
start vids	ARRAY [ANY-INTEGER]	Array of identifiers of destination vertices.
end vid	ANY-INTEGER	Identifier of the departure vertex.
end vids	ARRAY [ANY-INTEGER]	Array of identifiers of destination vertices.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Optional parameters<sup>1</sup>

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true the graph is considered <i>Directed</i></li> <li>When false the graph is considered as <i>Undirected</i>.</li> </ul>

With points optional parameters<sup>1</sup>

Parameter	Type	Default	Description
driving_side	CHAR	r	Value in [r, l] indicating if the driving side is: <ul style="list-style-type: none"> <li>r for right driving side</li> <li>l for left driving side</li> <li>Any other value will be considered as r</li> </ul>
details	BOOLEAN	false	<ul style="list-style-type: none"> <li>When true the results will include the points that are in the path.</li> <li>When false the results will not include the points that are in the path.</li> </ul>

Inner Queries<sup>1</sup>

Edges SQL<sup>1</sup>

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Restrictions SQL<sup>1</sup>

Column	Type	Description
path	ARRAY [ANY-INTEGER]	Sequence of edge identifiers that form a path that is not allowed to be taken. - Empty arrays or NULL arrays are ignored. - Arrays that have a NULL element will raise an exception.
Cost	ANY-NUMERICAL	Cost of taking the forbidden path.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Points SQL 1

Parameter	Type	Default	Description
pid	ANY-INTEGER	value	Identifier of the point.
			<ul style="list-style-type: none"><li>Use with positive value, as internally will be converted to negative value</li></ul>
			<ul style="list-style-type: none"><li>If column is present, it can not be NULL.</li></ul>
edge_id	ANY-INTEGER		<ul style="list-style-type: none"><li>If column is not present, a sequential negative <b>value</b> will be given automatically.</li></ul>
			Identifier of the "closest" edge to the point.
fraction	ANY-NUMERICAL		Value in <0,1> that indicates the relative position from the first end point of the edge.
			Value in [b, r, l, NULL] indicating if the point is:
side	CHAR	b	<ul style="list-style-type: none"><li>In the right <i>r</i>,</li></ul>
			<ul style="list-style-type: none"><li>In the left <i>l</i>,</li></ul>
			<ul style="list-style-type: none"><li>In both sides <i>b</i>, NULL</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL 1

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns 1

Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Path identifier.
		<ul style="list-style-type: none"><li>Has value 1 for the first of a path from <i>start_vid</i> to <i>end_vid</i>.</li></ul>
path_seq	INTEGER	Relative position in the path. Has value 1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
node	BIGINT	Identifier of the node in the path from <i>start_vid</i> to <i>end_vid</i> .
edge	BIGINT	Identifier of the edge used to go from <i>node</i> to the next node in the path sequence. -1 for the last node of the path.
cost	FLOAT	Cost to traverse from <i>node</i> using <i>edge</i> to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from <i>start_vid</i> to <i>node</i> .

Additional Examples 1

- [Use pgr\\_findCloseEdges for points on the fly](#)

- [Pass in front or visits.](#)
- [Show details on undirected graph.](#)

Use `pgr_findCloseEdges` for points on the fly!

Using `pgr_findCloseEdges`:

Find the routes from vertex \(-1\) to the two closest locations on the graph of point\((2.9, 1.8)\).

```
SELECT * FROM pgr_trsp_withPoints(
  $e$ SELECT * FROM edges $e$,
  $r$ SELECT id, path, cost FROM restrictions $r$,
  $p$ SELECT edge_id, round(fraction::numeric, 2) AS fraction, side
    FROM pgr_findCloseEdges(
      $$SELECT id, geom FROM edges$$,
      (SELECT ST_POINT(2.9, 1.8)),
      0.5, cap => 2)
  $p$,
  1, ARRAY[-1, -2],
  driving_side => 'r');
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -2 | 1 | 6 | 1 | 0
2 | 2 | 1 | -2 | 3 | 7 | 1 | 1
3 | 3 | 1 | -2 | 7 | 8 | 0.9 | 2
4 | 4 | 1 | -2 | -2 | -1 | 0 | 2.9
5 | 1 | 1 | -1 | 1 | 6 | 1 | 0
6 | 2 | 1 | -1 | 3 | 7 | 1 | 1
7 | 3 | 1 | -1 | 7 | 8 | 2 | 2
8 | 4 | 1 | -1 | 7 | 10 | 1 | 4
9 | 5 | 1 | -1 | 8 | 12 | 1 | 5
10 | 6 | 1 | -1 | 12 | 13 | 1 | 6
11 | 7 | 1 | -1 | 17 | 15 | 1 | 7
12 | 8 | 1 | -1 | 16 | 16 | 1 | 8
13 | 9 | 1 | -1 | 15 | 3 | 1 | 9
14 | 10 | 1 | -1 | 10 | 5 | 0.8 | 10
15 | 11 | 1 | -1 | -1 | -1 | 0 | 10.8
(15 rows)
```

- Point \(-1\) corresponds to the closest edge from point \((2.9, 1.8)\).
- Point \(-2\) corresponds to the next close edge from point \((2.9, 1.8)\).

[Pass in front or visits.](#)

Which path (if any) passes in front of point\((6\) or vertex \((11\) with right side driving topology.

```
SELECT ('(' || start_vid || ' => ' || end_vid || ') at ' || path_seq || 'th step:')::TEXT AS path_at,
CASE WHEN edge = -1 THEN 'visits'
ELSE 'passes in front of'
END AS status,
CASE WHEN node < 0 THEN 'Point'
ELSE 'Vertex'
END AS is_a,
abs(node) AS id
FROM pgr_trsp_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
  $$SELECT id, path, cost FROM restrictions$$,
  $$SELECT pid, edge_id, fraction, side FROM pointsOfInterest$$,
  ARRAY[5, -1], ARRAY[6, -3, -6, 10, 11],
  driving_side => 'r',
  details => true)
WHERE node IN (-6, 11);
path_at | status | is_a | id
-----+-----+-----+-----
(-1 => -6) at 4th step: | visits | Point | 6
(-1 => -3) at 4th step: | passes in front of | Point | 6
(-1 => 10) at 4th step: | passes in front of | Point | 6
(-1 => 11) at 4th step: | passes in front of | Point | 6
(-1 => 11) at 6th step: | visits | Vertex | 11
(5 => -6) at 3th step: | visits | Point | 6
(5 => -3) at 3th step: | passes in front of | Point | 6
(5 => 10) at 3th step: | passes in front of | Point | 6
(5 => 11) at 3th step: | passes in front of | Point | 6
(5 => 11) at 5th step: | visits | Vertex | 11
(10 rows)
```

[Show details on undirected graph.](#)

From point \((1\) and vertex \((6\) to point \((3\) to vertex \((1\) on an undirected graph, with details.

```
SELECT * FROM pgr_trsp_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
  $$SELECT id, path, cost FROM restrictions$$,
  $$SELECT pid, edge_id, fraction, side FROM pointsOfInterest$$,
  ARRAY[-1, 6], ARRAY[-3, 1],
  directed => false,
  details => true);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -3 | -1 | 1 | 0.6 | 0
2 | 2 | 1 | -3 | 6 | 4 | 0.7 | 0.6
3 | 3 | 1 | -3 | -6 | 4 | 0.3 | 1.3
4 | 4 | 1 | -3 | 7 | 10 | 1 | 1.6
5 | 5 | 1 | -3 | 8 | 12 | 0.6 | 2.6
6 | 6 | 1 | -3 | -3 | -1 | 0 | 3.2
7 | 1 | 1 | -1 | -1 | 1 | 0.6 | 0
8 | 2 | 1 | -1 | 1 | 6 | 4 | 0.7 | 0.6
9 | 3 | 1 | 1 | -6 | 4 | 0.3 | 1.3
10 | 4 | 1 | 1 | 7 | 7 | 1 | 1.6
11 | 5 | 1 | 1 | 3 | 6 | 0.7 | 2.6
12 | 6 | 1 | 1 | -4 | 6 | 0.3 | 3.3
13 | 7 | 1 | 1 | 1 | -1 | 0 | 3.6
14 | 1 | 6 | -3 | 6 | 4 | 0.7 | 0
15 | 2 | 6 | -3 | -6 | 4 | 0.3 | 0.7
16 | 3 | 6 | -3 | 7 | 10 | 1 | 1
17 | 4 | 6 | -3 | 8 | 12 | 0.6 | 2
18 | 5 | 6 | -3 | -3 | -1 | 0 | 2.6
19 | 1 | 6 | 1 | 6 | 4 | 0.7 | 0
20 | 2 | 6 | 1 | -6 | 4 | 0.3 | 0.7
21 | 3 | 6 | 1 | 7 | 7 | 1 | 1
22 | 4 | 6 | 1 | 3 | 6 | 0.7 | 2
23 | 5 | 6 | 1 | -4 | 6 | 0.3 | 2.7
24 | 6 | 6 | 1 | 1 | -1 | 0 | 3
(24 rows)
```

See Also

- [TRSP - Family of functions](#)
- [withPoints - Category](#)

- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

**pgr\_trspVia\_withPoints - Proposed**

pgr\_trspVia\_withPoints - Route that goes through a list of vertices and/or points with restrictions.

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Availability

- Version 3.4.0
  - New proposed function.

**Description**

Given a graph, a set of restriction on the graph edges, a set of points on the graphs edges and a list of vertices, this function is equivalent to finding the shortest path between  $(vertex\_i)$  and  $(vertex\_i+1)$  (where  $(vertex)$  can be a vertex or a point on the graph) for all  $(i < size\_of(via\_vertices))$  trying not to use restricted paths.

Route:

is a sequence of paths

Path:

is a section of the route.

The general algorithm is as follows:

- Build the Graph with the new points.
  - The points identifiers will be converted to negative values.
  - The vertices identifiers will remain positive.
- Execute a [pgr\\_withPointsVia - Proposed](#).
- For the set of paths of the solution that pass through a restriction then
  - Execute the **TRSP** algorithm with restrictions for the path.
  - NOTE** when this is done, U\_turn\_on\_edge flag is ignored.

Note

Do not use negative values on identifiers of the inner queries.

Boost Graph Inside

**Signatures**

**One Via**

pgr\_trspVia\_withPoints([Edges SQL](#), [Restrictions SQL](#), [Points SQL](#), **via vertices**, **[options]**)

**options:** [directed, strict, U\_turn\_on\_edge]

Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost, route\_agg\_cost)  
OR EMPTY SET

Example:

Find the route that visits the vertices  $(\{-6, 15, -5\})$  in that order on an directed graph.

```
SELECT * FROM pgr_trspVia_withPoints(
$$$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$,
$$$SELECT path, cost FROM restrictions$,
$$$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$,
ARRAY[-6, 15, -5]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -6 | 15 | 15 | 4 | 0.3 | 0 | 0
2 | 1 | 2 | -6 | 15 | 7 | 10 | 1 | 0.3 | 0.3
3 | 1 | 3 | -6 | 15 | 8 | 12 | 1 | 1.3 | 1.3
4 | 1 | 4 | -6 | 15 | 12 | 13 | 1 | 2.3 | 2.3
5 | 1 | 5 | -6 | 15 | 17 | 15 | 1 | 3.3 | 3.3
6 | 1 | 6 | -6 | 15 | 16 | 16 | 1 | 4.3 | 4.3
7 | 1 | 7 | -6 | 15 | 15 | -1 | 0 | 5.3 | 5.3
8 | 2 | 1 | 15 | -5 | 15 | 3 | 1 | 0 | 5.3
9 | 2 | 2 | 15 | -5 | 10 | 5 | 0.8 | 1 | 6.3
10 | 2 | 3 | 15 | -5 | -5 | -2 | 0 | 1.8 | 7.1
(10 rows)
```

**Parameters**

Parameter	Type	Default	Description
-----------	------	---------	-------------

Parameter	Type	Default	Description
<a href="#">Edges SQL</a>	TEXT		SQL query as described.
<a href="#">Points SQL</a>	TEXT		SQL query as described.
<b>via vertices</b>	ARRAY [ <b>ANY-INTEGER</b> ]		Array of ordered vertices identifiers that are going to be visited. <ul style="list-style-type: none"> <li>When positive it is considered a vertex identifier</li> <li>When negative it is considered a point identifier</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Optional parameters<sup>1</sup>

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true the graph is considered <i>Directed</i></li> <li>When false the graph is considered as <i>Undirected</i>.</li> </ul>

Via optional parameters<sup>1</sup>

Parameter	Type	Default	Description
strict	BOOLEAN	false	<ul style="list-style-type: none"> <li>When true if a path is missing stops and returns <b>EMPTY SET</b></li> <li>When false ignores missing paths returning all paths found</li> </ul>
U_turn_on_edge	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true departing from a visited vertex will not try to avoid</li> </ul>

With points optional parameters<sup>1</sup>

Parameter	Type	Default	Description
driving_side	CHAR	r	Value in [r, l] indicating if the driving side is: <ul style="list-style-type: none"> <li>r for right driving side</li> <li>l for left driving side</li> <li>Any other value will be considered as r</li> </ul>
details	BOOLEAN	false	<ul style="list-style-type: none"> <li>When true the results will include the points that are in the path.</li> <li>When false the results will not include the points that are in the path.</li> </ul>

Inner Queries<sup>1</sup>

Edges SQL<sup>1</sup>

Column	Type	Default	Description
id	<b>ANY-INTEGER</b>		Identifier of the edge.
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
target	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
cost	<b>ANY-NUMERICAL</b>		Weight of the edge (source, target)
reverse_cost	<b>ANY-NUMERICAL</b>	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Restrictions SQL<sup>1</sup>

Column	Type	Description
path	ARRAY [ANY-INTEGER]	Sequence of edge identifiers that form a path that is not allowed to be taken. - Empty arrays on NULL arrays are ignored. - Arrays that have a NULL element will raise an exception.
Cost	ANY-NUMERICAL	Cost of taking the forbidden path.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Points SQL [1](#)

Parameter	Type	Default	Description
			Identifier of the point. <ul style="list-style-type: none"> <li>Use with positive value, as internally will be converted to negative value</li> <li>If column is present, it can not be NULL.</li> <li>If column is not present, a sequential negative <b>value</b> will be given automatically.</li> </ul>
pid	ANY-INTEGER	value	
edge_id	ANY-INTEGER		Identifier of the “closest” edge to the point.
fraction	ANY-NUMERICAL		Value in <0,1> that indicates the relative position from the first end point of the edge.
side	CHAR	b	Value in [b, r, l, NULL] indicating if the point is: <ul style="list-style-type: none"> <li>In the right r,</li> <li>In the left l,</li> <li>In both sides b, NULL</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns [1](#)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Identifier of a path. Has value 1 for the first path.
path_seq	INTEGER	Relative position in the path. Has value 1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex of the path.
end_vid	BIGINT	Identifier of the ending vertex of the path.
node	BIGINT	Identifier of the node in the path from start_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go from node to the next node in the path sequence. <ul style="list-style-type: none"> <li>-1 for the last node of the path.</li> <li>-2 for the last node of the route.</li> </ul>
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.
route_agg_cost	FLOAT	Total cost from start_vid of seq = 1 to end_vid of the current seq.

Note

When start\_vid, end\_vid and node columns have negative values, the identifier is for a Point.

Additional Examples [1](#)

- [Use pgr\\_findCloseEdges for points on the fly](#)
- [Usage variations](#)



- [Aggregate cost of the third path.](#)
- [Route's aggregate cost of the route at the end of the third path.](#)
- [Nodes visited in the route.](#)
- [The aggregate costs of the route when the visited vertices are reached.](#)
- [Status of "passes in front" or "visits" of the nodes and points.](#)
- [Simulation of how algorithm works.](#)

[Use pgr\\_findCloseEdges for points on the fly!](#)

Using [pgr\\_findCloseEdges](#):

Visit from vertex \(-1\)| to the two locations on the graph of point(2.9, 1.8) in order of closeness to the graph.

```
SELECT * FROM pgr_trspVia_withPoints(
  $e$ SELECT * FROM edges $e$,
  $r$ SELECT path, cost FROM restrictions $r$,
  $p$ SELECT edge_id, round(fraction::numeric, 2) AS fraction, side
    FROM pgr_findCloseEdges(
      $$SELECT id, geom FROM edges$$,
      (SELECT ST_POINT(2.9, 1.8)),
      0.5, cap => 2)
  $p$,
  ARRAY[1, -1, -2], details => true);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | -1 | 1 | 6 | 1 | 0 | 0
2 | 1 | 2 | 1 | -1 | 3 | 7 | 1 | 1 | 1
3 | 1 | 3 | 1 | -1 | 7 | 8 | 0.9 | 2 | 2
4 | 1 | 4 | 1 | -1 | 2 | 8 | 0.1 | 2.9 | 2.9
5 | 1 | 5 | 1 | -1 | 11 | 8 | 1 | 3 | 3
6 | 1 | 6 | 1 | -1 | 7 | 10 | 1 | 4 | 4
7 | 1 | 7 | 1 | -1 | 8 | 12 | 1 | 5 | 5
8 | 1 | 8 | 1 | -1 | 12 | 13 | 1 | 6 | 6
9 | 1 | 9 | 1 | -1 | 17 | 15 | 1 | 7 | 7
10 | 1 | 10 | 1 | -1 | 16 | 16 | 1 | 8 | 8
11 | 1 | 11 | 1 | -1 | 15 | 3 | 1 | 9 | 9
12 | 1 | 12 | 1 | -1 | 10 | 5 | 0.8 | 10 | 10
13 | 1 | 13 | 1 | -1 | -1 | -1 | 0 | 10.8 | 10.8
14 | 2 | 1 | -1 | -2 | -1 | 5 | 0.2 | 0 | 10.8
15 | 2 | 2 | -1 | -2 | 11 | 8 | 1 | 0.2 | 11
16 | 2 | 3 | -1 | -2 | 7 | 8 | 0.9 | 1.2 | 12
17 | 2 | 4 | -1 | -2 | -2 | -2 | 0 | 2.1 | 12.9
(17 rows)
```

- Point \(-1\)| corresponds to the closest edge from point(2.9, 1.8).
- Point \(-2\)| corresponds to the next close edge from point(2.9, 1.8).
- Point \(-2\)| is visited on the route to from vertex \(-1\)| to Point \(-1\)| (See row where \((seq = 4)\)).

[Usage variations!](#)

All this examples are about the route that visits the vertices \((-6, 7, -4, 8, -2\)|) in that order on a directed graph.

```
SELECT * FROM pgr_trspVia_withPoints(
  $$$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
  $$$SELECT path, cost FROM restrictions$$,
  $$$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$$,
  ARRAY[-6, 7, -4, 8, -2]
);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | -6 | 7 | -6 | 4 | 0.3 | 0
2 | 1 | 2 | 1 | -6 | 7 | 7 | -1 | 0 | 0.3
3 | 2 | 1 | 7 | -4 | 7 | 7 | 1 | 0 | 0.3
4 | 2 | 2 | 7 | -4 | 3 | 6 | 1.3 | 1 | 1.3
5 | 2 | 3 | 7 | -4 | -4 | -1 | 0 | 2.3 | 2.6
6 | 3 | 1 | -4 | 8 | -4 | 6 | 0.7 | 0 | 2.6
7 | 3 | 2 | -4 | 8 | 3 | 7 | 1 | 0.7 | 3.3
8 | 3 | 3 | -4 | 8 | 7 | 4 | 0.6 | 1.7 | 4.3
9 | 3 | 4 | -4 | 8 | 7 | 10 | 1 | 2.3 | 4.9
10 | 3 | 5 | -4 | 8 | 8 | -1 | 0 | 3.3 | 5.9
11 | 4 | 1 | 8 | -2 | 8 | 10 | 1 | 0 | 5.9
12 | 4 | 2 | 8 | -2 | 7 | 8 | 1 | 1 | 6.9
13 | 4 | 3 | 8 | -2 | 11 | 9 | 1 | 2 | 7.9
14 | 4 | 4 | 8 | -2 | 16 | 15 | 0.4 | 3 | 8.9
15 | 4 | 5 | 8 | -2 | -2 | -2 | 0 | 3.4 | 9.3
(15 rows)
```

[Aggregate cost of the third path!](#)

```
SELECT agg_cost FROM pgr_trspVia_withPoints(
  $$$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
  $$$SELECT path, cost FROM restrictions$$,
  $$$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$$,
  ARRAY[-6, 7, -4, 8, -2]
)
WHERE path_id = 3 AND edge < 0;
agg_cost
-----
3.3
(1 row)
```

[Route's aggregate cost of the route at the end of the third path!](#)

```
SELECT route_agg_cost FROM pgr_trspVia_withPoints(
  $$$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
  $$$SELECT path, cost FROM restrictions$$,
  $$$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$$,
  ARRAY[-6, 7, -4, 8, -2]
)
WHERE path_id = 3 AND edge < 0;
route_agg_cost
-----
5.9
(1 row)
```

[Nodes visited in the route.](#)

```
SELECT row_number() over () as node_seq, node
FROM pgr_trspVia_withPoints(
  $$$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
  $$$SELECT path, cost FROM restrictions$$,
  $$$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$$,
  ARRAY[-6, 7, -4, 8, -2]
)
```

```
)
WHERE edge <> -1 ORDER BY seq;
node_seq | node
-----+-----
1 | -6
2 | 7
3 | 3
4 | -4
5 | 3
6 | 7
7 | 7
8 | 8
9 | 7
10 | 11
11 | 16
12 | -2
(12 rows)
```

The aggregate costs of the route when the visited vertices are reached.

```
SELECT path_id, route_agg_cost FROM pgr_trspVia_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$,
  $$SELECT path, cost FROM restrictions$,
  $$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$,
  ARRAY[-6, 7, -4, 8, -2]
)
WHERE edge < 0;
path_id | route_agg_cost
-----+-----
1 | 0.3
2 | 2.6
3 | 5.9
4 | 9.3
(4 rows)
```

Status of "passes in front" or "visits" of the nodes and points.

```
SELECT seq, route_agg_cost, node, agg_cost ,
CASE WHEN edge = -1 THEN $$visits$$
ELSE $$passes in front$$
END as status
FROM pgr_trspVia_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$,
  $$SELECT path, cost FROM restrictions$,
  $$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$,
  ARRAY[-6, 7, -4, 8, -2])
WHERE agg_cost <> 0 or seq = 1;
seq | route_agg_cost | node | agg_cost | status
-----+-----+-----+-----+-----
1 | 0 | -6 | 0 | passes in front
2 | 0.3 | 7 | 0.3 | visits
4 | 1.3 | 3 | 1 | passes in front
5 | 2.6 | -4 | 2.3 | visits
7 | 3.3 | 3 | 0.7 | passes in front
8 | 4.3 | 7 | 1.7 | passes in front
9 | 4.9 | 7 | 2.3 | passes in front
10 | 5.9 | 8 | 3.3 | visits
12 | 6.9 | 7 | 1 | passes in front
13 | 7.9 | 11 | 2 | passes in front
14 | 8.9 | 16 | 3 | passes in front
15 | 9.3 | -2 | 3.4 | passes in front
(12 rows)
```

Simulation of how algorithm works.

The algorithm performs a [pgr\\_withPointsVia - Proposed](#)

```
SELECT * FROM pgr_withPointsVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$,
  $$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$,
  ARRAY[-6, 15, -5]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -6 | 15 | -6 | 4 | 0.3 | 0 | 0
2 | 1 | 2 | -6 | 15 | 7 | 8 | 1 | 0.3 | 0.3
3 | 1 | 3 | -6 | 15 | 11 | 9 | 1 | 1.3 | 1.3
4 | 1 | 4 | -6 | 15 | 16 | 16 | 1 | 2.3 | 2.3
5 | 1 | 5 | -6 | 15 | 15 | -1 | 0 | 3.3 | 3.3
6 | 2 | 1 | 15 | -5 | 15 | 3 | 1 | 0 | 3.3
7 | 2 | 2 | 15 | -5 | 10 | 5 | 0.8 | 1 | 4.3
8 | 2 | 3 | 15 | -5 | -5 | -2 | 0 | 1.8 | 5.1
(8 rows)
```

Detects which of the paths pass through a restriction in this case is for the path\_id = 1 from -6 to 15 because the path(9 \rightarrow 16) is restricted.

Executes the [TRSP algorithm](#) for the conflicting paths.

```
SELECT 1 AS path_id, path_seq, start_vid, end_vid, node, edge, cost, agg_cost
FROM pgr_trsp_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$,
  $$SELECT path, cost FROM restrictions$,
  $$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$,
  -6, 15);
path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | -6 | 15 | -6 | 4 | 0.3 | 0
1 | 2 | -6 | 15 | 7 | 10 | 1 | 0.3
1 | 3 | -6 | 15 | 8 | 12 | 1 | 1.3
1 | 4 | -6 | 15 | 12 | 13 | 1 | 2.3
1 | 5 | -6 | 15 | 17 | 15 | 1 | 3.3
1 | 6 | -6 | 15 | 16 | 16 | 1 | 4.3
1 | 7 | -6 | 15 | 15 | -1 | 0 | 5.3
(7 rows)
```

From the [pgr\\_withPointsVia - Proposed](#) result it removes the conflicting paths and builds the solution with the results of the [pgr\\_trsp - Proposed](#) algorithm:

```
WITH
solutions AS (
  SELECT path_id, path_seq, start_vid, end_vid, node, edge, cost, agg_cost
  FROM pgr_withPointsVia(
    $$SELECT id, source, target, cost, reverse_cost FROM edges$,
    $$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$,
    ARRAY[-6, 15, -5] WHERE path_id != 1
  )
  UNION
  SELECT 1 AS path_id, path_seq, start_vid, end_vid, node, edge, cost, agg_cost
  FROM pgr_trsp_withPoints(
    $$SELECT id, source, target, cost, reverse_cost FROM edges$,
    $$SELECT path, cost FROM restrictions$,
    $$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$,
    -6, 15);
with_seq AS (
  SELECT row_number() over(ORDER BY path_id, path_seq) AS seq, *
```

```
FROM solutions),
aggregation AS (SELECT seq, SUM(cost) OVER(ORDER BY seq) AS route_agg_cost FROM with_seq)
SELECT with_seq.*, COALESCE(route_agg_cost, 0) AS route_agg_cost
FROM with_seq LEFT JOIN aggregation ON (with_seq.seq = aggregation.seq + 1);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -6 | 15 | -6 | 4 | 0.3 | 0 | 0
2 | 1 | 2 | -6 | 15 | 7 | 10 | 1 | 0.3 | 0.3
3 | 1 | 3 | -6 | 15 | 8 | 12 | 1 | 1.3 | 1.3
4 | 1 | 4 | -6 | 15 | 12 | 13 | 1 | 2.3 | 2.3
5 | 1 | 5 | -6 | 15 | 17 | 15 | 1 | 3.3 | 3.3
6 | 1 | 6 | -6 | 15 | 16 | 16 | 1 | 4.3 | 4.3
7 | 1 | 7 | -6 | 15 | 15 | -1 | 0 | 5.3 | 5.3
8 | 2 | 1 | 15 | -5 | 15 | 3 | 1 | 0 | 5.3
9 | 2 | 2 | 15 | -5 | 10 | 5 | 0.8 | 1 | 6.3
10 | 2 | 3 | 15 | -5 | -5 | -2 | 0 | 1.8 | 7.1
(10 rows)
```

Getting the same result as `pgr_trspVia_withPoints`:

```
SELECT * FROM pgr_trspVia_withPoints(
$$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
$$SELECT path, cost FROM restrictions$$,
$$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$$,
ARRAY[6, 15, -5]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -6 | 15 | -6 | 4 | 0.3 | 0 | 0
2 | 1 | 2 | -6 | 15 | 7 | 10 | 1 | 0.3 | 0.3
3 | 1 | 3 | -6 | 15 | 8 | 12 | 1 | 1.3 | 1.3
4 | 1 | 4 | -6 | 15 | 12 | 13 | 1 | 2.3 | 2.3
5 | 1 | 5 | -6 | 15 | 17 | 15 | 1 | 3.3 | 3.3
6 | 1 | 6 | -6 | 15 | 16 | 16 | 1 | 4.3 | 4.3
7 | 1 | 7 | -6 | 15 | 15 | -1 | 0 | 5.3 | 5.3
8 | 2 | 1 | 15 | -5 | 15 | 3 | 1 | 0 | 5.3
9 | 2 | 2 | 15 | -5 | 10 | 5 | 0.8 | 1 | 6.3
10 | 2 | 3 | 15 | -5 | -5 | -2 | 0 | 1.8 | 7.1
(10 rows)
```

Example 8:

Sometimes `U_turn_on_edge` flag is ignored when is set to false.

The first step, doing a [pgr\\_withPointsVia - Proposed](#) does consider not making a U turn on the same edge. But the path(9 → 16) (Rows 4 and 5) is restricted and the result is using it.

```
SELECT * FROM pgr_withPointsVia(
$$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
$$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$$,
ARRAY[6, 7, 6], U_turn_on_edge => false);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 6 | 7 | 6 | 4 | 1 | 0 | 0
2 | 1 | 2 | 6 | 7 | 7 | -1 | 0 | 1 | 1
3 | 2 | 1 | 7 | 6 | 7 | 8 | 1 | 0 | 1
4 | 2 | 2 | 7 | 6 | 11 | 9 | 1 | 1 | 2
5 | 2 | 3 | 7 | 6 | 16 | 16 | 1 | 2 | 3
6 | 2 | 4 | 7 | 6 | 15 | 3 | 1 | 3 | 4
7 | 2 | 5 | 7 | 6 | 10 | 2 | 1 | 4 | 5
8 | 2 | 6 | 7 | 6 | 6 | -2 | 0 | 5 | 6
(8 rows)
```

When executing the [pgr\\_trsp\\_withPoints - Proposed](#) algorithm for the conflicting path, there is no `U_turn_on_edge` flag.

```
SELECT 5 AS path_id, path_seq, start_vid, end_vid, node, edge, cost, agg_cost
FROM pgr_trsp_withPoints(
$$SELECT id, source, target, cost, reverse_cost FROM edges$$,
$$SELECT path, cost FROM restrictions$$,
$$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$$,
7, 6);
path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
5 | 1 | 7 | 6 | 7 | 4 | 1 | 0
5 | 2 | 7 | 6 | 6 | -1 | 0 | 1
(2 rows)
```

Therefore the result ignores the `U_turn_on_edge` flag when set to false. From the [pgr\\_withPointsVia - Proposed](#) result it removes the conflicting paths and builds the solution with the results of the [pgr\\_trsp - Proposed](#) algorithm. In this case a U turn is been done using the same edge.

```
SELECT * FROM pgr_trspVia_withPoints(
$$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
$$SELECT path, cost FROM restrictions$$,
$$SELECT pid, edge_id, side, fraction FROM pointsOfInterest$$,
ARRAY[6, 7, 6], U_turn_on_edge => false);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 6 | 7 | 6 | 4 | 1 | 0 | 0
2 | 1 | 2 | 6 | 7 | 7 | -1 | 0 | 1 | 1
3 | 2 | 1 | 7 | 6 | 7 | 4 | 1 | 0 | 1
4 | 2 | 2 | 7 | 6 | 6 | -2 | 0 | 1 | 2
(4 rows)
```

See Also

- [TRSP - Family of functions](#)
- [Via - Category](#)
- [withPoints - Category](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

`pgr_turnRestrictedPath` - Experimental

`pgr_turnRestrictedPath` Using Yen's algorithm Vertex - Vertex routing with restrictions

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.0.0
  - New experimental function.

Description

Using Yen's algorithm to obtain K shortest paths and analyze the paths to select the paths that do not use the restrictions

Boost Graph Inside

Signatures

`pgr_turnRestrictedPath(Edges SQL, Restrictions SQL, start vid, end vid, K, [options])`  
**options:** [directed, heap\_paths, stop\_on\_first, strict]  
Returns set of (seq, path\_id, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \3\ to vertex \8\ on a directed graph

```
SELECT * FROM pgr_turnRestrictedPath(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  3, 8, 3);
seq | path_id | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 3 | 7 | 1 | Infinity
2 | 1 | 2 | 7 | 10 | 1 | 1
3 | 1 | 3 | 8 | -1 | 0 | 2
(3 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	SQL query as described.
start vid	ANY-INTEGER	Identifier of the departure vertex.
end vid	ANY-INTEGER	Identifier of the destination vertex.
K	ANY-INTEGER	Number of required paths.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>• When true the graph is considered <i>Directed</i></li><li>• When false the graph is considered as <i>Undirected</i>.</li></ul>

KSP Optional parameters

Column	Type	Default	Description
--------	------	---------	-------------

Column	Type	Default	Description
heap_paths	BOOLEAN	false	<ul style="list-style-type: none"> <li>When false Returns at most K paths.</li> <li>When true all the calculated paths while processing are returned.</li> <li>Roughly, when the shortest path has N edges, the heap will contain about than <math>N * K</math> paths for small value of K and <math>K &gt; 5</math>.</li> </ul>

Special optional parameters

Column	Type	Default	Description
stop_on_first	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true stops on first path found that dos not violate restrictions</li> <li>When false returns at most K paths</li> </ul>
strict	BOOLEAN	false	<ul style="list-style-type: none"> <li>When true returns only paths that do not violate restrictions</li> <li>When false returns the paths found</li> </ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Restrictions SQL

Column	Type	Description
path	ARRAY [ANY-INTEGER]	Sequence of edge identifiers that form a path that is not allowed to be taken. - Empty arrays oNULL arrays are ignored. - Arrays that have a NULL element will raise an exception.
Cost	ANY-NUMERICAL	Cost of taking the forbidden path.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Path identifier. <ul style="list-style-type: none"> <li>Has value 1 for the first of a path fromstart_vid to end_vid.</li> </ul>
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.

Column	Type	Description
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Additional Examples¶

Example:

From vertex \3\ to \8\ with strict flag on.

No results because the only path available follows a restriction.

```
SELECT * FROM pgr_turnRestrictedPath(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  3, 8, 3,
  strict => true);
seq | path_id | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
(0 rows)
```

Example:

From vertex \3\ to vertex \8\ on an undirected graph

```
SELECT * FROM pgr_turnRestrictedPath(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  3, 8, 3,
  directed => false);
seq | path_id | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | 3 | 7 | 1 | 0
2 | 1 | 2 | 7 | 4 | 1 | 1
3 | 1 | 3 | 6 | 2 | 1 | 2
4 | 1 | 4 | 10 | 5 | 1 | 3
5 | 1 | 5 | 11 | 11 | 1 | 4
6 | 1 | 6 | 12 | 12 | 1 | 5
7 | 1 | 7 | 8 | -1 | 0 | 6
(7 rows)
```

Example:

From vertex \3\ to vertex \8\ with more alternatives

```
SELECT * FROM pgr_turnRestrictedPath(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT path, cost FROM restrictions$$,
  3, 8, 3,
  directed => false,
  heap_paths => true,
  stop_on_first => false);
seq | path_id | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | 3 | 7 | 1 | 0
2 | 1 | 2 | 7 | 4 | 1 | 1
3 | 1 | 3 | 6 | 2 | 1 | 2
4 | 1 | 4 | 10 | 5 | 1 | 3
5 | 1 | 5 | 11 | 11 | 1 | 4
6 | 1 | 6 | 12 | 12 | 1 | 5
7 | 1 | 7 | 8 | -1 | 0 | 6
8 | 2 | 1 | 3 | 7 | 1 | 0
9 | 2 | 2 | 7 | 8 | 1 | 1
10 | 2 | 3 | 11 | 9 | 1 | 2
11 | 2 | 4 | 16 | 15 | 1 | 3
12 | 2 | 5 | 17 | 13 | 1 | 4
13 | 2 | 6 | 12 | 12 | 1 | 5
14 | 2 | 7 | 8 | -1 | 0 | 6
(14 rows)
```

See Also¶

- [K shortest paths - Category](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Introduction¶

Road restrictions are a sequence of road segments that can not be taken in a sequential manner. Some restrictions are implicit on a directed graph, for example, one way roads where the wrong way edge is not even inserted on the graph. But normally on turns like no left turn or no right turn, hence the name turn restrictions, there are sometimes restrictions.

TRSP algorithm¶

The internal TRSP algorithm performs a lookahead over the dijkstra algorithm in order to find out if the attempted path has a restriction. This allows the algorithm to pass twice on the same vertex.

Parameters¶

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> query as described.

Parameter	Type	Description
<a href="#">Restrictions SQL</a>	TEXT	<a href="#">Restrictions SQL</a> query as described.
<b>via vertices</b>	ARRAY[ <b>ANY-INTEGER</b> ]	Array of ordered vertices identifiers that are going to be visited.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

#### Restrictions¶

On road networks, there are restrictions such as left or right turn restrictions, no U turn restrictions.

A restriction is a sequence of edges, called path and that path is to be avoided.

[images/with\\_restrictions.png](#)



#### Restrictions on the road network¶

These restrictions are represented on a table as follows:

```
/* -- r1 */
CREATE TABLE restrictions (
  id SERIAL PRIMARY KEY,
  path BIGINT[],
  cost FLOAT
);
/* -- r2 */
INSERT INTO restrictions (path, cost) VALUES
  (ARRAY[4, 7], 100),
  (ARRAY[8, 11], 100),
  (ARRAY[7, 10], 100),
  (ARRAY[3, 5, 9], 4),
  (ARRAY[9, 16], 100);
/* -- r3 */
SELECT * FROM restrictions;
/* -- r4 */
```

Note

The table has an identifier, which maybe is needed for the administration of the restrictions, but the algorithms do not need that information. If given it will be ignored.

#### Edges SQL¶

Column	Type	Default	Description
id	<b>ANY-INTEGER</b>		Identifier of the edge.
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
target	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
cost	<b>ANY-NUMERICAL</b>		Weight of the edge (source, target)
reverse_cost	<b>ANY-NUMERICAL</b>	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

#### Restrictions SQL¶

Column	Type	Description
path	ARRAY [ <b>ANY-INTEGER</b> ]	Sequence of edge identifiers that form a path that is not allowed to be taken. - Empty arrays on NULL arrays are ignored. - Arrays that have a NULL element will raise an exception.

Column	Type	Description
Cost	<b>ANY-NUMERICAL</b>	Cost of taking the forbidden path.
Where:		
ANY-INTEGER:		
SMALLINT, INTEGER, BIGINT		
ANY-NUMERICAL:		
SMALLINT, INTEGER, BIGINT, REAL, FLOAT		
See Also		
Indices and tables		
<ul style="list-style-type: none"> <li><a href="#">Index</a></li> <li><a href="#">Search Page</a></li> </ul>		
<a href="#">Transformation - Family of functions</a>		
<ul style="list-style-type: none"> <li><a href="#">pgr_lineGraph - Proposed</a> - Transformation algorithm for generating a Line Graph.</li> </ul>		
<a href="#">Coloring - Family of functions</a>		
<ul style="list-style-type: none"> <li><a href="#">pgr_sequentialVertexColoring - Proposed</a> - Vertex coloring algorithm using greedy approach.</li> </ul>		
<a href="#">Contraction - Family of functions</a>		
<ul style="list-style-type: none"> <li><a href="#">pgr_contractionDeadEnd - Proposed</a></li> <li><a href="#">pgr_contractionLinear - Proposed</a></li> </ul>		
<a href="#">Traversal - Family of functions</a>		
<ul style="list-style-type: none"> <li><a href="#">pgr_depthFirstSearch - Proposed</a> - Depth first search traversal of the graph.</li> </ul>		
<b>Traversal - Family of functions</b>		
<input type="checkbox"/> Proposed		
Warning		
Proposed functions for next mayor release.		
<ul style="list-style-type: none"> <li>They are not officially in the current release.</li> <li>They will likely officially be part of the next mayor release:             <ul style="list-style-type: none"> <li>The functions make use of ANY-INTEGER and ANY-NUMERICAL</li> <li>Name might not change. (But still can)</li> <li>Signature might not change. (But still can)</li> <li>Functionality might not change. (But still can)</li> <li>pgTap tests have being done. But might need more.</li> <li>Documentation might need refinement.</li> </ul> </li> <li><a href="#">pgr_depthFirstSearch - Proposed</a> - Depth first search traversal of the graph.</li> </ul>		
<input type="checkbox"/> Experimental		
Warning		
Possible server crash		
<ul style="list-style-type: none"> <li>These functions might create a server crash</li> </ul>		
Warning		
Experimental functions		
<ul style="list-style-type: none"> <li>They are not officially of the current release.</li> <li>They likely will not be officially be part of the next release:             <ul style="list-style-type: none"> <li>The functions might not make use of ANY-INTEGER and ANY-NUMERICAL</li> <li>Name might change.</li> <li>Signature might change.</li> <li>Functionality might change.</li> <li>pgTap tests might be missing.</li> <li>Might need c/c++ coding.</li> <li>May lack documentation.</li> <li>Documentation if any might need to be rewritten.</li> <li>Documentation examples might need to be automatically generated.</li> <li>Might need a lot of feedback from the community.</li> <li>Might depend on a proposed function of pgRouting</li> <li>Might depend on a deprecated function of pgRouting</li> </ul> </li> <li><a href="#">pgr_breadthFirstSearch - Experimental</a> - Breath first search traversal of the graph.</li> <li><a href="#">pgr_binaryBreadthFirstSearch - Experimental</a> - Breath first search traversal of the graph.</li> </ul>		
Additionally there are 2 categories under this family		



- [BFS - Category](#)
- [DFS - Category](#)

`pgr_depthFirstSearch` - **Proposed**

`pgr_depthFirstSearch` — Returns a depth first search traversal of the graph. The graph can be directed or undirected.

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Availability

- Version 3.3.0
  - Function promoted to proposed.
- Version 3.2.0
  - New experimental function.

Description

Depth First Search algorithm is a traversal algorithm which starts from a root vertex, goes as deep as possible, and backtracks once a vertex is reached with no adjacent vertices or with all visited adjacent vertices. The traversal continues until all the vertices reachable from the root vertex are visited.

The main Characteristics are:

- The implementation works for both **directed** and **undirected** graphs.
- Provides the Depth First Search traversal order from a root vertex or from a set of root vertices.
- An optional non-negative maximum depth parameter to limit the results up to a particular depth.
- For optimization purposes, any duplicated values in the *Root vids* are ignored.
- It does not produce the shortest path from a root vertex to a target vertex.
- The aggregate cost of traversal is not guaranteed to be minimal.
- The returned values are ordered in ascending order of *start\_vid*.
- Depth First Search Running time:  $\backslash(O(E + V))$

 Boost Graph Inside

Signatures

Summary

`pgr_depthFirstSearch`([Edges SQL](#), **root vid**, [**options**])  
`pgr_depthFirstSearch`([Edges SQL](#), **root vids**, [**options**])  
**options**: [directed, max\_depth]  
Returns set of (seq, depth, start\_vid, node, edge, cost, agg\_cost)

Single vertex

`pgr_depthFirstSearch`([Edges SQL](#), **root vid**, [**options**])  
**options**: [directed, max\_depth]  
Returns set of (seq, depth, start\_vid, node, edge, cost, agg\_cost)

Example:

From root vertex  $\backslash(6)$  on a **directed** graph with edges in ascending order ofid

```
SELECT * FROM pgr_depthFirstSearch(
'SELECT id, source, target, cost, reverse_cost FROM edges
ORDER BY id',
6);
```

seq	depth	start_vid	node	edge	cost	agg_cost
1	0	6	6	-1	0	0
2	1	6	5	1	1	1
3	1	6	7	4	1	1
4	2	6	3	7	1	2
5	3	6	1	6	1	3
6	2	6	11	8	1	2
7	3	6	16	9	1	3
8	4	6	17	15	1	4
9	4	6	15	16	1	4
10	5	6	10	3	1	5
11	3	6	12	11	1	3
12	2	6	8	10	1	2
13	3	6	9	14	1	3

(13 rows)

Multiple vertices

`pgr_depthFirstSearch`([Edges SQL](#), **root vids**, [**options**])  
**options**: [directed, max\_depth]  
Returns set of (seq, depth, start\_vid, node, edge, cost, agg\_cost)

Example:

From root vertices  $\backslash(\{12, 6\})$  on an **undirected** graph with **depth**  $\backslash(<= 2)$  and edges in ascending order ofid

```
SELECT * FROM pgr_depthFirstSearch(
  'SELECT id, source, target, cost, reverse_cost FROM edges
  ORDER BY id',
  ARRAY[12, 6], directed => false, max_depth => 2);
seq | depth | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 15 | 3 | 1 | 2
5 | 2 | 6 | 11 | 5 | 1 | 2
6 | 1 | 6 | 7 | 4 | 1 | 1
7 | 2 | 6 | 3 | 7 | 1 | 2
8 | 2 | 6 | 8 | 10 | 1 | 2
9 | 0 | 12 | 12 | -1 | 0 | 0
10 | 1 | 12 | 11 | 11 | 1 | 1
11 | 2 | 12 | 10 | 5 | 1 | 2
12 | 2 | 12 | 7 | 8 | 1 | 2
13 | 2 | 12 | 16 | 9 | 1 | 2
14 | 1 | 12 | 8 | 12 | 1 | 1
15 | 2 | 12 | 9 | 14 | 1 | 2
16 | 1 | 12 | 17 | 13 | 1 | 1
(16 rows)
```

Parameters

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below.
<b>root vid</b>	BIGINT	Identifier of the root vertex of the tree. <ul style="list-style-type: none"><li>When value is \0 then gets the spanning forest starting in aleatory nodes for each tree in the forest.</li></ul>
<b>root vids</b>	ARRAY [ <b>ANY-INTEGER</b> ]	Array of identifiers of the root vertices. <ul style="list-style-type: none"><li>\0 values are ignored</li><li>For optimization purposes, any duplicated value is ignored.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERIC:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT, NUMERIC

Optional parameters

Column	Type	Default	Description
<b>directed</b>	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

DFS optional parameters

Parameter	Type	Default	Description
<b>max_depth</b>	BIGINT	\(9223372036854775807\)	Upper limit of the depth of the tree. <ul style="list-style-type: none"><li>When negative throws an error.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
<b>id</b>	<b>ANY-INTEGER</b>		Identifier of the edge.
<b>source</b>	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
<b>target</b>	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
<b>cost</b>	<b>ANY-NUMERICAL</b>		Weight of the edge (source, target)
<b>reverse_cost</b>	<b>ANY-NUMERICAL</b>	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns1

Returns set of (seq, depth, start\_vid, node, edge, cost, agg\_cost)

Parameter	Type	Description
seq	BIGINT	Sequential value starting from \(\backslash\).
depth	BIGINT	Depth of the node. <ul style="list-style-type: none"><li>\(0\) when node = start_vid.</li></ul>
start_vid	BIGINT	Identifier of the root vertex.
node	BIGINT	Identifier of node reached using edge.
edge	BIGINT	Identifier of the edge used to arrive to node. <ul style="list-style-type: none"><li>\(-1\) when node = start_vid.</li></ul>
cost	FLOAT	Cost to traverse edge.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERIC:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT, NUMERIC

Additional Examples1

Example:

Same as [Single vertex](#) but with edges in descending order ofid.

```
SELECT * FROM pgr_depthFirstSearch(
'SELECT id, source, target, cost, reverse_cost FROM edges
ORDER BY id DESC',
6);
seq | depth | start_vid | node | edge | cost | agg_cost
-----
1 | 0 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 7 | 4 | 1 | 1
3 | 2 | 6 | 8 | 10 | 1 | 2
4 | 3 | 6 | 9 | 14 | 1 | 3
5 | 3 | 6 | 12 | 12 | 1 | 3
6 | 4 | 6 | 17 | 13 | 1 | 4
7 | 5 | 6 | 16 | 15 | 1 | 5
8 | 6 | 6 | 15 | 16 | 1 | 6
9 | 7 | 6 | 10 | 3 | 1 | 7
10 | 8 | 6 | 11 | 5 | 1 | 8
11 | 2 | 6 | 3 | 7 | 1 | 2
12 | 3 | 6 | 1 | 6 | 1 | 3
13 | 1 | 6 | 5 | 1 | 1 | 1
(13 rows)
```

The resulting traversal is different.

The left image shows the result with ascending order of ids and the right image shows with descending order of the edge identifiers.

ascending

descending

See Also1

- DFS - Category
- Sample Data
- Boost: Depth First Search
- Boost: Undirected DFS
- Wikipedia: Depth First Search algorithm

Indices and tables

- Index

- [Search Page](#)

pgr\_breadthFirstSearch - Experimental

pgr\_breadthFirstSearch — Returns the traversal order(s) using Breadth First Search algorithm.

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.0.0
  - New experimental function.

Description

Provides the Breadth First Search traversal order from a root vertex to a particular depth.

The main Characteristics are:

- The implementation will work on any type of graph.
- Provides the Breadth First Search traversal order from a source node to a target depth level.
- Running time:  $\backslash(O(E + V)\backslash)$

Boost Graph Inside

Signatures

Summary

pgr\_breadthFirstSearch([Edges SQL](#), root vid, [options])  
pgr\_breadthFirstSearch([Edges SQL](#), root vids, [options])  
**options:** [max\_depth, directed]  
Returns set of (seq, depth, start\_vid, node, edge, cost, agg\_cost)

Single vertex

pgr\_breadthFirstSearch([Edges SQL](#), root vid, [options])  
**options:** [max\_depth, directed]  
Returns set of (seq, depth, start\_vid, node, edge, cost, agg\_cost)

Example:

From root vertex  $\backslash(6)\backslash$  on a **directed** graph with edges in ascending order ofid

```
SELECT * FROM pgr_breadthFirstSearch(
'SELECT id, source, target, cost, reverse_cost
FROM edges ORDER BY id',
6);
seq | depth | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 7 | 4 | 1 | 1
4 | 2 | 6 | 3 | 7 | 1 | 2
5 | 2 | 6 | 11 | 8 | 1 | 2
6 | 2 | 6 | 8 | 10 | 1 | 2
7 | 3 | 6 | 1 | 6 | 1 | 3
8 | 3 | 6 | 16 | 9 | 1 | 3
9 | 3 | 6 | 12 | 11 | 1 | 3
10 | 3 | 6 | 9 | 14 | 1 | 3
11 | 4 | 6 | 17 | 15 | 1 | 4
12 | 4 | 6 | 15 | 16 | 1 | 4
13 | 5 | 6 | 10 | 3 | 1 | 5
(13 rows)
```

Multiple vertices

pgr\_breadthFirstSearch([Edges SQL](#), root vids, [options])  
**options:** [max\_depth, directed]  
Returns set of (seq, depth, start\_vid, node, edge, cost, agg\_cost)

Example:

From root vertices  $\backslash(12, 6\backslash)$  on an **undirected** graph with **depth**  $\backslash(<= 2\backslash)$  and edges in ascending order of id

```
SELECT * FROM pgr_breadthFirstSearch(
'SELECT id, source, target, cost, reverse_cost
FROM edges ORDER BY id',
ARRAY[12, 6], directed => false, max_depth => 2);
seq | depth | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 10 | 2 | 1 | 1
4 | 1 | 6 | 7 | 4 | 1 | 1
5 | 2 | 6 | 15 | 3 | 1 | 2
6 | 2 | 6 | 11 | 5 | 1 | 2
7 | 2 | 6 | 3 | 7 | 1 | 2
8 | 2 | 6 | 8 | 10 | 1 | 2
9 | 0 | 12 | 12 | -1 | 0 | 0
10 | 1 | 12 | 11 | 11 | 1 | 1
11 | 1 | 12 | 8 | 12 | 1 | 1
12 | 1 | 12 | 17 | 13 | 1 | 1
13 | 2 | 12 | 10 | 5 | 1 | 2
14 | 2 | 12 | 7 | 8 | 1 | 2
15 | 2 | 12 | 16 | 9 | 1 | 2
16 | 2 | 12 | 9 | 14 | 1 | 2
(16 rows)
```

Parameters

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below.
<b>root vid</b>	BIGINT	Identifier of the root vertex of the tree. <ul style="list-style-type: none"><li>When value is <math>\backslash(0\backslash)</math> then gets the spanning forest starting in aleatory nodes for each tree in the forest.</li></ul>
<b>root vids</b>	ARRAY [ <b>ANY-INTEGER</b> ]	Array of identifiers of the root vertices. <ul style="list-style-type: none"><li><math>\backslash(0\backslash)</math> values are ignored</li><li>For optimization purposes, any duplicated value is ignored.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERIC:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT, NUMERIC

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

DFS optional parameters

Parameter	Type	Default	Description
max_depth	BIGINT	$\backslash(9223372036854775807\backslash)$	Upper limit of the depth of the tree. <ul style="list-style-type: none"><li>When negative throws an error.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	<b>ANY-INTEGER</b>		Identifier of the edge.
source	<b>ANY-INTEGER</b>		Identifier of the first end point vertex of the edge.
target	<b>ANY-INTEGER</b>		Identifier of the second end point vertex of the edge.
cost	<b>ANY-NUMERICAL</b>		Weight of the edge (source, target)
reverse_cost	<b>ANY-NUMERICAL</b>	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

### Result columns

Parameter	Type	Description
-----------	------	-------------

Where:

SMALLINT, INTEGER, BIGINT

SMALLINT, INTEGER, BIGINT, REAL, FLOAT, NUMERIC

### Additional Examples

Same as Single vertex with edges in ascending order of id.

Example:

Same as Single vertex with edges in descending order of id.

The resulting traversal is different.

The left image shows the result with ascending order of ids and the right image shows with descending order of the edge identifiers.

**See Also**[1](#)

- [BFS - Category](#)
- [Sample Data](#)
- [Boost: Breadth First Search](#)
- [Wikipedia: Breadth First Search algorithm](#)

Indices and tables

- [Index](#)
- [Search Page](#)

`pgr_binaryBreadthFirstSearch` - **Experimental**

`pgr_binaryBreadthFirstSearch` — Returns the shortest path in a binary graph.

Any graph whose edge-weights belongs to the set  $\{0,X\}$ , where 'X' is any non-negative integer, is termed as a 'binary graph'.

☐ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.2.0
  - New experimental signature:
    - `pgr_binaryBreadthFirstSearch(Combinations)`
- Version 3.0.0
  - New experimental function.

**Description**

It is well-known that the shortest paths between a single source and all other vertices can be found using Breadth First Search in  $O(|E|)$  in an unweighted graph, i.e. the distance is the minimal number of edges that you need to traverse from the source to another vertex. We can interpret such a graph also as a weighted graph, where every edge has the weight  $\backslash(1\backslash)$ . If not all edges in graph have the same weight, that we need a more general algorithm, like Dijkstra's Algorithm which runs in  $\backslash(O(|E|\log|V|))\backslash$  time.

However if the weights are more constrained, we can use a faster algorithm. This algorithm, termed as 'Binary Breadth First Search' as well as '0-1 BFS', is a variation of the standard Breadth First Search problem to solve the SSSP (single-source shortest path) problem in  $\backslash(O(|E|)\backslash)$ , if the weights of each edge belongs to the set  $\{0,X\}$ , where 'X' is any non-negative real integer.

**The main Characteristics are:**

- Process is done only on 'binary graphs'. ('Binary Graph': Any graph whose edge-weights belongs to the set  $\{0,X\}$ , where 'X' is any non-negative real integer.)
- For optimization purposes, any duplicated value in the `start_vids` or `end_vids` are ignored.
- The returned values are ordered:
  - `start_vid` ascending
  - `end_vid` ascending
- Running time:  $\backslash(O(|\text{start\_vids}| * |E|)\backslash)$

Boost Graph Inside

**Signatures**

Summary

`pgr_binaryBreadthFirstSearch(Edges SQL, start vid, end vid, [directed])`  
`pgr_binaryBreadthFirstSearch(Edges SQL, start vid, end vids, [directed])`  
`pgr_binaryBreadthFirstSearch(Edges SQL, start vids, end vid, [directed])`  
`pgr_binaryBreadthFirstSearch(Edges SQL, start vids, end vids, [directed])`  
`pgr_binaryBreadthFirstSearch(Edges SQL, Combinations SQL, [directed])`  
Returns set of (seq, path\_seq, [start\_vid], [end\_vid], node, edge, cost, agg\_cost)  
OR EMPTY SET

**Note:** Using the [Sample Data](#) Network as all weights are same (i.e  $\backslash(1\backslash)$ )

**One to One**

pgr\_binaryBreadthFirstSearch([Edges SQL](#), start vid, end vid, [directed])  
Returns set of (seq, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertex \10\ on a **directed** graph

```
SELECT * FROM pgr_binaryBreadthFirstSearch(
  'SELECT id, source, target, cost, reverse_cost from edges',
  6, 10, true);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 4 | 1 | 0
2 | 2 | 7 | 8 | 1 | 1
3 | 3 | 11 | 9 | 1 | 2
4 | 4 | 16 | 16 | 1 | 3
5 | 5 | 15 | 3 | 1 | 4
6 | 6 | 10 | -1 | 0 | 5
(6 rows)
```

One to Many

pgr\_binaryBreadthFirstSearch([Edges SQL](#), start vid, end vids, [directed])  
Returns set of (seq, path\_seq, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \6\ to vertices \10, 17\ on a **directed** graph

```
SELECT * FROM pgr_binaryBreadthFirstSearch(
  'SELECT id, source, target, cost, reverse_cost from edges',
  6, ARRAY[10, 17]);
seq | path_seq | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 10 | 6 | 4 | 1 | 0
2 | 2 | 10 | 7 | 8 | 1 | 1
3 | 3 | 10 | 11 | 9 | 1 | 2
4 | 4 | 10 | 16 | 16 | 1 | 3
5 | 5 | 10 | 15 | 3 | 1 | 4
6 | 6 | 10 | 10 | -1 | 0 | 5
7 | 1 | 17 | 6 | 4 | 1 | 0
8 | 2 | 17 | 7 | 8 | 1 | 1
9 | 3 | 17 | 11 | 11 | 1 | 2
10 | 4 | 17 | 12 | 13 | 1 | 3
11 | 5 | 17 | 17 | -1 | 0 | 4
(11 rows)
```

Many to One

pgr\_binaryBreadthFirstSearch([Edges SQL](#), start vids, end vid, [directed])  
Returns set of (seq, path\_seq, start\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \6, 1\ to vertex \17\ on a **directed** graph

```
SELECT * FROM pgr_binaryBreadthFirstSearch(
  'SELECT id, source, target, cost, reverse_cost from edges',
  ARRAY[6, 1], 17);
seq | path_seq | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | 1 | 6 | 1 | 0
2 | 2 | 1 | 3 | 7 | 1 | 1
3 | 3 | 1 | 7 | 8 | 1 | 2
4 | 4 | 1 | 11 | 11 | 1 | 3
5 | 5 | 1 | 12 | 13 | 1 | 4
6 | 6 | 1 | 17 | -1 | 0 | 5
7 | 1 | 6 | 6 | 4 | 1 | 0
8 | 2 | 6 | 7 | 8 | 1 | 1
9 | 3 | 6 | 11 | 11 | 1 | 2
10 | 4 | 6 | 12 | 13 | 1 | 3
11 | 5 | 6 | 17 | -1 | 0 | 4
(11 rows)
```

Many to Many

pgr\_binaryBreadthFirstSearch([Edges SQL](#), start vids, end vids, [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \6, 1\ to vertices \10, 17\ on an **undirected** graph

```
SELECT * FROM pgr_binaryBreadthFirstSearch(
  'SELECT id, source, target, cost, reverse_cost from edges',
  ARRAY[6, 1], ARRAY[10, 17],
  directed => false);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 10 | 1 | 6 | 1 | 0
2 | 2 | 1 | 10 | 3 | 7 | 1 | 1
3 | 3 | 1 | 10 | 7 | 4 | 1 | 2
4 | 4 | 1 | 10 | 6 | 2 | 1 | 3
5 | 5 | 1 | 10 | 10 | -1 | 0 | 4
6 | 1 | 1 | 17 | 1 | 6 | 1 | 0
7 | 2 | 1 | 17 | 3 | 7 | 1 | 1
8 | 3 | 1 | 17 | 7 | 8 | 1 | 2
9 | 4 | 1 | 17 | 11 | 11 | 1 | 3
10 | 5 | 1 | 17 | 12 | 13 | 1 | 4
11 | 6 | 1 | 17 | 17 | -1 | 0 | 5
12 | 1 | 6 | 10 | 6 | 2 | 1 | 0
13 | 2 | 6 | 10 | 10 | -1 | 0 | 1
14 | 1 | 6 | 17 | 6 | 4 | 1 | 0
15 | 2 | 6 | 17 | 7 | 8 | 1 | 1
16 | 3 | 6 | 17 | 11 | 11 | 1 | 2
17 | 4 | 6 | 17 | 12 | 13 | 1 | 3
18 | 5 | 6 | 17 | 17 | -1 | 0 | 4
(18 rows)
```

Combinations

pgr\_binaryBreadthFirstSearch([Edges SQL](#), [Combinations SQL](#), [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET



Example:

Using a combinations table on an **undirected** graph

The combinations table:

```
SELECT source, target FROM combinations;
source | target
-----+-----
5 | 6
5 | 10
6 | 5
6 | 15
6 | 14
(5 rows)
```

The query:

```
SELECT * FROM pgr_binaryBreadthFirstSearch(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT source, target FROM combinations',
false);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 5 | 6 | 5 | 1 | 1 | 0
2 | 2 | 5 | 6 | 6 | -1 | 0 | 1
3 | 1 | 5 | 10 | 5 | 1 | 1 | 0
4 | 2 | 5 | 10 | 6 | 2 | 1 | 1
5 | 3 | 5 | 10 | 10 | -1 | 0 | 2
6 | 1 | 6 | 5 | 6 | 1 | 1 | 0
7 | 2 | 6 | 5 | 5 | -1 | 0 | 1
8 | 1 | 6 | 15 | 6 | 2 | 1 | 0
9 | 2 | 6 | 15 | 10 | 3 | 1 | 1
10 | 3 | 6 | 15 | 15 | -1 | 0 | 2
(10 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
source	<b>ANY-INTEGER</b>	Identifier of the departure vertex.
target	<b>ANY-INTEGER</b>	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Set of (seq, path\_id, path\_seq [, start\_vid] [, end\_vid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Path identifier. <ul style="list-style-type: none"> <li>Has value <b>1</b> for the first of a path fromstart_vid to end_vid.</li> </ul>
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex. Returned when multiple starting vetrices are in the query. <ul style="list-style-type: none"> <li><a href="#">Many to One</a></li> <li><a href="#">Many to Many</a></li> <li><a href="#">Combinations</a></li> </ul>
end_vid	BIGINT	Identifier of the ending vertex. Returned when multiple ending vertices are in the query. <ul style="list-style-type: none"> <li><a href="#">One to Many</a></li> <li><a href="#">Many to Many</a></li> <li><a href="#">Combinations</a></li> </ul>
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go fromnode to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Additional Examples

Example:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_binaryBreadthFirstSearch(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)');
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
3 | 1 | 6 | 10 | 6 | 4 | 1 | 0
4 | 2 | 6 | 10 | 7 | 8 | 1 | 1
5 | 3 | 6 | 10 | 11 | 9 | 1 | 2
6 | 4 | 6 | 10 | 16 | 16 | 1 | 3
7 | 5 | 6 | 10 | 15 | 3 | 1 | 4
8 | 6 | 6 | 10 | 10 | -1 | 0 | 5
9 | 1 | 12 | 10 | 12 | 13 | 1 | 0
10 | 2 | 12 | 10 | 17 | 15 | 1 | 1
11 | 3 | 12 | 10 | 16 | 16 | 1 | 2
12 | 4 | 12 | 10 | 15 | 3 | 1 | 3
13 | 5 | 12 | 10 | 10 | -1 | 0 | 4
(13 rows)
```

See Also

- [Sample Data](#)
- [Boost: Breadth First Search](#)
- [https://cp-algorithms.com/graph/01\\_bfs.html](https://cp-algorithms.com/graph/01_bfs.html)
- [https://en.wikipedia.org/wiki/Dijkstra%27s\\_algorithm#Specialized\\_variants](https://en.wikipedia.org/wiki/Dijkstra%27s_algorithm#Specialized_variants)

Indices and tables

- [Index](#)
- [Search Page](#)

See Also

Indices and tables

- [Index](#)

- [Search Page](#)

## Coloring - Family of functions¶

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.
- [pgr\\_sequentialVertexColoring - Proposed](#) - Vertex coloring algorithm using greedy approach.

☐ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting
- [pgr\\_bipartite - Experimental](#) - Bipartite graph algorithm using a DFS-based coloring approach.
- [pgr\\_edgeColoring - Experimental](#) - Edge Coloring algorithm using Vizing's theorem.

## pgr\_sequentialVertexColoring - Proposed¶

pgr\_sequentialVertexColoring — Returns the vertex coloring of an undirected graph, using greedy approach.

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Availability

- Version 3.3.0
  - Function promoted to proposed.
- Version 3.2.0
  - New experimental function.

## Description¶

Sequential vertex coloring algorithm is a graph coloring algorithm in which color identifiers are assigned to the vertices of a graph in a sequential manner, such that no edge connects two identically colored vertices.

The main Characteristics are:

- The implementation is applicable only for **undirected** graphs.
- Provides the color to be assigned to all the vertices present in the graph.
- Color identifiers values are in the Range\([1, |V|]\)
- The algorithm tries to assign the least possible color to every vertex.
- Efficient graph coloring is an NP-Hard problem, and therefore, this algorithm does not always produce optimal coloring. It follows a greedy strategy by iterating through all the vertices sequentially, and assigning the smallest possible color that is not used by its neighbors, to each vertex.
- The returned rows are ordered in ascending order of the vertex value.
- Sequential Vertex Coloring Running Time:  $\mathcal{O}(|V| \cdot (d + k))$ 
  - where  $|V|$  is the number of vertices,
  - $d$  is the maximum degree of the vertices in the graph,
  - $k$  is the number of colors used.

Boost Graph Inside

Signatures

pgr\_sequentialVertexColoring([Edges SQL](#))  
Returns set of (vertex\_id, color\_id)  
OR EMPTY SET

Example:

Graph coloring of pgRouting [Sample Data](#)

```
SELECT * FROM pgr_sequentialVertexColoring(
'SELECT id, source, target, cost, reverse_cost FROM edges
ORDER BY id'
);
```

vertex_id	color_id
1	1
2	1
3	2
4	2
5	1
6	2
7	1
8	2
9	1
10	1
11	2
12	1
13	1
14	2
15	2
16	1
17	2

(17 rows)

Parameters

Parameter Type	Description
----------------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (vertex\_id, color\_id)

Column Type	Description
-------------	-------------

Column Type Description

vertex\_id BIGINT Identifier of the vertex.

Identifier of the color of the vertex.

color\_id BIGINT

- The minimum value of color is 1.

See Also

- [Sample Data](#)
- [Boost: Sequential Vertex Coloring](#)
- [Wikipedia: Graph coloring](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_bipartite - Experimental

pgr\_bipartite — Disjoint sets of vertices such that no two vertices within the same set are adjacent.

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.2.0
  - New experimental function.

Description

A bipartite graph is a graph with two sets of vertices which are connected to each other, but not within themselves. A bipartite graph is possible if the graph coloring is possible using two colors such that vertices in a set are colored with the same color.

The main Characteristics are:

- The algorithm works in undirected graph only.
- The returned values are not ordered.
- The algorithm checks graph is bipartite or not. If it is bipartite then it returns the node along with two colors 0 and 1 which represents two different sets.
- If graph is not bipartite then algorithm returns empty set.
- Running time:  $O(V + E)$

Boost Graph Inside

Signatures

pgr\_bipartite(Edges SQL)

Returns set of (vertex\_id, color\_id)

OR EMPTY SET

Example:

When the graph is bipartite

```
SELECT * FROM pgr_bipartite(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$
) ORDER BY vertex_id;
vertex_id | color_id
```

```
-----+-----
1 | 0
2 | 0
3 | 1
```

4		1
5		0
6		1
7		0
8		1
9		0
10		0
11		1
12		0
13		0
14		1
15		1
16		0
17		1

(17 rows)

Parameters

Parameter	Type	Description
-----------	------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (vertex\_id, color\_id)

Column	Type	Description
--------	------	-------------

vertex_id	BIGINT	Identifier of the vertex.
		Identifier of the color of the vertex.
color_id	BIGINT	<ul style="list-style-type: none"><li>The minimum value of color is 1.</li></ul>

Additional Example

Example:

The odd length cyclic graph can not be bipartite.

The edge  $(5 \rightarrow 1)$  will make subgraph with vertices  $\{1, 3, 7, 6, 5\}$  an odd length cyclic graph, as the cycle has 5 vertices.

```
INSERT INTO edges (source, target, cost, reverse_cost) VALUES
(5, 1, 1, 1);
INSERT 0 1
```

Edges in blue represent odd length cycle subgraph.

[\\_images/bipartite.png](#)

```
SELECT * FROM pgr_bipartite(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$
```

```
);
vertex_id | color_id
-----+-----
(0 rows)
```

See Also

- [Boost: is\\_bipartite](#)
- [Wikipedia: bipartite graph](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_edgeColoring - Experimental

pgr\_edgeColoring — Returns the edge coloring of undirected and loop-free graphs

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.3.0
  - New experimental function.

Description

Edge Coloring is an algorithm used for coloring of the edges for the vertices in the graph. It is an assignment of colors to the edges of the graph so that no two adjacent edges have the same color.

The main Characteristics are:

- The implementation is for **undirected** and **loop-free** graphs
  - loop free:
    - no self-loops and no parallel edges.
- Provides the color to be assigned to all the edges present in the graph.
- At most  $(\Delta + 1)$  colors are used, where  $\Delta$  is the degree of the graph.
  - This is optimal for some graphs, and by Vizing's theorem it uses at most one color more than the optimal for all others.
  - When the graph is bipartite
    - the chromatic number  $\chi(G)$  (minimum number of colors needed for proper edge coloring of graph) is equal to the degree  $\Delta$  of the graph,  $(\chi(G) = \Delta)$
- The algorithm tries to assign the least possible color to every edge.
  - Does not always produce optimal coloring.
- The returned rows are ordered in ascending order of the edge identifier.
- Efficient graph coloring is an NP-Hard problem, and therefore:
  - In this implemtentation the running time:  $O(|E| * |V|)$ 
    - where  $|E|$  is the number of edges in the graph,
    - $|V|$  is the number of vertices in the graph.

Boost Graph Inside

Signatures

pgr\_edgeColoring([Edges SQL](#))  
Returns set of (edge\_id, color\_id)  
OR EMPTY SET

Example:

Graph coloring of pgRouting[Sample Data](#)

```
SELECT * FROM pgr_edgeColoring(
'SELECT id, source, target, cost, reverse_cost FROM edges
ORDER BY id'
);
edge_id | color_id
-----+-----
1 | 3
2 | 2
3 | 3
4 | 4
5 | 4
6 | 1
7 | 2
8 | 1
9 | 2
10 | 5
11 | 5
12 | 3
13 | 2
14 | 1
15 | 3
16 | 1
17 | 1
18 | 1
(18 rows)
```

Parameters<sup>1</sup>

Parameter	Type	Description
-----------	------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below.

Inner Queries<sup>1</sup>

Edges SQL<sup>1</sup>

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns<sup>1</sup>

Returns set of (edge\_id, color\_id)

Column	Type	Description
edge_id	BIGINT	Identifier of the edge.
color_id	BIGINT	Identifier of the color of the edge. <ul style="list-style-type: none"><li>The minimum value of color is 1.</li></ul>

See Also<sup>1</sup>

- [Sample Data](#)
- [Boost: Edge Coloring](#)
- [Wikipedia: Graph coloring](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Result columns<sup>1</sup>

Returns set of (vertex\_id, color\_id)

Column	Type	Description
vertex_id	BIGINT	Identifier of the vertex.



Column	Type	Description
		Identifier of the color of the vertex.
color_id	BIGINT	<ul style="list-style-type: none"> <li>The minimum value of color is 1.</li> </ul>

Returns set of (edge\_id, color\_id)

Column	Type	Description
		Identifier of the edge.
		Identifier of the color of the edge.
color_id	BIGINT	<ul style="list-style-type: none"> <li>The minimum value of color is 1.</li> </ul>

See Also

- [Boost](#):

Indices and tables

- [Index](#)
- [Search Page](#)

categories

[Cost - Category](#)

- [pgr\\_withPointsCost - Proposed](#)

[Cost Matrix - Category](#)

- [pgr\\_withPointsCostMatrix - proposed](#)

[Driving Distance - Category](#)

- [pgr\\_withPointsDD - Proposed](#) - Driving Distance based on pgr\_withPoints

[K shortest paths - Category](#)

- [pgr\\_withPointsKSP - Proposed](#) - Yen's algorithm based on pgr\_withPoints

[Via - Category](#)

- [pgr\\_dijkstraVia - Proposed](#)
- [pgr\\_withPointsVia - Proposed](#)
- [pgr\\_trspVia - Proposed](#)
- [pgr\\_trspVia\\_withPoints - Proposed](#)

[withPoints - Category](#)

- [withPoints - Family of functions](#) - Functions based on Dijkstra algorithm.
- From the [TRSP - Family of functions](#):
  - [pgr\\_trsp\\_withPoints - Proposed](#) - Vertex/Point routing with restrictions.
  - [pgr\\_trspVia\\_withPoints - Proposed](#) - Via Vertex/point routing with restrictions.

withPoints - Family of functions

When points are also given as input:

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.
- [pgr\\_withPoints - Proposed](#) - Route from/to points anywhere on the graph.
- [pgr\\_withPointsCost - Proposed](#) - Costs of the shortest paths.
- [pgr\\_withPointsCostMatrix - proposed](#) - Costs of the shortest paths.
- [pgr\\_withPointsKSP - Proposed](#) - K shortest paths.
- [pgr\\_withPointsDD - Proposed](#) - Driving distance.
- [pgr\\_withPointsVia - Proposed](#) - Via routing

pgr\_withPoints - Returns the shortest path in a graph with additional temporary vertices.

❏ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Availability

- Version 3.2.0
  - New proposed signature:
    - pgr\_withPoints(Combinations)
- Version 2.2.0
  - New proposed function.

Description🔗

Modify the graph to include points defined by points\_sql. Using Dijkstra algorithm, find the shortest path

The main characteristics are:

- Process is done only on edges with positive costs.
- Vertices of the graph are:
  - **positive** when it belongs to the edges\_sql
  - **negative** when it belongs to the points\_sql
- Values are returned when there is a path.
  - When the starting vertex and ending vertex are the same, there is no path. - The agg\_cost the non included values (v, v) is 0
  - When the starting vertex and ending vertex are the different and there is no path: - The agg\_cost the non included values (u, v) is ∞
- For optimization purposes, any duplicated value in the start\_vids or end\_vids are ignored.
- The returned values are ordered: - start\_vid ascending - end\_vid ascending
- Running time:  $\backslash(O((\text{start\_vids})\times(V\log V + E)))\backslash$

Boost Graph Inside

Signatures🔗

Summary

pgr\_withPoints([Edges SQL](#), [Points SQL](#), start vid, end vid, [options])  
pgr\_withPoints([Edges SQL](#), [Points SQL](#), start vid, end vids, [options])  
pgr\_withPoints([Edges SQL](#), [Points SQL](#), start vids, end vid, [options])  
pgr\_withPoints([Edges SQL](#), [Points SQL](#), start vids, end vids, [options])  
pgr\_withPoints([Edges SQL](#), [Points SQL](#), [Combinations SQL](#), [options])  
**options:** [directed, driving\_side, details])  
Returns set of (seq, path\_seq, [start\_pid], [end\_pid], node, edge, cost, agg\_cost)  
OR EMPTY SET

One to One🔗

pgr\_withPoints([Edges SQL](#), [Points SQL](#), start vid, end vid, [options])  
**options:** [directed, driving\_side, details])  
Returns set of (seq, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From point \1\ to vertex \10\ with details

```
SELECT * FROM pgr_withPoints(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
-1, 10,
details => true);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -1 | 1 | 0.6 | 0
2 | 2 | 6 | 4 | 0.7 | 0.6
3 | 3 | 3 | -6 | 4 | 0.3 | 1.3
4 | 4 | 4 | 7 | 8 | 1 | 1.6
5 | 5 | 5 | 11 | 9 | 1 | 2.6
6 | 6 | 6 | 16 | 16 | 1 | 3.6
7 | 7 | 7 | 15 | 3 | 1 | 4.6
8 | 8 | 8 | 10 | -1 | 0 | 5.6
(8 rows)
```

One to Many🔗

pgr\_withPoints([Edges SQL](#), [Points SQL](#), start vid, end vids, [options])  
**options:** [directed, driving\_side, details])  
Returns set of (seq, path\_seq, end\_pid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From point \1\ to point \3\ and vertex \7\ on an undirected graph

```
SELECT * FROM pgr_withPoints(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  'SELECT pid, edge_id, fraction, side from pointsOfInterest',
  -1, ARRAY[-3, 7],
  directed => false);
seq | path_seq | end_pid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | -3 | -1 | 1 | 0.6 | 0
2 | 2 | -3 | 6 | 4 | 1 | 0.6
3 | 3 | -3 | 7 | 10 | 1 | 1.6
4 | 4 | -3 | 8 | 12 | 0.6 | 2.6
5 | 5 | -3 | -3 | -1 | 0 | 3.2
6 | 1 | 7 | -1 | 1 | 0.6 | 0
7 | 2 | 7 | 6 | 4 | 1 | 0.6
8 | 3 | 7 | -1 | 0 | 1.6
(8 rows)
```

Many to One

pgr\_withPoints([Edges SQL](#), [Points SQL](#), start vids, end vid, [options])  
options: [directed, driving\_side, details])  
Returns set of (seq, path\_seq, start\_pid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:  
From point \1\ and vertex \6\ to point \3\

```
SELECT * FROM pgr_withPoints(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  'SELECT pid, edge_id, fraction, side from pointsOfInterest',
  ARRAY[-1, 6], -3);
seq | path_seq | start_pid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -1 | -1 | 1 | 0.6 | 0
2 | 2 | 2 | -1 | 6 | 4 | 1 | 0.6
3 | 3 | 3 | -1 | 7 | 10 | 1 | 1.6
4 | 4 | 4 | -1 | 8 | 12 | 0.6 | 2.6
5 | 5 | 5 | -1 | -3 | -1 | 0 | 3.2
6 | 1 | 6 | 6 | 4 | 1 | 0
7 | 2 | 6 | 7 | 10 | 1 | 1
8 | 3 | 6 | 8 | 12 | 0.6 | 2
9 | 4 | 6 | -3 | -1 | 0 | 2.6
(9 rows)
```

Many to Many

pgr\_withPoints([Edges SQL](#), [Points SQL](#), start vids, end vids, [options])  
options: [directed, driving\_side, details])  
Returns set of (seq, path\_seq, start\_pid, end\_pid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:  
From point \1\ and vertex \6\ to point \3\ and vertex \1\

```
SELECT * FROM pgr_withPoints(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  'SELECT pid, edge_id, fraction, side from pointsOfInterest',
  ARRAY[-1, 6], ARRAY[-3, 1]);
seq | path_seq | start_pid | end_pid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -1 | -3 | -1 | 1 | 0.6 | 0
2 | 2 | 2 | -1 | -3 | 6 | 4 | 1 | 0.6
3 | 3 | 3 | -1 | -3 | 7 | 10 | 1 | 1.6
4 | 4 | 4 | -1 | -3 | 8 | 12 | 0.6 | 2.6
5 | 5 | 5 | -1 | -3 | -3 | -1 | 0 | 3.2
6 | 1 | 1 | -1 | 1 | -1 | 1 | 0.6 | 0
7 | 2 | 2 | -1 | 1 | 6 | 4 | 1 | 0.6
8 | 3 | 3 | -1 | 1 | 7 | 7 | 1 | 1.6
9 | 4 | 4 | -1 | 1 | 3 | 6 | 1 | 2.6
10 | 5 | 5 | -1 | 1 | 1 | -1 | 0 | 3.6
11 | 1 | 6 | -3 | 6 | 4 | 1 | 0
12 | 2 | 6 | -3 | 7 | 10 | 1 | 1
13 | 3 | 6 | -3 | 8 | 12 | 0.6 | 2
14 | 4 | 6 | -3 | -3 | -1 | 0 | 2.6
15 | 1 | 6 | 1 | 6 | 4 | 1 | 0
16 | 2 | 6 | 1 | 7 | 7 | 1 | 1
17 | 3 | 6 | 1 | 3 | 6 | 1 | 2
18 | 4 | 6 | 1 | 1 | -1 | 0 | 3
(18 rows)
```

Combinations

pgr\_withPoints([Edges SQL](#), [Points SQL](#), [Combinations SQL](#), [options])  
options: [directed, driving\_side, details])  
Returns set of (seq, path\_seq, start\_pid, end\_pid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:  
Two combinations  
From point \1\ to vertex \10\, and from vertex \6\ to point \3\ with **right** side driving.

```
SELECT * FROM pgr_withPoints(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  'SELECT pid, edge_id, fraction, side from pointsOfInterest',
  'SELECT * FROM (VALUES (-1, 10), (6, -3)) AS combinations(source, target)',
  driving_side => 'r', details => true);
seq | path_seq | start_pid | end_pid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -1 | 10 | -1 | 1 | 0.4 | 0
2 | 2 | 2 | -1 | 10 | 5 | 1 | 1 | 0.4
3 | 3 | 3 | -1 | 10 | 6 | 4 | 0.7 | 1.4
4 | 4 | 4 | -1 | 10 | -6 | 4 | 0.3 | 2.1
5 | 5 | 5 | -1 | 10 | 7 | 8 | 1 | 2.4
6 | 6 | 6 | -1 | 10 | 11 | 9 | 1 | 3.4
7 | 7 | 7 | -1 | 10 | 16 | 16 | 1 | 4.4
8 | 8 | 8 | -1 | 10 | 15 | 3 | 1 | 5.4
9 | 9 | 9 | -1 | 10 | 10 | -1 | 0 | 6.4
10 | 1 | 6 | -3 | 6 | 4 | 0.7 | 0
11 | 2 | 6 | -3 | -6 | 4 | 0.3 | 0.7
12 | 3 | 6 | -3 | 7 | 10 | 1 | 1
13 | 4 | 6 | -3 | 8 | 12 | 0.6 | 2
14 | 5 | 6 | -3 | -3 | -1 | 0 | 2.6
(14 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Points SQL</a>	TEXT	<a href="#">Points SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
<b>start vid</b>	BIGINT	Identifier of the starting vertex of the path. Negative value is for point's identifier.
<b>start vids</b>	ARRAY[BIGINT]	Array of identifiers of starting vertices. Negative values are for point's identifiers.
<b>end vid</b>	BIGINT	Identifier of the ending vertex of the path. Negative value is for point's identifier.
<b>end vids</b>	ARRAY[BIGINT]	Array of identifiers of ending vertices. Negative values are for point's identifiers.

Optional parameters<sup>1</sup>

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true the graph is considered <i>Directed</i></li> <li>When false the graph is considered as <i>Undirected</i>.</li> </ul>

With points optional parameters<sup>1</sup>

Parameter	Type	Default	Description
driving_side	CHAR	b	Value in [r, l, b] indicating if the driving side is: <ul style="list-style-type: none"> <li>r for right driving side.</li> <li>l for left driving side.</li> <li>b for both.</li> </ul>
details	BOOLEAN	false	<ul style="list-style-type: none"> <li>When true the results will include the points that are in the path.</li> <li>When false the results will not include the points that are in the path.</li> </ul>

Inner Queries<sup>1</sup>

Edges SQL<sup>1</sup>

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Points SQL<sup>1</sup>

Parameter	Type	Default	Description
			Identifier of the point.
pid	ANY-INTEGER	value	<ul style="list-style-type: none"> <li>Use with positive value, as internally will be converted to negative value</li> <li>If column is present, it can not be NULL.</li> <li>If column is not present, a sequential negative <b>value</b> will be given automatically.</li> </ul>

Parameter	Type	Default	Description
edge_id	<b>ANY-INTEGER</b>		Identifier of the "closest" edge to the point.
fraction	<b>ANY-NUMERICAL</b>		Value in <0,1> that indicates the relative position from the first end point of the edge.
side	CHAR	b	Value in [b, r, l, NULL] indicating if the point is: <ul style="list-style-type: none"> <li>• In the right r,</li> <li>• In the left l,</li> <li>• In both sides b, NULL</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations [SQL ¶](#)

Parameter	Type	Description
source	<b>ANY-INTEGER</b>	Identifier of the departure vertex.
target	<b>ANY-INTEGER</b>	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns [¶](#)

Returns set of (seq, path\_seq [, start\_pid] [, end\_pid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_seq	INTEGER	Relative position in the path. <ul style="list-style-type: none"> <li>• 1 For the first row of the path.</li> </ul>
start_pid	BIGINT	Identifier of a starting vertex/point of the path. <ul style="list-style-type: none"> <li>• When positive is the identifier of the starting vertex.</li> <li>• When negative is the identifier of the starting point.</li> <li>• Returned on <a href="#">Many to One</a> and <a href="#">Many to Many</a></li> </ul>
end_pid	BIGINT	Identifier of an ending vertex/point of the path. <ul style="list-style-type: none"> <li>• When positive is the identifier of the ending vertex.</li> <li>• When negative is the identifier of the ending point.</li> <li>• Returned on <a href="#">One to Many</a> and <a href="#">Many to Many</a></li> </ul>
node	BIGINT	Identifier of the node in the path from start_pid to end_pid. <ul style="list-style-type: none"> <li>• When positive is the identifier of the a vertex.</li> <li>• When negative is the identifier of the a point.</li> </ul>
edge	BIGINT	Identifier of the edge used to go from node to the next node in the path sequence. <ul style="list-style-type: none"> <li>• -1 for the last row of the path.</li> </ul>
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence. <ul style="list-style-type: none"> <li>• 0 For the first row of the path.</li> </ul>
agg_cost	FLOAT	Aggregate cost from start_vid to node. <ul style="list-style-type: none"> <li>• 0 For the first row of the path.</li> </ul>

Additional Examples [¶](#)

- [Use pgr\\_findCloseEdges](#) in the Points SQL.
- [Usage variations](#)
  - [Passes in front or visits with right side driving.](#)
  - [Passes in front or visits with left side driving.](#)

Use [pgr\\_findCloseEdges](#) in the [Points SQL](#).

Find the routes from vertex \1\ to the two closest locations on the graph of point(2.9, 1.8).

```
SELECT * FROM pgr_withPoints(
  $e$ SELECT * FROM edges $e$,
  $p$ SELECT edge_id, round(fraction::numeric, 2) AS fraction, side
    FROM pgr_findCloseEdges(
      $$SELECT id, geom FROM edges $$,
      (SELECT ST_POINT(2.9, 1.8)),
      0.5, cap => 2)
  $p$,
  1, ARRAY[-1, -2]);
seq | path_seq | end_pid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | -2 | 1 | 6 | 1 | 0
2 | 2 | -2 | 3 | 7 | 1 | 1
3 | 3 | -2 | 7 | 8 | 0.9 | 2
4 | 4 | -2 | -2 | -1 | 0 | 2.9
5 | 1 | -1 | 1 | 6 | 1 | 0
6 | 2 | -1 | 3 | 7 | 1 | 1
7 | 3 | -1 | 7 | 8 | 1 | 2
8 | 4 | -1 | 11 | 9 | 1 | 3
9 | 5 | -1 | 16 | 16 | 1 | 4
10 | 6 | -1 | 15 | 3 | 1 | 5
11 | 7 | -1 | 10 | 5 | 0.8 | 6
12 | 8 | -1 | -1 | -1 | 0 | 6.8
(12 rows)
```

- Point \-1\ corresponds to the closest edge from point(2.9, 1.8).
- Point \-2\ corresponds to the next close edge from point(2.9, 1.8).

[Usage variations](#)

All the examples are about traveling from point \1\ and vertex \5\ to points \2, 3, 6\ and vertices \10, 11\

```
SELECT *
FROM pgr_withPoints(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  'SELECT pid, edge_id, fraction, side from pointsOfInterest',
  ARRAY[5, -1], ARRAY[-2, -3, -6, 10, 11],
  driving_side => 'r', details => true);
seq | path_seq | start_pid | end_pid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | -1 | -6 | -1 | 1 | 0.4 | 0
2 | 2 | -1 | -6 | 5 | 1 | 1 | 0.4
3 | 3 | -1 | -6 | 6 | 4 | 0.7 | 1.4
4 | 4 | -1 | -6 | -1 | 0 | 2.1
5 | 1 | -1 | -3 | -1 | 1 | 0.4 | 0
6 | 2 | -1 | -3 | 5 | 1 | 1 | 0.4
7 | 3 | -1 | -3 | 6 | 4 | 0.7 | 1.4
8 | 4 | -1 | -3 | -6 | 4 | 0.3 | 2.1
9 | 5 | -1 | -3 | 7 | 10 | 1 | 2.4
10 | 6 | -1 | -3 | 8 | 12 | 0.6 | 3.4
11 | 7 | -1 | -3 | -3 | -1 | 0 | 4
12 | 1 | -1 | -2 | -1 | 1 | 0.4 | 0
13 | 2 | -1 | -2 | 5 | 1 | 1 | 0.4
14 | 3 | -1 | -2 | 6 | 4 | 0.7 | 1.4
15 | 4 | -1 | -2 | -6 | 4 | 0.3 | 2.1
16 | 5 | -1 | -2 | 7 | 8 | 1 | 2.4
17 | 6 | -1 | -2 | 11 | 9 | 1 | 3.4
18 | 7 | -1 | -2 | 16 | 15 | 0.4 | 4.4
19 | 8 | -1 | -2 | -2 | -1 | 0 | 4.8
20 | 1 | -1 | 10 | -1 | 1 | 0.4 | 0
21 | 2 | -1 | 10 | 5 | 1 | 1 | 0.4
22 | 3 | -1 | 10 | 6 | 4 | 0.7 | 1.4
23 | 4 | -1 | 10 | -6 | 4 | 0.3 | 2.1
24 | 5 | -1 | 10 | 7 | 8 | 1 | 2.4
25 | 6 | -1 | 10 | 11 | 9 | 1 | 3.4
26 | 7 | -1 | 10 | 16 | 16 | 1 | 4.4
27 | 8 | -1 | 10 | 15 | 3 | 1 | 5.4
28 | 9 | -1 | 10 | 10 | -1 | 0 | 6.4
29 | 1 | -1 | 11 | -1 | 1 | 0.4 | 0
30 | 2 | -1 | 11 | 5 | 1 | 1 | 0.4
31 | 3 | -1 | 11 | 6 | 4 | 0.7 | 1.4
32 | 4 | -1 | 11 | -6 | 4 | 0.3 | 2.1
33 | 5 | -1 | 11 | 7 | 8 | 1 | 2.4
34 | 6 | -1 | 11 | 11 | -1 | 0 | 3.4
35 | 1 | 5 | -6 | 5 | 1 | 1 | 0
36 | 2 | 5 | -6 | 6 | 4 | 0.7 | 1
37 | 3 | 5 | -6 | -6 | -1 | 0 | 1.7
38 | 1 | 5 | -3 | 5 | 1 | 1 | 0
39 | 2 | 5 | -3 | 6 | 4 | 0.7 | 1
40 | 3 | 5 | -3 | -6 | 4 | 0.3 | 1.7
41 | 4 | 5 | -3 | 7 | 10 | 1 | 2
42 | 5 | 5 | -3 | 8 | 12 | 0.6 | 3
43 | 6 | 5 | -3 | -3 | -1 | 0 | 3.6
44 | 1 | 5 | -2 | 5 | 1 | 1 | 0
45 | 2 | 5 | -2 | 6 | 4 | 0.7 | 1
46 | 3 | 5 | -2 | -6 | 4 | 0.3 | 1.7
47 | 4 | 5 | -2 | 7 | 8 | 1 | 2
48 | 5 | 5 | -2 | 11 | 9 | 1 | 3
49 | 6 | 5 | -2 | 16 | 15 | 0.4 | 4
50 | 7 | 5 | -2 | -2 | -1 | 0 | 4.4
51 | 1 | 5 | 10 | 5 | 1 | 1 | 0
52 | 2 | 5 | 10 | 6 | 4 | 0.7 | 1
53 | 3 | 5 | 10 | -6 | 4 | 0.3 | 1.7
54 | 4 | 5 | 10 | 7 | 8 | 1 | 2
55 | 5 | 5 | 10 | 11 | 9 | 1 | 3
56 | 6 | 5 | 10 | 16 | 16 | 1 | 4
57 | 7 | 5 | 10 | 15 | 3 | 1 | 5
58 | 8 | 5 | 10 | 10 | -1 | 0 | 6
59 | 1 | 5 | 11 | 5 | 1 | 1 | 0
60 | 2 | 5 | 11 | 6 | 4 | 0.7 | 1
61 | 3 | 5 | 11 | -6 | 4 | 0.3 | 1.7
62 | 4 | 5 | 11 | 7 | 8 | 1 | 2
63 | 5 | 5 | 11 | 11 | -1 | 0 | 3
(63 rows)
```

[Passes in front or visits with right side driving](#)

For point \6\ and vertex \11\.

```
SELECT (start_pid || ' -> ' || end_pid || ' at ' || path_seq || 'th step')::TEXT AS path_at,
CASE WHEN edge = -1 THEN 'visits'
  ELSE 'passes in front of'
END AS status,
CASE WHEN node < 0 THEN 'Point'
  ELSE 'Vertex'
END AS is_a,
abs(node) AS id
FROM pgr_withPoints(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  'SELECT pid, edge_id, fraction, side from pointsOfInterest',
```

```
ARRAY[5, -1], ARRAY[-2, -3, -6, 10, 11],
driving_side => 'r', details => true)
WHERE node IN (-6, 11);
path_at | status | is_a | id
-----+-----+-----+-----
-1 -> -6 at 4th step | visits | Point | 6
-1 -> -3 at 4th step | passes in front of | Point | 6
-1 -> -2 at 4th step | passes in front of | Point | 6
-1 -> -2 at 6th step | passes in front of | Vertex | 11
-1 -> 10 at 4th step | passes in front of | Point | 6
-1 -> 10 at 6th step | passes in front of | Vertex | 11
-1 -> 11 at 4th step | passes in front of | Point | 6
-1 -> 11 at 6th step | visits | Vertex | 11
5 -> -6 at 3th step | visits | Point | 6
5 -> -3 at 3th step | passes in front of | Point | 6
5 -> -2 at 3th step | passes in front of | Point | 6
5 -> -2 at 5th step | passes in front of | Vertex | 11
5 -> 10 at 3th step | passes in front of | Point | 6
5 -> 10 at 5th step | passes in front of | Vertex | 11
5 -> 11 at 3th step | passes in front of | Point | 6
5 -> 11 at 5th step | visits | Vertex | 11
(16 rows)
```

Passes in front or visits with left side driving.¶

For point \(-6\)\) and vertex \(\{11\}\).

```
SELECT (start_pid || ' => ' || end_pid || ' at ' || path_seq || 'th step')::TEXT AS path_at,
CASE WHEN edge = -1 THEN 'visits'
ELSE 'passes in front of'
END AS status,
CASE WHEN node < 0 THEN 'Point'
ELSE 'Vertex'
END AS is_a,
abs(node) as id
FROM pgr_withPoints(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[5, -1], ARRAY[-2, -3, -6, 10, 11],
driving_side => 'l', details => true)
WHERE node IN (-6, 11);
path_at | status | is_a | id
-----+-----+-----+-----
-1 => -6 at 3th step | visits | Point | 6
-1 => -3 at 3th step | passes in front of | Point | 6
-1 => -2 at 3th step | passes in front of | Point | 6
-1 => -2 at 5th step | passes in front of | Vertex | 11
-1 => 10 at 3th step | passes in front of | Point | 6
-1 => 10 at 5th step | passes in front of | Vertex | 11
-1 => 11 at 3th step | passes in front of | Point | 6
-1 => 11 at 5th step | visits | Vertex | 11
5 => -6 at 4th step | visits | Point | 6
5 => -3 at 4th step | passes in front of | Point | 6
5 => -2 at 4th step | passes in front of | Point | 6
5 => -2 at 6th step | passes in front of | Vertex | 11
5 => 10 at 4th step | passes in front of | Point | 6
5 => 10 at 6th step | passes in front of | Vertex | 11
5 => 11 at 4th step | passes in front of | Point | 6
5 => 11 at 6th step | visits | Vertex | 11
(16 rows)
```

See Also¶

- [withPoints - Family of functions](#)
- [withPoints - Category](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_withPointsCost - Proposed¶

pgr\_withPointsCost - Calculates the shortest path and returns only the aggregate cost of the shortest path found, for the combination of points given.

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Availability

- Version 3.2.0
  - New proposed signature:
    - pgr\_withPointsCost(Combinations)
- Version 2.2.0
  - New proposed function.

Description¶

Modify the graph to include points defined by points\_sql. Using Dijkstra algorithm, return only the aggregate cost of the shortest path found.

The main characteristics are:

- It does not return a path.
- Returns the sum of the costs of the shortest path for pair combination of vertices in the modified graph.
- Vertices of the graph are:
  - **positive** when it belongs to the edges\_sql
  - **negative** when it belongs to the points\_sql
- Process is done only on edges with positive costs.
- Values are returned when there is a path.
  - The returned values are in the form of a set of (start\_vid, end\_vid, agg\_cost).
  - When the starting vertex and ending vertex are the same, there is no path.
    - The agg\_cost in the non included values (v, v) is 0
  - When the starting vertex and ending vertex are the different and there is no path.
    - The agg\_cost in the non included values (u, v) is \(\infty\)
- If the values returned are stored in a table, the unique index would be the pair (start\_vid, end\_vid).
- For **undirected** graphs, the results are **symmetric**.
  - The agg\_cost of (u, v) is the same as for (v, u).
- For optimization purposes, any duplicated value in the start\_vids or end\_vids is ignored.
- The returned values are ordered:
  - start\_vid ascending
  - end\_vid ascending
- Running time:  $\mathcal{O}(|start\_vids| \times (V \log V + E))$

Boost Graph Inside

Signatures

Summary

pgr\_withPointsCost([Edges SQL](#), 'Points SQL'\_, **start vid, end vid, [options]**)  
pgr\_withPointsCost([Edges SQL](#), 'Points SQL'\_, **start vid, end vids, [options]**)  
pgr\_withPointsCost([Edges SQL](#), 'Points SQL'\_, **start vids, end vid, [options]**)  
pgr\_withPointsCost([Edges SQL](#), 'Points SQL'\_, **start vids, end vids, [options]**)  
pgr\_withPointsCost([Edges SQL](#), 'Points SQL'\_, [Combinations SQL](#), **[options]**)  
**options:** [directed, driving\_side]  
Returns set of (start\_pid, end\_pid, agg\_cost)  
OR EMPTY SET

Note

There is no **details** flag, unlike the other members of the withPoints family of functions.

One to One

pgr\_withPointsCost([Edges SQL](#), 'Points SQL'\_, **start vid, end vid, [options]**)  
**options:** [directed, driving\_side]  
Returns set of (start\_pid, end\_pid, agg\_cost)  
OR EMPTY SET

Example:

From point \(\backslash1\backslash\) to vertex \(\backslash10\backslash\) with defaults

```
SELECT * FROM pgr_withPointsCost(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
-1, 10);
start_pid | end_pid | agg_cost
-----+-----+-----
-1 | 10 | 5.6
(1 row)
```

One to Many

pgr\_withPointsCost([Edges SQL](#), [Points SQL](#), **start vid, end vids, [options]**)  
**options:** [directed, driving\_side]  
Returns set of (start\_pid, end\_pid, agg\_cost)  
OR EMPTY SET

Example:

From point \(\backslash1\backslash\) to point \(\backslash3\backslash\) and vertex \(\backslash7\backslash\) on an undirected graph

```
SELECT * FROM pgr_withPointsCost(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
-1, ARRAY[-3, 7],
directed => false);
start_pid | end_pid | agg_cost
-----+-----+-----
-1 | -3 | 3.2
-1 | 7 | 1.6
(2 rows)
```

Many to One

pgr\_withPointsCost([Edges SQL](#), [Points SQL](#), **start vids, end vid, [options]**)  
**options:** [directed, driving\_side]  
Returns set of (start\_pid, end\_pid, agg\_cost)  
OR EMPTY SET

Example:

From point \(\backslash1\backslash\) and vertex \(\backslash6\backslash\) to point \(\backslash3\backslash\)

```
SELECT * FROM pgr_withPointsCost(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
```



```
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[-1, 6], -3);
start_pid | end_pid | agg_cost
-----+-----+-----
-1 | -3 | 3.2
6 | -3 | 2.6
(2 rows)
```

Many to Many

pgr\_withPointsCost([Edges SQL](#), [Points SQL](#), start vids, end vids, [options])  
**options:** [directed, driving\_side]  
Returns set of (start\_pid, end\_pid, agg\_cost)  
OR EMPTY SET

Example:

From point \1 and vertex \6 to point \3 and vertex \1

```
SELECT * FROM pgr_withPointsCost(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[-1, 6], ARRAY[-3, 1]);
start_pid | end_pid | agg_cost
-----+-----+-----
-1 | -3 | 3.2
-1 | 1 | 3.6
6 | -3 | 2.6
6 | 1 | 3
(4 rows)
```

Combinations

pgr\_withPointsCost([Edges SQL](#), [Points SQL](#), [Combinations SQL](#), [options])  
**options:** [directed, driving\_side]  
Returns set of (start\_pid, end\_pid, agg\_cost)  
OR EMPTY SET

Example:

Two combinations

From point \1 to vertex \10, and from vertex \6 to point \3 with **right** side driving.

```
SELECT * FROM pgr_withPointsCost(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
'SELECT * FROM (VALUES (-1, 10), (6, -3)) AS combinations(source, target)',
driving_side => 'r');
start_pid | end_pid | agg_cost
-----+-----+-----
-1 | 10 | 6.4
6 | -3 | 2.6
(2 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Points SQL</a>	TEXT	<a href="#">Points SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path. Negative value is for point's identifier.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices. Negative values are for point's identifiers.
end vid	BIGINT	Identifier of the ending vertex of the path. Negative value is for point's identifier.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices. Negative values are for point's identifiers.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

With points optional parameters

Parameter	Type	Default	Description
driving_side	CHAR	b	Value in [r, l, b] indicating if the driving side is: <ul style="list-style-type: none"><li>r for right driving side.</li><li>l for left driving side.</li><li>b for both.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGERS		Identifier of the edge.
source	ANY-INTEGERS		Identifier of the first end point vertex of the edge.
target	ANY-INTEGERS		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGERS:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Points SQL

Parameter	Type	Default	Description
pid	ANY-INTEGERS	value	Identifier of the point. <ul style="list-style-type: none"><li>Use with positive value, as internally will be converted to negative value</li><li>If column is present, it can not be NULL.</li><li>If column is not present, a sequential negative <b>value</b> will be given automatically.</li></ul>
edge_id	ANY-INTEGERS		Identifier of the "closest" edge to the point.
fraction	ANY-NUMERICAL		Value in <0,1> that indicates the relative position from the first end point of the edge.
side	CHAR	b	Value in [b, r, l, NULL] indicating if the point is: <ul style="list-style-type: none"><li>In the right r,</li><li>In the left l,</li><li>In both sides b, NULL</li></ul>

Where:

ANY-INTEGERS:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
source	ANY-INTEGERS	Identifier of the departure vertex.
target	ANY-INTEGERS	Identifier of the arrival vertex.

Where:

ANY-INTEGERS:

SMALLINT, INTEGER, BIGINT

Result columns

Column	Type	Description
		Identifier of the starting vertex or point.
start_pid	BIGINT	<ul style="list-style-type: none"><li>When positive: is a vertex's identifier.</li><li>When negative: is a point's identifier.</li></ul>

Column	Type	Description
		Identifier of the ending vertex or point.
end_pid	BIGINT	<ul style="list-style-type: none"><li>When positive: is a vertex's identifier.</li><li>When negative: is a point's identifier.</li></ul>

agg\_cost    FLOAT    Aggregate cost from start\_vid to end\_vid.

Additional Examples¶

- [Use pgr\\_findCloseEdges](#) in the Points SQL.
- [Right side driving topology](#)
- [Left side driving topology](#)
- [Does not matter driving side driving topology](#)

Use [pgr\\_findCloseEdges](#) in the [Points SQL ¶](#)

Find the cost of the routes from vertex\(-1\) to the two closest locations on the graph of point\(*2.9, 1.8*\).

```
SELECT * FROM pgr_withPointsCost(
  $$ SELECT * FROM edges $$,
  $$ SELECT edge_id, round(fraction::numeric, 2) AS fraction, side
    FROM pgr_findCloseEdges(
      $$ SELECT id, geom FROM edges $$,
      (SELECT ST_POINT(2.9, 1.8)),
      0.5, cap => 2)
  $$,
  1, ARRAY[-1, -2]);
start_pid | end_pid | agg_cost
-----+-----+-----
1 | -2 | 2.9
1 | -1 | 6.8
(2 rows)
```

- Point \(-1\) corresponds to the closest edge from point\(*2.9, 1.8*\).
- Point \(-2\) corresponds to the next close edge from point\(*2.9, 1.8*\).
- Being close to the graph does not mean have a shorter route.

[Right side driving topology¶](#)

Traveling from point \(-1\) and vertex \((5)\) to points \(\{(2, 3, 6)\}\) and vertices \(\{(10, 11)\}\)

```
SELECT * FROM pgr_withPointsCost(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  'SELECT pid, edge_id, fraction, side from pointsOfInterest',
  ARRAY[5, -1], ARRAY[-2, -3, -6, 10, 11],
  driving_side => 'r');
start_pid | end_pid | agg_cost
-----+-----+-----
-1 | -6 | 2.1
-1 | -3 | 4
-1 | -2 | 4.8
-1 | 10 | 6.4
-1 | 11 | 3.4
5 | -6 | 1.7
5 | -3 | 3.6
5 | -2 | 4.4
5 | 10 | 6
5 | 11 | 3
(10 rows)
```

[Left side driving topology¶](#)

Traveling from point \(-1\) and vertex \((5)\) to points \(\{(2, 3, 6)\}\) and vertices \(\{(10, 11)\}\)

```
SELECT * FROM pgr_withPointsCost(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  'SELECT pid, edge_id, fraction, side from pointsOfInterest',
  ARRAY[5, -1], ARRAY[-2, -3, -6, 10, 11],
  driving_side => 'l');
start_pid | end_pid | agg_cost
-----+-----+-----
-1 | -6 | 1.3
-1 | -3 | 3.2
-1 | -2 | 5.2
-1 | 10 | 5.6
-1 | 11 | 2.6
5 | -6 | 1.7
5 | -3 | 3.6
5 | -2 | 5.6
5 | 10 | 6
5 | 11 | 3
(10 rows)
```

[Does not matter driving side driving topology¶](#)

Traveling from point \(-1\) and vertex \((5)\) to points \(\{(2, 3, 6)\}\) and vertices \(\{(10, 11)\}\)

```
SELECT * FROM pgr_withPointsCost(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  'SELECT pid, edge_id, fraction, side from pointsOfInterest',
  ARRAY[5, -1], ARRAY[-2, -3, -6, 10, 11]);
start_pid | end_pid | agg_cost
-----+-----+-----
-1 | -6 | 1.3
-1 | -3 | 3.2
-1 | -2 | 4
-1 | 10 | 5.6
-1 | 11 | 2.6
5 | -6 | 1.7
5 | -3 | 3.6
5 | -2 | 4.4
5 | 10 | 6
5 | 11 | 3
(10 rows)
```

[Sample Data](#)

See Also

- [withPoints - Family of functions](#)

Indices and tables

- [Index](#)
- [Search Page](#)

`pgr_withPointsCostMatrix` - [proposed](#)

`pgr_withPointsCostMatrix` - Calculates a cost matrix using [pgr\\_withPoints - Proposed](#).

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Availability

- Version 2.2.0
  - New proposed function.

Description

Using Dijkstra algorithm, calculate and return a cost matrix.

Dijkstra's algorithm, conceived by Dutch computer scientist Edsger Dijkstra in 1956. It is a graph search algorithm that solves the shortest path problem for a graph with non-negative edge path costs, producing a shortest path from a starting vertex to an ending vertex. This implementation can be used with a directed graph and an undirected graph.

The main Characteristics are:

- Can be used as input to [pgr\\_TSP](#).
  - Use directly when the resulting matrix is symmetric and there is no  $\infty$  value.
  - It will be the users responsibility to make the matrix symmetric.
    - By using geometric or harmonic average of the non symmetric values.
    - By using max or min the non symmetric values.
    - By setting the upper triangle to be the mirror image of the lower triangle.
    - By setting the lower triangle to be the mirror image of the upper triangle.
  - It is also the users responsibility to fix an  $\infty$  value.
- Each function works as part of the family it belongs to.
- It does not return a path.
- Returns the sum of the costs of the shortest path for pair combination of nodes in the graph.
- Process is done only on edges with positive costs.
- Values are returned when there is a path.
  - When the starting vertex and ending vertex are the same, there is no path.
    - The aggregate cost in the non included values  $(v, v)$  is 0.
  - When the starting vertex and ending vertex are the different and there is no path.
    - The aggregate cost in the non included values  $(u, v)$  is  $\infty$ .
- Let be the case the values returned are stored in a table:
  - The unique index would be the pair:  $(start\_vid, end\_vid)$ .
- Depending on the function and its parameters, the results can be symmetric.
  - The aggregate cost of  $(u, v)$  is the same as for  $(v, u)$ .
- Any duplicated value in the **start vids** are ignored.
- The returned values are ordered:
  - start\_vid ascending
  - end\_vid ascending

Boost Graph Inside

Signatures

`pgr_withPointsCostMatrix`([Edges SQL](#), [Points SQL](#), **start vids**, [options])

**options:** [directed, driving\_side]

Returns set of (start\_vid, end\_vid, agg\_cost)

OR EMPTY SET

Note

There is no **details** flag, unlike the other members of the withPoints family of functions.

Example:

- Cost matrix for points  $\backslash(\{1, 6\})$  and vertices  $\backslash(\{10, 11\})$  on an **undirected** graph
- Returning a **symmetrical** cost matrix
  - Using the default side value on the **points\_sql** query
  - Using the default driving\_side value

```
SELECT * FROM pgr_withPointsCostMatrix(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction from pointsOfInterest',
array[-1, 10, 11, -6], directed := false);
start_vid | end_vid | agg_cost
-----+-----+-----
-6 | -1 | 1.3
-6 | 10 | 1.7
-6 | 11 | 1.3
-1 | -6 | 1.3
-1 | 10 | 1.6
-1 | 11 | 2.6
10 | -6 | 1.7
10 | -1 | 1.6
10 | 11 | 1
11 | -6 | 1.3
11 | -1 | 2.6
11 | 10 | 1
(12 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Points SQL</a>	TEXT	<a href="#">Points SQL</a> as described below

**start vids** ARRAY[BIGINT] Array of identifiers of starting vertices.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>• When true the graph is considered <i>Directed</i></li><li>• When false the graph is considered as <i>Undirected</i>.</li></ul>

With points optional parameters

Parameter	Type	Default	Description
driving_side	CHAR	b	Value in [r, l, b] indicating if the driving side is: <ul style="list-style-type: none"><li>• r for right driving side.</li><li>• l for left driving side.</li><li>• b for both.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Points SQL

Parameter	Type	Default	Description
			Identifier of the point.
pid	ANY-INTEGER	value	<ul style="list-style-type: none"><li>Use with positive value, as internally will be converted to negative value</li><li>If column is present, it can not be NULL.</li><li>If column is not present, a sequential negative <b>value</b> will be given automatically.</li></ul>
edge_id	ANY-INTEGER		Identifier of the “closest” edge to the point.
fraction	ANY-NUMERICAL		Value in <0,1> that indicates the relative position from the first end point of the edge.
			Value in [b, r, l, NULL] indicating if the point is:
side	CHAR	b	<ul style="list-style-type: none"><li>In the right <i>r</i>,</li><li>In the left <i>l</i>,</li><li>In both sides <i>b</i>, NULL</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns¶

Set of (start\_vid, end\_vid, agg\_cost)

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex.
end_vid	BIGINT	Identifier of the ending vertex.
agg_cost	FLOAT	Aggregate cost from start_vid to end_vid.

Note

When start\_vid or end\_vid columns have negative values, the identifier is for a Point.

Additional Examples¶

- Use [pgr\\_findCloseEdges](#) in the Points SQL.
- Use with [pgr\\_TSP](#).

Use [pgr\\_findCloseEdges](#) in the [Points SQL](#).¶

Find the matrix cost of the routes from vertex\(-1\) and the two closest locations on the graph of point(2.9, 1.8).

```
SELECT * FROM pgr_withPointsCostMatrix(
  $$ SELECT * FROM edges $$,
  $$ SELECT edge_id, round(fraction::numeric, 2) AS fraction, side
     FROM pgr_findCloseEdges(
        $$ SELECT id, geom FROM edges $$,
        (SELECT ST_POINT(2.9, 1.8)),
        0.5, cap => 2)
  $$,
  ARRAY[5, 10, -1, -2]);
start_vid | end_vid | agg_cost
-----+-----+-----
-2 | -1 | 3.9
-2 | 5 | 2.9
-2 | 10 | 3.1
-1 | -2 | 0.3
-1 | 5 | 3.2
-1 | 10 | 3.2
5 | -2 | 2.9
5 | -1 | 6.8
5 | 10 | 6
10 | -2 | 1.1
10 | -1 | 0.8
10 | 5 | 2
(12 rows)
```

- Point \(-1\) corresponds to the closest edge from point(2.9, 1.8).
- Point \(-2\) corresponds to the next close edge from point(2.9, 1.8).

Use with [pgr\\_TSP](#).¶

```
SELECT * FROM pgr_TSP(
  $$
  SELECT * FROM pgr_withPointsCostMatrix(
    'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
    'SELECT pid, edge_id, fraction from pointsOfInterest',
    array[-1, 10, 11, -6], directed := false);
  $$
);
NOTICE: pgr_TSP no longer solving with simulated annealing
HINT: Ignoring annealing parameters
seq | node | cost | agg_cost
-----+-----+-----+-----
1 | -6 | 0 | 0
2 | -1 | 1.3 | 1.3
3 | 10 | 1.6 | 2.9
4 | 11 | 1 | 3.9
```

See Also

- [withPoints - Family of functions](#)
- [Cost Matrix - Category](#)
- [Traveling Sales Person - Family of functions](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_withPointsKSP - Proposed

pgr\_withPointsKSP — Yen's algorithm for K shortest paths using Dijkstra.

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Version 3.6.0

- Standardizing output columns to (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)
- pgr\_withPointsKSP(One to One)
  - Signature change: driving\_side parameter changed from named optional to unnamed compulsory **driving side**.
  - Added start\_vid and end\_vid result columns.
- New proposed signatures:
  - pgr\_withPointsKSP(One to Many)
  - pgr\_withPointsKSP(Many to One)
  - pgr\_withPointsKSP(Many to Many)
  - pgr\_withPointsKSP(Combinations)
- Deprecated signature
  - pgr\_withpointsksp(text,text,bigint,bigint,integer,boolean,boolean,char,boolean)``

Version 2.2.0

- New proposed function.

Description

Modifies the graph to include the points defined in the [Points SQL](#) and using Yen algorithm, finds the \K\ shortest paths.

Boost Graph Inside

Signatures

pgr\_withPointsKSP([Edges SQL](#), [Points SQL](#), start\_vid, end\_vid, K, driving\_side, [options])  
pgr\_withPointsKSP([Edges SQL](#), [Points SQL](#), start\_vid, end\_vids, K, driving\_side, [options])  
pgr\_withPointsKSP([Edges SQL](#), [Points SQL](#), start\_vids, end\_vid, K, driving\_side, [options])  
pgr\_withPointsKSP([Edges SQL](#), [Points SQL](#), start\_vids, end\_vids, K, driving\_side, [options])  
pgr\_withPointsKSP([Edges SQL](#), [Points SQL](#), [Combinations SQL](#), K, driving\_side, [options])  
**options:** [directed, heap\_paths, details]  
Returns set of (seq, path\_id, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

One to One

pgr\_withPointsKSP([Edges SQL](#), [Points SQL](#), start\_vid, end\_vid, K, driving\_side, [options])  
**options:** [directed, heap\_paths, details]  
Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

- Get 2 paths from Point \1\ to point \2\ on a directed graph with **left** side driving.
- For a directed graph.
- No details are given about distance of other points of the query.
- No heap paths are returned.

```
SELECT * FROM pgr_withPointsKSP(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
-1, -2, 2, 'l');
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
```

1	1	1	-1	-2	-1	1	0.6	0
2	1	2	-1	-2	6	4	1	0.6
3	1	3	-1	-2	7	8	1	1.6
4	1	4	-1	-2	11	11	1	2.6
5	1	5	-1	-2	12	13	1	3.6
6	1	6	-1	-2	17	15	0.6	4.6
7	1	7	-1	-2	-2	-1	0	5.2
8	2	1	-1	-2	-1	1	0.6	0
9	2	2	-1	-2	6	4	1	0.6
10	2	3	-1	-2	7	8	1	1.6
11	2	4	-1	-2	11	9	1	2.6
12	2	5	-1	-2	16	15	1.6	3.6
13	2	6	-1	-2	-2	-1	0	5.2

(13 rows)

#### One to Many

pgr\_withPointsKSP([Edges SQL](#), [Points SQL](#), start vid, end vids, K, driving\_side, [options])

**options:** [directed, heap\_paths, details]

Returns set of (seq, path\_id, path\_seq, node, edge, cost, agg\_cost)

OR EMPTY SET

Example:

Get 2 paths from point \{1\} to point \{3\} and vertex \{7\} on an undirected graph

```
SELECT * FROM pgr_withPointsKSP(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
-1, ARRAY[-3, 7], 2, 'B',
directed => false);
```

seq	path_id	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
-----	---------	----------	-----------	---------	------	------	------	----------

1	1	1	-1	-3	-1	1	0.6	0
2	1	2	-1	-3	6	4	1	0.6
3	1	3	-1	-3	7	10	1	1.6
4	1	4	-1	-3	8	12	0.6	2.6
5	1	5	-1	-3	-3	-1	0	3.2
6	2	1	-1	-3	-1	1	0.6	0
7	2	2	-1	-3	6	4	1	0.6
8	2	3	-1	-3	7	8	1	1.6
9	2	4	-1	-3	11	11	1	2.6
10	2	5	-1	-3	12	12	0.4	3.6
11	2	6	-1	-3	-3	-1	0	4
12	3	1	-1	7	-1	1	0.6	0
13	3	2	-1	7	6	4	1	0.6
14	3	3	-1	7	7	-1	0	1.6
15	4	1	-1	7	-1	1	0.6	0
16	4	2	-1	7	6	2	1	0.6
17	4	3	-1	7	10	5	1	1.6
18	4	4	-1	7	11	8	1	2.6
19	4	5	-1	7	7	-1	0	3.6

(19 rows)

#### Many to One

pgr\_withPointsKSP([Edges SQL](#), [Points SQL](#), start vids, end vid, K, driving\_side, [options])

**options:** [directed, heap\_paths, details]

Returns set of (seq, path\_id, path\_seq, node, edge, cost, agg\_cost)

OR EMPTY SET

Example:

Get a path from point \{1\} and vertex \{6\} to point \{3\} on a **directed** graph with **right** side driving and **details** set to **True**

```
SELECT * FROM pgr_withPointsKSP(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[-1, 6], -3, 1, 'r', details=> true);
```

seq	path_id	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
-----	---------	----------	-----------	---------	------	------	------	----------

1	1	1	-1	-3	-1	1	0.4	0
2	1	2	-1	-3	5	1	1	0.4
3	1	3	-1	-3	6	4	0.7	1.4
4	1	4	-1	-3	-6	4	0.3	2.1
5	1	5	-1	-3	7	10	1	2.4
6	1	6	-1	-3	8	12	0.6	3.4
7	1	7	-1	-3	-3	-1	0	4
8	2	1	6	-3	6	4	0.7	0
9	2	2	6	-3	-6	4	0.3	0.7
10	2	3	6	-3	7	10	1	1
11	2	4	6	-3	8	12	0.6	2
12	2	5	6	-3	-3	-1	0	2.6

(12 rows)

#### Many to Many

pgr\_withPointsKSP([Edges SQL](#), [Points SQL](#), start vids, end vids, K, driving\_side, [options])

**options:** [directed, heap\_paths, details]

Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)

OR EMPTY SET

Example:

Get a path from point \{1\} and vertex \{6\} to point \{3\} and vertex \{1\} on a **directed** graph with **left** side driving and **heap\_paths** set to **True**

```
SELECT * FROM pgr_withPointsKSP(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[-1, 6], ARRAY[-3, 1], 1, 'l', heap_paths => true);
```

seq	path_id	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
-----	---------	----------	-----------	---------	------	------	------	----------

1	1	1	-1	-3	-1	1	0.6	0
2	1	2	-1	-3	6	4	1	0.6
3	1	3	-1	-3	7	10	1	1.6
4	1	4	-1	-3	8	12	0.6	2.6
5	1	5	-1	-3	-3	-1	0	3.2
6	2	1	-1	1	-1	1	0.6	0
7	2	2	-1	1	6	4	1	0.6
8	2	3	-1	1	7	7	1	1.6
9	2	4	-1	1	3	6	1	2.6
10	2	5	-1	1	1	-1	0	3.6
11	3	1	6	-3	6	4	1	0
12	3	2	6	-3	7	10	1	1
13	3	3	6	-3	8	12	0.6	2
14	3	4	6	-3	-3	-1	0	2.6
15	4	1	6	1	6	4	1	0
16	4	2	6	1	7	7	1	1
17	4	3	6	1	3	6	1	2
18	4	4	6	1	1	-1	0	3

(18 rows)



Combinations¶

pgr\_withPointsKSP([Edges SQL](#), [Points SQL](#), [Combinations SQL](#), K, driving\_side, [options])  
**options:** [directed, heap\_paths, details]  
Returns set of (seq, path\_id, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table on an **directed** graph

```
SELECT * FROM pgr_withPointsKSP(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
'SELECT * FROM (VALUES (-1, 10), (6, -3)) AS combinations(source, target)',
2, 'r', details => true);
```

seq	path_id	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	1	-1	10	-1	1	0.4	0
2	1	2	-1	10	5	1	1	0.4
3	1	3	-1	10	6	4	0.7	1.4
4	1	4	-1	10	-6	4	0.3	2.1
5	1	5	-1	10	7	8	1	2.4
6	1	6	-1	10	11	9	1	3.4
7	1	7	-1	10	16	16	1	4.4
8	1	8	-1	10	15	3	1	5.4
9	1	9	-1	10	10	-1	0	6.4
10	2	1	-1	10	-1	1	0.4	0
11	2	2	-1	10	5	1	1	0.4
12	2	3	-1	10	6	4	0.7	1.4
13	2	4	-1	10	-6	4	0.3	2.1
14	2	5	-1	10	7	8	1	2.4
15	2	6	-1	10	11	11	1	3.4
16	2	7	-1	10	12	13	1	4.4
17	2	8	-1	10	17	15	1	5.4
18	2	9	-1	10	16	16	1	6.4
19	2	10	-1	10	15	3	1	7.4
20	2	11	-1	10	10	-1	0	8.4
21	3	1	6	-3	6	4	0.7	0
22	3	2	6	-3	-6	4	0.3	0.7
23	3	3	6	-3	7	10	1	1
24	3	4	6	-3	8	12	0.6	2
25	3	5	6	-3	-3	-1	0	2.6

Parameters¶

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> query as described.
<a href="#">Points SQL</a>	TEXT	<a href="#">Points SQL</a> query as described.
start vid	ANY-INTEGER	Identifier of the departure vertex. <ul style="list-style-type: none"><li>Negative values represent a point</li></ul>
end vid	ANY-INTEGER	Identifier of the destination vertex. <ul style="list-style-type: none"><li>Negative values represent a point</li></ul>
K	ANY-INTEGER	Number of required paths
driving_side	CHAR	Value in [r, R, l, L, b, B] indicating if the driving side is: <ul style="list-style-type: none"><li>[r, R] for right driving side (for directed graph only)</li><li>[l, L] for left driving side (for directed graph only)</li><li>[b, B] for both (only for undirected graph)</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Optional parameters¶

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

KSP Optional parameters¶

Column	Type	Default	Description
heap_paths	BOOLEAN	false	<ul style="list-style-type: none"><li>When false Returns at most K paths.</li><li>When true all the calculated paths while processing are returned.</li><li>Roughly, when the shortest path hasN edges, the heap will contain about thanN * K paths for small value ofK and K &gt; 5.</li></ul>

withPointsKSP optional parameters¶

Parameter	Type	Default	Description
details	BOOLEAN	false	<ul style="list-style-type: none"> <li>When true the results will include the points that are in the path.</li> <li>When false the results will not include the points that are in the path.</li> </ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Points SQL

Parameter	Type	Default	Description
			Identifier of the point.
pid	ANY-INTEGER	value	<ul style="list-style-type: none"> <li>Use with positive value, as internally will be converted to negative value</li> <li>If column is present, it can not be NULL.</li> <li>If column is not present, a sequential negative <b>value</b> will be given automatically.</li> </ul>
edge_id	ANY-INTEGER		Identifier of the "closest" edge to the point.
fraction	ANY-NUMERICAL		Value in <0,1> that indicates the relative position from the first end point of the edge.
side	CHAR	b	Value in [b, r, l, NULL] indicating if the point is: <ul style="list-style-type: none"> <li>In the right r,</li> <li>In the left l,</li> <li>In both sides b, NULL</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns

Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
Path identifier.		
path_id	INTEGER	<ul style="list-style-type: none"><li>Has value 1 for the first of a path from start_vid to end_vid</li></ul>
path_seq	INTEGER	Relative position in the path. Has value 1 for the beginning of a path.
node	BIGINT	Identifier of the node in the path from start_vid to end_vid
edge	BIGINT	Identifier of the edge used to go from node to the next node in the path sequence. -1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence. <ul style="list-style-type: none"><li>\(0\)</li></ul>
agg_cost	FLOAT	Aggregate cost from start vid to node.

Additional Examples¶

- Use pgr\_findCloseEdges in the Points SQL.
- Left driving side
- Right driving side

Use pgr\_findCloseEdges in the Points SQL.¶

Get \(\) paths using left side driving topology, from vertex\(\) to the closest location on the graph of point(2.9, 1.8).

```
SELECT * FROM pgr_withPointsKSP(
  $e$ SELECT * FROM edges $e$,
  $p$ SELECT edge_id, round(fraction::numeric, 2) AS fraction, side
    FROM pgr_findCloseEdges(
      $$SELECT id, geom FROM edges $$,
      (SELECT ST_POINT(2.9, 1.8)),
      0.5, cap => 2)
  $p$,
  1, -1, 2, 'r');
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | -1 | 1 | 6 | 1 | 0
2 | 1 | 2 | 1 | -1 | 3 | 7 | 1 | 1
3 | 1 | 3 | 1 | -1 | 7 | 8 | 1 | 2
4 | 1 | 4 | 1 | -1 | 11 | 9 | 1 | 3
5 | 1 | 5 | 1 | -1 | 16 | 16 | 1 | 4
6 | 1 | 6 | 1 | -1 | 15 | 3 | 1 | 5
7 | 1 | 7 | 1 | -1 | 10 | 5 | 0.8 | 6
8 | 1 | 8 | 1 | -1 | -1 | -1 | 0 | 6.8
9 | 2 | 1 | 1 | -1 | 1 | 6 | 1 | 0
10 | 2 | 2 | 1 | -1 | 3 | 7 | 1 | 1
11 | 2 | 3 | 1 | -1 | 7 | 10 | 1 | 2
12 | 2 | 4 | 1 | -1 | 8 | 12 | 1 | 3
13 | 2 | 5 | 1 | -1 | 12 | 13 | 1 | 4
14 | 2 | 6 | 1 | -1 | 17 | 15 | 1 | 5
15 | 2 | 7 | 1 | -1 | 16 | 16 | 1 | 6
16 | 2 | 8 | 1 | -1 | 15 | 3 | 1 | 7
17 | 2 | 9 | 1 | -1 | 10 | 5 | 0.8 | 8
18 | 2 | 10 | 1 | -1 | -1 | -1 | 0 | 8.8
(18 rows)
```

- Point \(-\)

Left driving side¶

Get \(\) paths using left side driving topology, from point\(\) to point\(\) with details.

```
SELECT * FROM pgr_withPointsKSP(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  'SELECT pid, edge_id, fraction, side from pointsOfInterest',
  -1, -3, 2, 'l', details => true);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | -3 | -1 | 1 | 0.6 | 0
2 | 1 | 2 | 1 | -3 | 6 | 4 | 0.7 | 0.6
3 | 1 | 3 | 1 | -3 | -6 | 4 | 0.3 | 1.3
4 | 1 | 4 | 1 | -3 | 7 | 10 | 1 | 1.6
5 | 1 | 5 | 1 | -3 | 8 | 12 | 0.6 | 2.6
6 | 1 | 6 | 1 | -3 | -3 | -1 | 0 | 3.2
(6 rows)
```

Right driving side¶

Get \(\) paths using right side driving topology from, point\(\) to point\(\) with heap paths and details.

```
SELECT * FROM pgr_withPointsKSP(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  'SELECT pid, edge_id, fraction, side from pointsOfInterest',
  -1, -2, 2, 'r',
  heap_paths => true, details => true);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | -2 | -1 | 1 | 0.4 | 0
2 | 1 | 2 | 1 | -2 | 5 | 1 | 1 | 0.4
3 | 1 | 3 | 1 | -2 | 6 | 4 | 0.7 | 1.4
4 | 1 | 4 | 1 | -2 | -6 | 4 | 0.3 | 2.1
5 | 1 | 5 | 1 | -2 | 7 | 8 | 1 | 2.4
6 | 1 | 6 | 1 | -2 | 11 | 9 | 1 | 3.4
7 | 1 | 7 | 1 | -2 | 16 | 15 | 0.4 | 4.4
8 | 1 | 8 | 1 | -2 | -2 | -1 | 0 | 4.8
9 | 2 | 1 | 1 | -2 | -1 | 1 | 0.4 | 0
10 | 2 | 2 | 1 | -2 | 5 | 1 | 1 | 0.4
11 | 2 | 3 | 1 | -2 | 6 | 4 | 0.7 | 1.4
12 | 2 | 4 | 1 | -2 | -6 | 4 | 0.3 | 2.1
13 | 2 | 5 | 1 | -2 | 7 | 8 | 1 | 2.4
14 | 2 | 6 | 1 | -2 | 11 | 11 | 1 | 3.4
15 | 2 | 7 | 1 | -2 | 12 | 13 | 1 | 4.4
```

16	2	8	-1	-2	17	15	1	5.4
17	2	9	-1	-2	16	15	0.4	6.4
18	2	10	-1	-2	-2	-1	0	6.8
19	3	1	-1	-2	-1	1	0.4	0
20	3	2	-1	-2	5	1	1	0.4
21	3	3	-1	-2	6	4	0.7	1.4
22	3	4	-1	-2	-6	4	0.3	2.1
23	3	5	-1	-2	7	10	1	2.4
24	3	6	-1	-2	8	12	0.6	3.4
25	3	7	-1	-2	-3	12	0.4	4
26	3	8	-1	-2	12	13	1	4.4
27	3	9	-1	-2	17	15	1	5.4
28	3	10	-1	-2	16	15	0.4	6.4
29	3	11	-1	-2	-2	-1	0	6.8

(29 rows)

See Also

- [withPoints - Family of functions](#)
- [K shortest paths - Category](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_withPointsDD - Proposed

pgr\_withPointsDD - Returns the driving **distance** from a starting point.

Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Availability

Version 3.6.0

- Signature change: `driving_side` parameter changed from named optional to unnamed compulsory **driving side**.
  - `pgr_withPointsDD`(Single vertex)
  - `pgr_withPointsDD`(Multiple vertices)
- Standardizing output columns to (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
  - `pgr_withPointsDD`(Single vertex)
    - Added depth, pred and start\_vid column.
  - `pgr_withPointsDD`(Multiple vertices)
    - Added depth, pred columns.
- When `details` is false:
  - Only points that are visited are removed, that is, points reached within the distance are included
- Deprecated signatures
  - `pgr_withpointsdd`(text,text,bigint,double precision,boolean,character,boolean)
  - `pgr_withpointsdd`(text,text,array,double precision,boolean,character,boolean,boolean)

Version 2.2.0

- New proposed function.

Description

Modify the graph to include points and using Dijkstra algorithm, extracts all the nodes and points that have costs less than or equal to the value `distance` from the starting point. The edges extracted will conform the corresponding spanning tree.

Boost Graph Inside

Signatures

`pgr_withPointsDD`([Edges SQL](#), [Points SQL](#), **root vid**, **distance**, **driving side**, [options A])  
`pgr_withPointsDD`([Edges SQL](#), [Points SQL](#), **root vids**, **distance**, **driving side**, [options B])  
**options A:** [directed, details]  
**options B:** [directed, details, equicost]  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)  
OR EMPTY SET

Single vertex

`pgr_withPointsDD`([Edges SQL](#), [Points SQL](#), **root vid**, **distance**, **driving side**, [options])  
**options:** [directed, details]  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Right side driving topology, from point\1\ within a distance of\3.3\ with details.

```
SELECT * FROM pgr_withPointsDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
-1, 3.3, 'r',
details => true);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | -1 | -1 | -1 | -1 | 0 | 0
2 | 1 | -1 | -1 | 5 | 1 | 0.4 | 0.4
3 | 2 | -1 | 5 | 6 | 1 | 1 | 1.4
4 | 3 | -1 | 6 | -6 | 4 | 0.7 | 2.1
5 | 4 | -1 | -6 | 7 | 4 | 0.3 | 2.4
(5 rows)
```

Multiple vertices1

pgr\_withPointsDD([Edges SQL](#), [Points SQL](#), root vids, distance, driving side, [options])  
options: [directed, details, equicost]  
Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From point \1\ and vertex \16\ within a distance of \3.3\ with equicost on a directed graph

```
SELECT * FROM pgr_withPointsDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[-1, 16], 3.3, 'r',
equicost => true);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | -1 | -1 | -1 | -1 | 0 | 0
2 | 1 | -1 | -1 | 6 | 1 | 0.6 | 0.6
3 | 2 | -1 | 6 | 7 | 4 | 1 | 1.6
4 | 2 | -1 | 6 | 5 | 1 | 1 | 1.6
5 | 3 | -1 | 7 | 3 | 7 | 1 | 2.6
6 | 3 | -1 | 7 | 8 | 10 | 1 | 2.6
7 | 4 | -1 | 8 | -3 | 12 | 0.6 | 3.2
8 | 4 | -1 | 3 | -4 | 6 | 0.7 | 3.3
9 | 0 | 16 | 16 | 16 | -1 | 0 | 0
10 | 1 | 16 | 16 | 11 | 9 | 1 | 1
11 | 1 | 16 | 16 | 15 | 16 | 1 | 1
12 | 1 | 16 | 16 | 17 | 15 | 1 | 1
13 | 2 | 16 | 15 | 10 | 3 | 1 | 2
14 | 2 | 16 | 11 | 12 | 11 | 1 | 2
(14 rows)
```

Parameters1

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Points SQL</a>	TEXT	<a href="#">Points SQL</a> as described below
Root vid	BIGINT	Identifier of the root vertex of the tree. <ul style="list-style-type: none"><li>Negative values represent a point</li></ul>
		Array of identifiers of the root vertices. <ul style="list-style-type: none"><li>Negative values represent a point</li></ul>
Root vids	ARRAY [ANY-INTEGER]	<ul style="list-style-type: none"><li>\0\ values are ignored</li><li>For optimization purposes, any duplicated value is ignored.</li></ul>
distance	FLOAT	Upper limit for the inclusion of a node in the result. <ul style="list-style-type: none"><li>Value in [r, R, l, L, b, B] indicating if the driving side is:<ul style="list-style-type: none"><li>r, R for right driving side,</li><li>l, L for left driving side.</li><li>b, B for both.</li></ul></li><li>Valid values differ for directed and undirected graphs:<ul style="list-style-type: none"><li>In directed graphs: [r, R, l, L].</li><li>In undirected graphs: [b, B].</li></ul></li></ul>
driving side	CHAR	

Where:

ANY-INTEGER:  
SMALLINT, INTEGER, BIGINT

Optional parameters1

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

With points optional parameters1

Parameter	Type	Default	Description
details	BOOLEAN	false	<ul style="list-style-type: none"> <li>When true the results will include the points that are in the path.</li> <li>When false the results will not include the points that are in the path.</li> </ul>

Driving distance optional parameters

Column	Type	Default	Description
equicost	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true the node will only appear in the closeststart_vid list. Tie brakes are arbitrary.</li> <li>When false which resembles several calls using the single vertex signature.</li> </ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Points SQL

Parameter	Type	Default	Description
pid	ANY-INTEGER	value	Identifier of the point. <ul style="list-style-type: none"> <li>Use with positive value, as internally will be converted to negative value</li> <li>If column is present, it can not be NULL.</li> <li>If column is not present, a sequential negativevalue will be given automatically.</li> </ul>
edge_id	ANY-INTEGER		Identifier of the "closest" edge to the point.
fraction	ANY-NUMERICAL		Value in <0,1> that indicates the relative position from the first end point of the edge.
side	CHAR	b	Value in [b, r, l, NULL] indicating if the point is: <ul style="list-style-type: none"> <li>In the rightr,</li> <li>In the leftl,</li> <li>In both sidesb, NULL</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

Parameter	Type	Description
seq	BIGINT	Sequential value starting from \{1\}.

Parameter	Type	Description
Depth of the node.		
depth	BIGINT	<ul style="list-style-type: none"><li><math>\backslash(0)</math> when node = start_vid.</li><li><math>\backslash(\text{depth}-1)</math> is the depth of pred</li></ul>
start_vid	BIGINT	Identifier of the root vertex.
Predecessor of node.		
pred	BIGINT	<ul style="list-style-type: none"><li>When node = start_vid then has the value node.</li></ul>
node	BIGINT	Identifier of node reached using edge.
Identifier of the edge used to arrive from pred to node.		
edge	BIGINT	<ul style="list-style-type: none"><li><math>\backslash(-1)</math> when node = start_vid.</li></ul>
cost	FLOAT	Cost to traverse edge.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Additional Examples¶

- Use [pgr\\_findCloseEdges](#) in the Points SQL.
- [Driving side does not matter](#)

Use [pgr\\_findCloseEdges](#) in the [Points SQL ¶](#)

Find the driving distance from the two closest locations on the graph of point(2.9, 1.8).

```
SELECT * FROM pgr_withPointsDD(
  $e$ SELECT * FROM edges $e$,
  $p$ SELECT edge_id, round(fraction::numeric, 2) AS fraction, side
    FROM pgr_findCloseEdges(
      $$SELECT id, geom FROM edges$,
      (SELECT ST_POINT(2.9, 1.8)),
      0.5, cap => 2)
  $p$,
  ARRAY[-1, -2], 2.3, 'r',
  details => true);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | -2 | -2 | -1 | 0 | 0
2 | 1 | -2 | -2 | 11 | 8 | 0.1 | 0.1
3 | 2 | -2 | 11 | 16 | 9 | 1 | 1.1
4 | 2 | -2 | 11 | 12 | 11 | 1 | 1.1
5 | 2 | -2 | 11 | 7 | 8 | 1 | 1.1
6 | 3 | -2 | 12 | 17 | 13 | 1 | 2.1
7 | 3 | -2 | 16 | 15 | 16 | 1 | 2.1
8 | 3 | -2 | 7 | 8 | 10 | 1 | 2.1
9 | 3 | -2 | 7 | 6 | 4 | 1 | 2.1
10 | 3 | -2 | 7 | 3 | 7 | 1 | 2.1
11 | 0 | -1 | -1 | -1 | -1 | 0 | 0
12 | 1 | -1 | -1 | 11 | 5 | 0.2 | 0.2
13 | 2 | -1 | 11 | 7 | 8 | 1 | 1.2
14 | 2 | -1 | 11 | 16 | 9 | 1 | 1.2
15 | 2 | -1 | 11 | 12 | 11 | 1 | 1.2
16 | 3 | -1 | 7 | -2 | 8 | 0.9 | 2.1
17 | 3 | -1 | 7 | 3 | 7 | 1 | 2.2
18 | 3 | -1 | 7 | 6 | 4 | 1 | 2.2
19 | 3 | -1 | 7 | 8 | 10 | 1 | 2.2
20 | 3 | -1 | 16 | 15 | 16 | 1 | 2.2
21 | 3 | -1 | 12 | 17 | 13 | 1 | 2.2
(21 rows)
```

- Point  $\backslash(-1)$  corresponds to the closest edge from point $\backslash((2.9, 1.8))$ .
- Point  $\backslash(-2)$  corresponds to the next close edge from point $\backslash((2.9, 1.8))$ .

[Driving side does not matter¶](#)

From point  $\backslash(1)$  within a distance of  $\backslash(3.3)$ , does not matter driving side, with details.

```
SELECT * FROM pgr_withPointsDD(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  'SELECT pid, edge_id, fraction, side from pointsOfInterest',
  -1, 3.3, 'b',
  directed => false,
  details => true);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | -1 | -1 | -1 | -1 | 0 | 0
2 | 1 | -1 | -1 | 5 | 1 | 0.4 | 0.4
3 | 1 | -1 | -1 | 6 | 1 | 0.6 | 0.6
4 | 2 | -1 | 6 | -6 | 4 | 0.7 | 1.3
5 | 2 | -1 | 6 | 10 | 2 | 1 | 1.6
6 | 3 | -1 | -6 | 7 | 4 | 0.3 | 1.6
7 | 3 | -1 | 10 | -5 | 5 | 0.8 | 2.4
8 | 3 | -1 | 10 | 15 | 3 | 1 | 2.6
9 | 4 | -1 | 7 | 3 | 7 | 1 | 2.6
10 | 4 | -1 | 7 | 8 | 10 | 1 | 2.6
11 | 4 | -1 | 7 | 11 | 8 | 1 | 2.6
12 | 5 | -1 | 8 | -3 | 12 | 0.6 | 3.2
13 | 5 | -1 | 3 | -4 | 6 | 0.7 | 3.3
(13 rows)
```

See Also¶

- [pgr\\_drivingDistance](#)
- [Sample Data](#)

Indices and tables

- [Index](#)

- [Search Page](#)

pgr\_withPointsVia - Proposed

pgr\_withPointsVia - Route that goes through a list of vertices and/or points.

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

Availability

- Version 3.4.0
  - New proposed function.

Description

Given a graph, a set of points on the graphs edges and a list of vertices, this function is equivalent to finding the shortest path between(vertex\_i) and (vertex\_{i+1}) (where (vertex\)) can be a vertex or a point on the graph) for all (i < size\_of(via\;vertices)).

Route:

is a sequence of paths.

Path:

is a section of the route.

The general algorithm is as follows:

- Build the Graph with the new points.
  - The points identifiers will be converted to negative values.
  - The vertices identifiers will remain positive.
- Execute a [pgr\\_dijkstraVia - Proposed](#).

 Boost Graph Inside

Signatures

One Via

pgr\_withPointsVia([Edges SQL](#), [Points SQL](#), **via vertices**, **[options]**)

**options:** [directed, strict, U\_turn\_on\_edge]

Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost, route\_agg\_cost)  
OR EMPTY SET

Example:

Find the route that visits the vertices\({ -6, 15, -1}\)in that order on a**directed** graph.

```
SELECT * FROM pgr_withPointsVia(
'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[-6, 15, -1]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -6 | 15 | -6 | 4 | 0.3 | 0 | 0
2 | 1 | 2 | -6 | 15 | 7 | 8 | 0.3 | 0.3 | 0.3
3 | 1 | 3 | -6 | 15 | 11 | 9 | 1 | 1.3 | 1.3
4 | 1 | 4 | -6 | 15 | 16 | 16 | 1 | 2.3 | 2.3
5 | 1 | 5 | -6 | 15 | 15 | -1 | 0 | 3.3 | 3.3
6 | 2 | 1 | 15 | -1 | 15 | 3 | 1 | 0 | 3.3
7 | 2 | 2 | 15 | -1 | 10 | 2 | 1 | 1 | 4.3
8 | 2 | 3 | 15 | -1 | 6 | 1 | 0.6 | 2 | 5.3
9 | 2 | 4 | 15 | -1 | -1 | -2 | 0 | 2.6 | 5.9
(9 rows)
```

Parameters

Parameter	Type	Default	Description
<a href="#">Edges SQL</a>	TEXT		SQL query as described.
<a href="#">Points SQL</a>	TEXT		SQL query as described.
<b>via vertices</b>	ARRAY [ <b>ANY-INTEGER</b> ]		Array of ordered vertices identifiers that are going to be visited. <ul style="list-style-type: none"><li>• When positive it is considered a vertex identifier</li><li>• When negative it is considered a point identifier</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT



ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li></ul>
			<ul style="list-style-type: none"><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Via optional parameters

Parameter	Type	Default	Description
strict	BOOLEAN	false	<ul style="list-style-type: none"><li>When true if a path is missing stops and returns <b>EMPTY SET</b></li></ul>
			<ul style="list-style-type: none"><li>When false ignores missing paths returning all paths found</li></ul>
U_turn_on_edge	BOOLEAN	true	<ul style="list-style-type: none"><li>When true departing from a visited vertex will not try to avoid</li></ul>

With points optional parameters

Parameter	Type	Default	Description
driving_side	CHAR	b	Value in [r, l, b] indicating if the driving side is: <ul style="list-style-type: none"><li>r for right driving side.</li></ul>
			<ul style="list-style-type: none"><li>l for left driving side.</li></ul>
			<ul style="list-style-type: none"><li>b for both.</li></ul>
details	BOOLEAN	false	<ul style="list-style-type: none"><li>When true the results will include the points that are in the path.</li></ul>
			<ul style="list-style-type: none"><li>When false the results will not include the points that are in the path.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source)
			<ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Points SQL

Parameter	Type	Default	Description
pid	ANY-INTEGER	value	Identifier of the point.
			<ul style="list-style-type: none"><li>Use with positive value, as internally will be converted to negative value</li></ul>
			<ul style="list-style-type: none"><li>If column is present, it can not be NULL.</li></ul>
edge_id	ANY-INTEGER		<ul style="list-style-type: none"><li>If column is not present, a sequential negative <b>value</b> will be given automatically.</li></ul>
			Identifier of the "closest" edge to the point.
fraction	ANY-NUMERICAL		Value in <0,1> that indicates the relative position from the first end point of the edge.

Parameter	Type	Default	Description
			Value in [b, r, l, NULL] indicating if the point is: <ul style="list-style-type: none"><li>In the right <i>r</i>,</li><li>In the left <i>l</i>,</li><li>In both sides <i>b</i>, NULL</li></ul>
side	CHAR	b	

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_id	INTEGER	Identifier of a path. Has value 1 for the first path.
path_seq	INTEGER	Relative position in the path. Has value 1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex of the path.
end_vid	BIGINT	Identifier of the ending vertex of the path.
node	BIGINT	Identifier of the node in the path from start_vid to end_vid.
		Identifier of the edge used to go from node to the next node in the path sequence.
edge	BIGINT	<ul style="list-style-type: none"><li>-1 for the last node of the path.</li><li>-2 for the last node of the route.</li></ul>
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.
route_agg_cost	FLOAT	Total cost from start_vid of seq = 1 to end_vid of the current seq.

Note

When start\_vid, end\_vid and node columns have negative values, the identifier is for a Point.

Additional Examples

- Use [pgr\\_findCloseEdges](#) in the Points SQL
- Usage variations
  - Aggregate cost of the third path.
  - Route's aggregate cost of the route at the end of the third path.
  - Nodes visited in the route.
  - The aggregate costs of the route when the visited vertices are reached.
  - Status of "passes in front" or "visits" of the nodes and points.

Use [pgr\\_findCloseEdges](#) in the Points SQL

Visit from vertex \(-1\\) to the two locations on the graph of point(2.9, 1.8) in order of closeness to the graph.

```
SELECT * FROM pgr_withPointsVia(
  $e$ SELECT * FROM edges $e$,
  $p$ SELECT edge_id, round(fraction::numeric, 2) AS fraction, side
    FROM pgr_findCloseEdges(
      $$SELECT id, geom FROM edges$,
      (SELECT ST_POINT(2.9, 1.8)),
      0.5, cap => 2)
  $p$,
  ARRAY[1, -1, -2], details => true);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | -1 | 1 | 6 | 1 | 0 | 0
2 | 1 | 2 | 1 | -1 | 3 | 7 | 1 | 1 | 1
3 | 1 | 3 | 1 | -1 | 7 | 8 | 0.9 | 2 | 2
4 | 1 | 4 | 1 | -1 | 2 | 8 | 0.1 | 2.9 | 2.9
5 | 1 | 5 | 1 | -1 | 11 | 9 | 1 | 3 | 3
6 | 1 | 6 | 1 | -1 | 16 | 16 | 1 | 4 | 4
7 | 1 | 7 | 1 | -1 | 15 | 3 | 1 | 5 | 5
8 | 1 | 8 | 1 | -1 | 10 | 5 | 0.8 | 6 | 6
9 | 1 | 9 | 1 | -1 | -1 | 0 | 6.8 | 6.8 | 6.8
10 | 2 | 1 | -1 | -2 | -1 | 5 | 0.2 | 0 | 6.8
11 | 2 | 2 | -1 | -2 | 11 | 8 | 0.1 | 0.2 | 7
12 | 2 | 3 | -1 | -2 | -2 | -2 | 0 | 0.3 | 7.1
(12 rows)
```

- Point \(-1\\) corresponds to the closest edge from point(2.9, 1.8).
- Point \(-2\\) corresponds to the next close edge from point(2.9, 1.8).

- Point \(-2\)) is visited on the route to from vertex\(\{1\}) to Point\(-1\)) (See row where \(\text{seq} = 4\)).

[Usage variations¶](#)

All this examples are about the route that visits the vertices\(\{-1, 7, -3, 16, 15\}\) in that order on a **directed** graph.

```
SELECT * FROM pgr_withPointsVia(
'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[-1, 7, -3, 16, 15]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -1 | 7 | 7 | -1 | 0.6 | 0 | 0
2 | 1 | 2 | -1 | 7 | 6 | 4 | 1 | 0.6 | 0.6
3 | 1 | 3 | -1 | 7 | 7 | -1 | 0 | 1.6 | 1.6
4 | 2 | 1 | 7 | -3 | 7 | 10 | 1 | 0 | 1.6
5 | 2 | 2 | 7 | -3 | 8 | 12 | 0.6 | 1 | 2.6
6 | 2 | 3 | 7 | -3 | -3 | -1 | 0 | 1.6 | 3.2
7 | 3 | 1 | -3 | 16 | -3 | 12 | 0.4 | 0 | 3.2
8 | 3 | 2 | -3 | 16 | 12 | 13 | 1 | 0.4 | 3.6
9 | 3 | 3 | -3 | 16 | 17 | 15 | 1 | 1.4 | 4.6
10 | 3 | 4 | -3 | 16 | 16 | -1 | 0 | 2.4 | 5.6
11 | 4 | 1 | 16 | 15 | 16 | 16 | 1 | 0 | 5.6
12 | 4 | 2 | 16 | 15 | 15 | -2 | 0 | 1 | 6.6
(12 rows)
```

[Aggregate cost of the third path¶](#)

```
SELECT agg_cost FROM pgr_withPointsVia(
'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[-1, 7, -3, 16, 15])
WHERE path_id = 3 AND edge < 0;
agg_cost
-----
2.4
(1 row)
```

[Route's aggregate cost of the route at the end of the third path¶](#)

```
SELECT route_agg_cost FROM pgr_withPointsVia(
'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[-1, 7, -3, 16, 15])
WHERE path_id = 3 AND edge < 0;
route_agg_cost
-----
5.6
(1 row)
```

[Nodes visited in the route.¶](#)

```
SELECT row_number() over () as node_seq, node
FROM pgr_withPointsVia(
'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[-1, 7, -3, 16, 15])
WHERE edge <> -1 ORDER BY seq;
node_seq | node
-----+-----
1 | -1
2 | 6
3 | 7
4 | 8
5 | -3
6 | 12
7 | 17
8 | 16
9 | 15
(9 rows)
```

[The aggregate costs of the route when the visited vertices are reached¶](#)

```
SELECT path_id, route_agg_cost FROM pgr_withPointsVia(
'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[-1, 7, -3, 16, 15])
WHERE edge < 0;
path_id | route_agg_cost
-----+-----
1 | 1.6
2 | 3.2
3 | 5.6
4 | 6.6
(4 rows)
```

[Status of "passes in front" or "visits" of the nodes and points.¶](#)

```
SELECT seq, node,
CASE WHEN edge = -1 THEN 'visits'
ELSE 'passes in front'
END as status
FROM pgr_withPointsVia(
'SELECT id, source, target, cost, reverse_cost FROM edges order by id',
'SELECT pid, edge_id, fraction, side from pointsOfInterest',
ARRAY[-1, 7, -3, 16, 15], details => true)
WHERE agg_cost <> 0 or seq = 1;
seq | node | status
-----+-----+-----
1 | -1 | passes in front
2 | 6 | passes in front
3 | -6 | passes in front
4 | 7 | visits
6 | 8 | passes in front
7 | -3 | visits
9 | 12 | passes in front
10 | 17 | passes in front
11 | -2 | passes in front
12 | 16 | visits
14 | 15 | passes in front
(11 rows)
```

[See Also¶](#)

- [withPoints - Family of functions](#)
- [Via - Category](#)

- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Introduction¶

This family of functions belongs to the [withPoints - Category](#) and the functions that compose them are based one way or another on dijkstra algorithm.

Depending on the name:

- pgr\_withPoints is pgr\_dijkstra **with points**
- pgr\_withPointsCost is pgr\_dijkstraCost **with points**
- pgr\_withPointsCostMatrix is pgr\_dijkstraCostMatrix **with points**
- pgr\_withPointsKSP is pgr\_ksp **with points**
- pgr\_withPointsDD is pgr\_drivingDistance **with points**
- pgr\_withPointsvia is pgr\_dijkstraVia **with points**

Parameters¶

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Points SQL</a>	TEXT	<a href="#">Points SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path. Negative value is for point's identifier.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices. Negative values are for point's identifiers.
end vid	BIGINT	Identifier of the ending vertex of the path. Negative value is for point's identifier.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices. Negative values are for point's identifiers.

Optional parameters¶

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

With points optional parameters¶

Parameter	Type	Default	Description
driving_side	CHAR	b	Value in [r, l, b] indicating if the driving side is: <ul style="list-style-type: none"><li>r for right driving side.</li><li>l for left driving side.</li><li>b for both.</li></ul>
details	BOOLEAN	false	<ul style="list-style-type: none"><li>When true the results will include the points that are in the path.</li><li>When false the results will not include the points that are in the path.</li></ul>

Inner Queries¶

Edges SQL¶

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)

Column	Type	Default	Description
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Points [SQL¶](#)

Parameter	Type	Default	Description
			Identifier of the point. <ul style="list-style-type: none"> <li>Use with positive value, as internally will be converted to negative value</li> <li>If column is present, it can not be NULL.</li> <li>If column is not present, a sequential negative <b>value</b> will be given automatically.</li> </ul>
pid	ANY-INTEGER	value	
edge_id	ANY-INTEGER		Identifier of the “closest” edge to the point.
fraction	ANY-NUMERICAL		Value in <0,1> that indicates the relative position from the first end point of the edge.
			Value in [b, r, l, NULL] indicating if the point is: <ul style="list-style-type: none"> <li>In the right r,</li> <li>In the left l,</li> <li>In both sides b, NULL</li> </ul>
side	CHAR	b	

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations [SQL¶](#)

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Advanced Documentation [¶](#)


Contents

- [About points](#)
- [Driving side](#)
  - [Right driving side](#)
  - [Left driving side](#)
  - [Driving side does not matter](#)
- [Creating temporary vertices](#)
  - [On a right hand side driving network](#)
  - [On a left hand side driving network](#)
  - [When driving side does not matter](#)

[About points¶](#)

For this section the following city (see [Sample Data](#)) some interesting points such as restaurant, supermarket, post office, etc. will be used as example.

[\\_images/Fig1-originalData.png](#)



- The graph is **directed**
- Red arrows show the (source, target) of the edge on the edge table
- Blue arrows show the (target, source) of the edge on the edge table
- Each point location shows where it is located with relation of the edge(source, target)
  - On the right for points **2** and **4**.
  - On the left for points **1**, **3** and **5**.
  - On both sides for point **6**.

The representation on the data base follows the [Points SQL](#) description, and for this example:

```
SELECT pid, edge_id, fraction, side FROM pointsOfInterest;
pid | edge_id | fraction | side
-----
1 | 1 | 0.4 | l
4 | 6 | 0.3 | r
3 | 12 | 0.6 | l
2 | 15 | 0.4 | r
5 | 5 | 0.8 | l
6 | 4 | 0.7 | b
(6 rows)
```

[Driving side¶](#)

In the following images:

- The squared vertices are the temporary vertices,
- The temporary vertices are added according to the driving side,
- visually showing the differences on how depending on the driving side the data is interpreted.

[Right driving side¶](#)

[\\_images/rightDrivingSide.png](#)



- Point **1** located on edge (6, 5)
- Point **2** located on edge (16, 17)
- Point **3** located on edge (8, 12)
- Point **4** located on edge (1, 3)
- Point **5** located on edge (10, 11)
- Point **6** located on edges (6, 7) and (7, 6)

[Left driving side¶](#)

[\\_images/leftDrivingSide.png](#)



- Point 1 located on edge (5, 6)
- Point 2 located on edge (17, 16)
- Point 3 located on edge (8, 12)
- Point 4 located on edge (3, 1)
- Point 5 located on edge (10, 11)
- Point 6 located on edges (6, 7) and (7, 6)

[Driving side does not matter¶](#)

- Like having all points to be considered in both sides
- Preferred usage on **undirected** graphs
- On the [TRSP - Family of functions](#) this option is not valid

[\\_images/noMatterDrivingSide.png](#)



- Point 1 located on edge (5, 6) and (6, 5)
- Point 2 located on edge (17, 16) and 16, 17
- Point 3 located on edge (8, 12)
- Point 4 located on edge (3, 1) and (1, 3)
- Point 5 located on edge (10, 11)
- Point 6 located on edges (6, 7) and (7, 6)

[Creating temporary vertices¶](#)

This section will demonstrate how a temporary vertex is created internally on the graph.

Problem

For edge:

```
SELECT id, source, target, cost, reverse_cost
FROM edges WHERE id = 15;
id | source | target | cost | reverse_cost
-----+-----+-----+-----+-----
15 | 16 | 17 | 1 | 1
(1 row)
```

insert point:

```
SELECT pid, edge_id, fraction, side
FROM pointsOfInterest WHERE pid = 2;
pid | edge_id | fraction | side
-----+-----+-----+-----
2 | 15 | 0.4 | r
(1 row)
```

[On a right hand side driving network¶](#)

Right driving side

[\\_images/rightDrivingSide.png](#)



- Arrival to point -2 can be achieved only via vertex **16**.
- Does not affect edge (17, 16), therefore the edge is kept.
- It only affects the edge (16, 17), therefore the edge is removed.
- Create two new edges:
  - Edge (16, -2) with cost 0.4 (original cost \* fraction ==  $\backslash(1 * 0.4\backslash)$ )
  - Edge (-2, 17) with cost 0.6 (the remaining cost)
- The total cost of the additional edges is equal to the original cost.
- If more points are on the same edge, the process is repeated recursively.

[On a left hand side driving network](#)

Left driving side

[\\_images/leftDrivingSide.png](#)



- Arrival to point -2 can be achieved only via vertex **17**.
- Does not affect edge (16, 17), therefore the edge is kept.
- It only affects the edge (17, 16), therefore the edge is removed.
- Create two new edges:
  - Work with the original edge (16, 17) as the fraction is a fraction of the original:
    - Edge (16, -2) with cost 0.4 (original cost \* fraction ==  $\backslash(1 * 0.4\backslash)$ )
    - Edge (-2, 17) with cost 0.6 (the remaining cost)
    - If more points are on the same edge, the process is repeated recursively.
  - Flip the Edges and add them to the graph:
    - Edge (17, -2) becomes (-2, 16) with cost 0.4 and is added to the graph.
    - Edge (-2, 16) becomes (17, -2) with cost 0.6 and is added to the graph.
- The total cost of the additional edges is equal to the original cost.

[When driving side does not matter](#)



[\\_images/noMatterDrivingSide.png](#)



- Arrival to point -2 can be achieved via vertices **16** or **17**.
- Affects the edges (16, 17) and (17, 16), therefore the edges are removed.
- Create four new edges:
  - Work with the original edge (16, 17) as the fraction is a fraction of the original:
    - Edge (16, -2) with cost 0.4 (original cost \* fraction == (1 \* 0.4))
    - Edge (-2, 17) with cost 0.6 (the remaining cost)
    - If more points are on the same edge, the process is repeated recursively.
  - Flip the Edges and add all the edges to the graph:
    - Edge (16, -2) is added to the graph.
    - Edge (-2, 17) is added to the graph.
    - Edge (16, -2) becomes (-2, 16) with cost 0.4 and is added to the graph.
    - Edge (-2, 17) becomes (17, -2) with cost 0.6 and is added to the graph.

See Also

- [withPoints - Category](#)

Indices and tables

- [Index](#)
- [Search Page](#)

See Also

- [Experimental Functions](#)

Indices and tables

- [Index](#)
- [Search Page](#)

## Experimental Functions

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Families

[Flow - Family of functions](#)

- [pgr\\_maxFlowMinCost - Experimental](#) - Details of flow and cost on edges.
- [pgr\\_maxFlowMinCost\\_Cost - Experimental](#) - Only the Min Cost calculation.

#### [Chinese Postman Problem - Family of functions \(Experimental\)](#)

- [pgr\\_chinesePostman - Experimental](#)
- [pgr\\_chinesePostmanCost - Experimental](#)

#### [Coloring - Family of functions](#)

- [pgr\\_bipartite - Experimental](#) - Bipartite graph algorithm using a DFS-based coloring approach.
- [pgr\\_edgeColoring - Experimental](#) - Edge Coloring algorithm using Vizing's theorem.

#### [Contraction - Family of functions](#)

- [pgr\\_contractionHierarchies - Experimental](#)

#### [Transformation - Family of functions](#)

- [pgr\\_lineGraphFull - Experimental](#) - Transformation algorithm for generating a Line Graph out of each vertex in the input graph.

#### [Traversal - Family of functions](#)

- [pgr\\_breadthFirstSearch - Experimental](#) - Breath first search traversal of the graph.
- [pgr\\_binaryBreadthFirstSearch - Experimental](#) - Breath first search traversal of the graph.

#### [Components - Family of functions](#)

- [pgr\\_makeConnected - Experimental](#) - Details of edges to make graph connected.

#### [Ordering - Family of functions](#)

- [pgr\\_cuthillMcKeeOrdering - Experimental](#) - Return reverse Cuthill-McKee ordering of an undirected graph.
- [pgr\\_topologicalSort - Experimental](#) - Linear ordering of the vertices for directed acyclic graph.

#### [Metrics - Family of functions](#)

- [pgr\\_betweennessCentrality - Experimental](#) - Calculates relative betweenness centrality using Brandes Algorithm

#### [TRSP - Family of functions](#)

- [pgr\\_turnRestrictedPath - Experimental](#) - Routing with restrictions.

### Chinese Postman Problem - Family of functions (Experimental)

- [pgr\\_chinesePostman - Experimental](#)
- [pgr\\_chinesePostmanCost - Experimental](#)

#### [pgr\\_chinesePostman - Experimental](#)

`pgr_chinesePostman` — Calculates the shortest circuit path which contains every edge in a directed graph and starts and ends on the same vertex.

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.0.0
  - New experimental function.

#### Description

The main characteristics are:

- Process is done only on edges with **positive** costs.
- Running time:  $O(E * (E + V * \log V))$
- Graph must be connected.

- Returns EMPTY SET on a disconnected graph

Boost Graph Inside

Signatures

pgr\_chinesePostman([Edges SQL](#))

Returns set of (seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

```
SELECT * FROM pgr_chinesePostman(
'SELECT id, source, target, cost, reverse_cost
FROM edges WHERE id < 17');
seq | node | edge | cost | agg_cost
```

1	1	6	1	0
2	3	7	1	1
3	7	4	1	2
4	6	4	1	3
5	7	8	1	4
6	11	8	1	5
7	7	10	1	6
8	8	12	1	7
9	12	13	1	8
10	17	15	1	9
11	16	15	1	10
12	17	15	1	11
13	16	16	1	12
14	15	16	1	13
15	16	9	1	14
16	11	11	1	15
17	12	13	1	16
18	17	15	1	17
19	16	16	1	18
20	15	3	1	19
21	10	5	1	20
22	11	9	1	21
23	16	16	1	22
24	15	3	1	23
25	10	2	1	24
26	6	1	1	25
27	5	1	1	26
28	6	4	1	27
29	7	10	1	28
30	8	14	1	29
31	9	14	1	30
32	8	10	1	31
33	7	7	1	32
34	3	6	1	33
35	1	-1	0	34

(35 rows)

Parameters

Parameter	Type	Description
-----------	------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below.

Inner Queries

Edges SQL

An Edges SQL that represents a **directed** graph with the following columns

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
			Weight of the edge (target, source)
reverse_cost	ANY-NUMERICAL	-1	<ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, node, edge, cost, agg\_cost)

Column	Type	Description
seq	INT	Sequential value starting from 1
node	BIGINT	Identifier of the node in the path fromstart_vid to end_vid.

Column	Type	Description
edge	BIGINT	Identifier of the edge used to go from node to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_v to node.

See Also

- [Chinese Postman Problem - Family of functions \(Experimental\)](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_chinesePostmanCost - Experimental

pgr\_chinesePostmanCost — Calculates the minimum costs of a circuit path which contains every edge in a directed graph and starts and ends on the same vertex.

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.0.0
  - New experimental function.

Description

The main characteristics are:

- Process is done only on edges with **positive** costs.
- Running time:  $\mathcal{O}(E * (E + V * \log V))$
- Graph must be connected.
- Return value when the graph if disconnected

Boost Graph Inside

Signatures

pgr\_chinesePostmanCost([Edges SQL](#))  
RETURNS FLOAT

Example:

```
SELECT * FROM pgr_chinesePostmanCost(
  'SELECT id, source, target, cost, reverse_cost
  FROM edges WHERE id < 17');
pgr_chinesepostmancost
```

(1 row)

Parameters

Parameter	Type	Description
-----------	------	-------------

Parameter TypeDescription

Edges SQL TEXT Edges SQL as described below.

Inner Queries1

Edges SQL1

An Edges SQL that represents a **directed** graph with the following columns

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns1

Column	Type	Description
pgr_chinesepostmancost	FLOAT	Minimum costs of a circuit path.

See Also1

- Chinese Postman Problem - Family of functions (Experimental)
- Sample Data

Indices and tables

- Index
- Search Page

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Description1

The main characteristics are:

- Process is done only on edges with **positive** costs.

- Running time:  $\mathcal{O}(E * (E + V * \log V))$
- Graph must be connected.

Parameters¶

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below.

Inner Queries¶

Edges SQL¶

An Edges SQL that represents a **directed** graph with the following columns

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
			Weight of the edge (target, source)
reverse_cost	ANY-NUMERICAL	-1	<ul style="list-style-type: none"> <li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

See Also¶

Indices and tables

- [Index](#)
- [Search Page](#)

Transformation - Family of functions¶

☐ Proposed

Warning

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.
- [pgr\\_lineGraph - Proposed](#) - Transformation algorithm for generating a Line Graph.

☐ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.

- May lack documentation.
- Documentation if any might need to be rewritten.
- Documentation examples might need to be automatically generated.
- Might need a lot of feedback from the community.
- Might depend on a proposed function of pgRouting
- Might depend on a deprecated function of pgRouting
- [pgr\\_lineGraphFull - Experimental](#) - Transformation algorithm for generating a Line Graph out of each vertex in the input graph.

**pgr\_lineGraph - Proposed**

pgr\_lineGraph — Transforms the given graph into its corresponding edge-based graph.

☐ Proposed

**Warning**

Proposed functions for next mayor release.

- They are not officially in the current release.
- They will likely officially be part of the next mayor release:
  - The functions make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might not change. (But still can)
  - Signature might not change. (But still can)
  - Functionality might not change. (But still can)
  - pgTap tests have being done. But might need more.
  - Documentation might need refinement.

**Availability**

- Version 3.7.0
  - Function promoted to proposed.
  - Works for directed and undirected graphs.
- Version 2.5.0
  - New experimental function.

**Description**

Given a graph  $G$ , its line graph  $L(G)$  is a graph such that:

- Each vertex of  $L(G)$  represents an edge of  $G$ .
- Two vertices of  $L(G)$  are adjacent if and only if their corresponding edges share a common endpoint in  $G$ .

**The main characteristics are:**

- Works for directed and undirected graphs.
- The cost and reverse\_cost columns of the result represent existence of the edge.
- When the graph is directed the result is directed.
  - To get the complete Line Graph use unique identifiers on the double way edges (See [Additional Examples](#)).
- When the graph is undirected the result is undirected.
  - The reverse\_cost is always  $-1$ .

Boost Graph Inside

**Signatures**

pgr\_lineGraph([Edges SQL](#), [directed])  
Returns set of (seq, source, target, cost, reverse\_cost)  
OR EMPTY SET

**Example:**

For an undirected graph with edges  $\{2,4,5,8\}$

```
SELECT * FROM pgr_lineGraph(
'SELECT id, source, target, cost, reverse_cost
FROM edges WHERE id IN (2,4,5,8)',
false);
seq | source | target | cost | reverse_cost
-----+-----+-----+-----+-----
1 | 2 | 4 | 1 | -1
2 | 2 | 5 | 1 | -1
3 | 4 | 8 | 1 | -1
4 | 5 | 8 | 1 | -1
(4 rows)
```

**Parameters**

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below.

**Optional parameters**

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, source, target, cost, reverse\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1. <ul style="list-style-type: none"><li>Gives a local identifier for the edge</li></ul>
source	BIGINT	Identifier of the source vertex of the current edge. <ul style="list-style-type: none"><li>When <i>negative</i>: the source is the reverse edge in the original graph.</li></ul>
target	BIGINT	Identifier of the target vertex of the current edge. <ul style="list-style-type: none"><li>When <i>negative</i>: the target is the reverse edge in the original graph.</li></ul>
cost	FLOAT	Weight of the edge (source, target). <ul style="list-style-type: none"><li>When <i>negative</i>: edge (source, target) does not exist, therefore it's not part of the graph.</li></ul>
reverse_cost	FLOAT	Weight of the edge (target, source). <ul style="list-style-type: none"><li>When <i>negative</i>: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

- Additional Examples
- Representation as directed with shared edge identifiers
    - Line Graph of a directed graph represented with shared edges
  - Representation as directed with unique edge identifiers
    - Line Graph of a directed graph represented with unique edges

Given the following directed graph

$G(V,E) = G(\{1,2,3,4\},\{ 1 \rightarrow 2, 1 \rightarrow 4, 2 \rightarrow 3, 3 \rightarrow 1, 3 \rightarrow 2, 3 \rightarrow 4, 4 \rightarrow 3\})$

Representation as directed with shared edge identifiers

For the simplicity, the design of the edges table on the database, has the edge's identifiers are represented with 3 digits:

hundreds:

the source vertex

tens:

always 0, acts as a separator

units:

the target vertex



In this image,

- Single or double head arrows represent one edge (row) on the edges table.
- The numbers in the yellow shadow are the edge identifiers.

Two pair of edges share the same identifier when the reverse\_cost column is used.

- Edges  $\{(2 \rightarrow 3, 3 \rightarrow 2)\}$  are represented with one edge row with  $(id=203)$ .
- Edges  $\{(3 \rightarrow 4, 4 \rightarrow 3)\}$  are represented with one edge row with  $(id=304)$ .

The graph can be created as follows:

```
CREATE TABLE edges_shared (  
  id BIGINT,  
  source BIGINT,  
  target BIGINT,  
  cost FLOAT,  
  reverse_cost FLOAT,  
  geom geometry  
);  
CREATE TABLE  
INSERT INTO edges_shared (id, source, target, cost, reverse_cost, geom) VALUES  
(102, 1, 2, 1, -1, ST_MakeLine(ST_POINT(0, 2), ST_POINT(2, 2))),  
(104, 1, 4, 1, -1, ST_MakeLine(ST_POINT(0, 2), ST_POINT(0, 0))),  
(301, 3, 1, 1, -1, ST_MakeLine(ST_POINT(2, 0), ST_POINT(0, 2))),  
(203, 2, 3, 1, 1, ST_MakeLine(ST_POINT(2, 2), ST_POINT(2, 0))),  
(304, 3, 4, 1, 1, ST_MakeLine(ST_POINT(0, 0), ST_POINT(2, 0)));
```

Line Graph of a directed graph represented with shared edges

```
SELECT seq, source, target, cost, reverse_cost  
FROM pgr_lineGraph(  
  'SELECT id, source, target, cost, reverse_cost FROM edges_shared',  
  true);  
seq | source | target | cost | reverse_cost  
-----  
1 | 102 | 203 | 1 | -1  
2 | 104 | 304 | 1 | -1  
3 | 203 | 203 | 1 | 1  
4 | 203 | 301 | 1 | -1  
5 | 203 | 304 | 1 | 1  
6 | 301 | 102 | 1 | -1  
7 | 301 | 104 | 1 | -1  
8 | 304 | 301 | 1 | -1  
9 | 304 | 304 | 1 | 1  
(9 rows)
```

- The result is a directed graph.
- For  $(seq=4)$  from  $(203 \rightarrow 304)$  represent two edges
- For all the other values of seq represent one edge.
- The cost and reverse\_cost values represent the existence of the edge.
  - When positive: the edge exists.
  - When negative: the edge does not exist.

Representation as directed with unique edge identifiers

For the simplicity, the design of the edges table on the database, has the edge's identifiers are represented with 3 digits:  
hundreds:

the source vertex

tens:

always 0, acts as a separator

units:

the target vertex

In this image,

- Single head arrows represent one edge (row) on the edges table.
- There are no double head arrows
- The numbers in the yellow shadow are the edge identifiers.

Two pair of edges share the same ending nodes and the reverse\_cost column is not used.

- Edges  $\{(2 \rightarrow 3, 3 \rightarrow 2)\}$  are represented with two edges  $(id=203)$  and  $(id=302)$  respectively.
- Edges  $\{(3 \rightarrow 4, 4 \rightarrow 3)\}$  are represented with two edges  $(id=304)$  and  $(id=403)$  respectively.

The graph can be created as follows:

```
CREATE TABLE edges_unique (  
  id BIGINT,  
  source BIGINT,  
  target BIGINT,  
  cost FLOAT,  
  geom geometry  
);  
CREATE TABLE  
INSERT INTO edges_unique (id, source, target, cost, geom) VALUES  
(102, 1, 2, 1, ST_MakeLine(ST_POINT(0, 2), ST_POINT(2, 2))),  
(104, 1, 4, 1, ST_MakeLine(ST_POINT(0, 2), ST_POINT(0, 0))),  
(301, 3, 1, 1, ST_MakeLine(ST_POINT(2, 0), ST_POINT(0, 2))),  
(203, 2, 3, 1, ST_MakeLine(ST_POINT(2, 2), ST_POINT(2, 0))),  
(304, 3, 4, 1, ST_MakeLine(ST_POINT(2, 0), ST_POINT(0, 0))),  
(302, 3, 2, 1, ST_MakeLine(ST_POINT(2, 0), ST_POINT(2, 2))),  
(403, 4, 3, 1, ST_MakeLine(ST_POINT(0, 0), ST_POINT(2, 0)));
```

Line Graph of a directed graph represented with unique edges

```
SELECT seq, source, target, cost, reverse_cost  
FROM pgr_lineGraph(  
  'SELECT id, source, target, cost, reverse_cost FROM edges_unique',  
  true);
```

```
'SELECT id, source, target, cost FROM edges_unique',
true);
seq | source | target | cost | reverse_cost
-----+-----+-----+-----+-----
1 | 102 | 203 | 1 | -1
2 | 104 | 403 | 1 | -1
3 | 203 | 301 | 1 | -1
4 | 203 | 304 | 1 | -1
5 | 301 | 102 | 1 | -1
6 | 301 | 104 | 1 | -1
7 | 302 | 203 | 1 | 1
8 | 304 | 403 | 1 | 1
9 | 403 | 301 | 1 | -1
10 | 403 | 302 | 1 | -1
(10 rows)
```

- The result is a directed graph.
- For  $\backslash(seq=7)$  from  $\backslash(203 \rightarrow 302)$  represent two edges.
- For  $\backslash(seq=8)$  from  $\backslash(304 \rightarrow 403)$  represent two edges.
- For all the other values of seq represent one edge.
- The cost and reverse\_cost values represent the existence of the edge.
  - When positive: the edge exists.
  - When negative: the edge does not exist.

See Also

- wikipedia: [Line Graph](#)
- mathworld: [Line Graph](#)
- [Sample Data](#)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_lineGraphFull - Experimental

pgr\_lineGraphFull — Transforms a given graph into a new graph where all of the vertices from the original graph are converted to line graphs.

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 2.6.0
  - New experimental function.

Description

pgr\_lineGraphFull, converts original directed graph to a directed line graph by converting each vertex to a complete graph and keeping all the original edges. The new connecting edges have a cost 0 and go between the adjacent original edges, respecting the directionality.

A possible application of the resulting graph is “**routing with two edge restrictions**”:

- Setting a cost of using the vertex when routing between edges on the connecting edge
- Forbid the routing between two edges by removing the connecting edge

This is possible because each of the intersections (vertices) in the original graph are now complete graphs that have a new edge for each possible turn across that intersection.

The main characteristics are:

- This function is for **directed** graphs.
- Results are undefined when a negative vertex id is used in the input graph.
- Results are undefined when a duplicated edge id is used in the input graph.

- Running time: TBD

Boost Graph Inside

Signatures1

Summary

pgr\_lineGraphFull([Edges SQL](#))  
Returns set of (seq, source, target, cost, edge)  
OR EMPTY SET

Example:

Full line graph of subgraph of edges \(\{4, 7, 8, 10\}\)

```
SELECT * FROM pgr_lineGraphFull(  
  $$SELECT id, source, target, cost, reverse_cost  
  FROM edges  
  WHERE id IN (4, 7, 8, 10)$$);  
seq | source | target | cost | edge
```

1	-1	7	1	4
2	6	-1	0	0
3	-2	6	1	-4
4	-3	3	1	-7
5	-4	11	1	8
6	-5	8	1	10
7	7	-2	0	0
8	7	-3	0	0
9	7	-4	0	0
10	7	-5	0	0
11	-6	-2	0	0
12	-6	-3	0	0
13	-6	-4	0	0
14	-6	-5	0	0
15	-7	-2	0	0
16	-7	-3	0	0
17	-7	-4	0	0
18	-7	-5	0	0
19	-8	-2	0	0
20	-8	-3	0	0
21	-8	-4	0	0
22	-8	-5	0	0
23	-9	-6	1	7
24	3	-9	0	0
25	-10	-7	1	-8
26	11	-10	0	0
27	-11	-8	1	-10
28	8	-11	0	0

(28 rows)

Parameters1

Parameter	Type	Description
-----------	------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below.

Inner Queries1

Edges SQL1

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns1

Returns set of (seq, source, target, cost, edge)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1. <ul style="list-style-type: none"><li>Gives a local identifier for the edge</li></ul>
source	BIGINT	Identifier of the source vertex of the current edge. <ul style="list-style-type: none"><li>When <i>negative</i>: the source is the reverse edge in the original graph.</li></ul>

Column	Type	Description
target	BIGINT	Identifier of the target vertex of the current edge. <ul style="list-style-type: none"> <li>When <i>negative</i>: the target is the reverse edge in the original graph.</li> </ul>
cost	FLOAT	Weight of the edge (source, target). <ul style="list-style-type: none"> <li>When <i>negative</i>: edge (source, target) does not exist, therefore it's not part of the graph.</li> </ul>
reverse_cost	FLOAT	Weight of the edge (target, source). <ul style="list-style-type: none"> <li>When <i>negative</i>: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

#### Additional Examples¶

- [The data](#)
- [The transformation¶](#)
- [Creating table that identifies transformed vertices](#)
  - [Store edge results](#)
  - [Create the mapping table](#)
  - [Filling the mapping table](#)
- [Adding a soft restriction](#)
  - [Identifying the restriction](#)
  - [Adding a value to the restriction](#)
- [Simplifying leaf vertices](#)
  - [Using the vertex map give the leaf verices their original value.](#)
  - [Removing self loops on leaf nodes](#)
- [Complete routing graph](#)
  - [Add edges from the original graph](#)
  - [Add the newly calculated edges](#)
- [Using the routing graph](#)
- The examples of this section are based on the [Sample Data](#) network.
- The examples include the subgraph including edges 4, 7, 8, and 10 with `reverse_cost`.

#### [The data¶](#)

This example displays how this graph transformation works to create additional edges for each possible turn in a graph.

```
SELECT id, source, target, cost, reverse_cost
FROM edges
WHERE id IN (4, 7, 8, 10);
id | source | target | cost | reverse_cost
-----+-----+-----+-----+-----
4 | 6 | 7 | 1 | 1
7 | 3 | 7 | 1 | -4
8 | 7 | 11 | 1 | 8
10 | 7 | 8 | 1 | 10
(4 rows)
```



#### [The transformation¶](#)

```
SELECT * FROM pgr_lineGraphFull(
  $$SELECT id, source, target, cost, reverse_cost
  FROM edges
  WHERE id IN (4, 7, 8, 10)$$);
seq | source | target | cost | edge
-----+-----+-----+-----+-----
1 | -1 | 7 | 1 | 4
2 | 6 | -1 | 0 | 0
3 | -2 | 6 | 1 | -4
4 | -3 | 3 | 1 | -7
5 | -4 | 11 | 1 | 8
6 | -5 | 8 | 1 | 10
7 | 7 | -2 | 0 | 0
8 | 7 | -3 | 0 | 0
9 | 7 | -4 | 0 | 0
10 | 7 | -5 | 0 | 0
11 | -6 | -2 | 0 | 0
12 | -6 | -3 | 0 | 0
13 | -6 | -4 | 0 | 0
14 | -6 | -5 | 0 | 0
15 | -7 | -2 | 0 | 0
16 | -7 | -3 | 0 | 0
17 | -7 | -4 | 0 | 0
18 | -7 | -5 | 0 | 0
19 | -8 | -2 | 0 | 0
20 | -8 | -3 | 0 | 0
21 | -8 | -4 | 0 | 0
22 | -8 | -5 | 0 | 0
23 | -9 | -6 | 1 | 7
24 | 3 | -9 | 0 | 0
25 | -10 | -7 | 1 | -8
26 | 11 | -10 | 0 | 0
27 | -11 | -8 | 1 | -10
28 | 8 | -11 | 0 | 0
(28 rows)
```



In the transformed graph, all of the edges from the original graph are still present (yellow), but we now have additional edges for every turn that could be made across vertex 7 (orange).

Creating table that identifies transformed vertices¶

The vertices in the transformed graph are each created by splitting up the vertices in the original graph. Unless a vertex in the original graph is a leaf vertex, it will generate more than one vertex in the transformed graph. One of the newly created vertices in the transformed graph will be given the same vertex identifier as the vertex that it was created from in the original graph, but the rest of the newly created vertices will have negative vertex ids.

Following is an example of how to generate a table that maps the ids of the newly created vertices with the original vertex that they were created from

Store edge results¶

The first step is to store the results of the pgr\_lineGraphFull call into a table

```
SELECT seq AS id, source, target, cost, edge
INTO lineGraph_edges
FROM pgr_lineGraphFull(
  $$SELECT id, source, target, cost, reverse_cost
  FROM edges
  WHERE id IN (4, 7, 8, 10)$$);
SELECT 28
```

Create the mapping table¶

From the original graph's vertex information

```
SELECT id, NULL::BIGINT original_id
INTO vertex_map
FROM vertices;
SELECT 17
```

Add the new vertices

```
INSERT INTO vertex_map (id)
(SELECT id
FROM pgr_extractVertices(
  $$SELECT id, source, target FROM lineGraph_edges$$) WHERE id < 0);
INSERT 0 11
```

Filling the mapping table¶

The positive vertex identifiers are the original identifiers

```
UPDATE vertex_map
SET original_id = id
WHERE id > 0;
UPDATE 17
```

Inspecting the vertices map

```
SELECT *
FROM vertex_map ORDER BY id DESC;
id | original_id
----+-----
17 | 17
16 | 16
15 | 15
14 | 14
13 | 13
12 | 12
11 | 11
10 | 10
9 | 9
8 | 8
7 | 7
6 | 6
5 | 5
4 | 4
3 | 3
2 | 2
1 | 1
-1 |
-2 |
-3 |
-4 |
-5 |
-6 |
-7 |
-8 |
-9 |
-10 |
-11 |
(28 rows)
```

The self loops happen when there is no cost traveling to the target and the source has an original value.

```
SELECT *, source AS targets_original_id
FROM lineGraph_edges
WHERE cost = 0 and source > 0;
id | source | target | cost | edge | targets_original_id
----+-----+-----+-----+-----+-----
2 | 6 | -1 | 0 | 0 | 6
7 | 7 | -2 | 0 | 0 | 7
8 | 7 | -3 | 0 | 0 | 7
9 | 7 | -4 | 0 | 0 | 7
10 | 7 | -5 | 0 | 0 | 7
24 | 3 | -9 | 0 | 0 | 3
26 | 11 | -10 | 0 | 0 | 11
28 | 8 | -11 | 0 | 0 | 8
(8 rows)
```

Updating values from self loops

```
WITH
self_loops AS (
  SELECT DISTINCT source, target, source AS targets_original_id
  FROM lineGraph_edges
  WHERE cost = 0 and source > 0)
UPDATE vertex_map SET original_id = targets_original_id
FROM self_loops WHERE target = id;
UPDATE 8
```

Inspecting the vertices table

```
SELECT *
FROM vertex_map WHERE id < 0
ORDER BY id DESC;
id | original_id
----+-----
-1 | 6
-2 | 7
-3 | 7
-4 | 7
-5 | 7
-6 |
```

-7		
-8		
-9		3
-10		11
-11		8

(11 rows)

Updating from inner self loops

```

WITH
assigned_vertices
AS (SELECT id, original_id
  FROM vertex_map
  WHERE original_id IS NOT NULL),
cross_edges
AS (SELECT DISTINCT e.source, v.original_id AS source_original_id
  FROM lineGraph_edges AS e
  JOIN vertex_map AS v ON (e.target = v.id)
  WHERE source NOT IN (SELECT id FROM assigned_vertices)
)
UPDATE vertex_map SET original_id = source_original_id
FROM cross_edges WHERE source = id;
UPDATE 3

```

Inspecting the vertices map

```

SELECT *
FROM vertex_map WHERE id < 0
ORDER BY id DESC;
id | original_id
-----+-----
-1 | 6
-2 | 7
-3 | 7
-4 | 7
-5 | 7
-6 | 7
-7 | 7
-8 | 7
-9 | 3
-10 | 11
-11 | 8
(11 rows)

```

[Adding a soft restriction¶](#)

A soft restriction going from vertex 6 to vertex 3 using edges 4 -> 7 is wanted.

[Identifying the restriction¶](#)

Running a [pgr\\_dijkstraNear - Proposed](#) the edge with cost 0, edge 8, is where the cost will be increased

```

SELECT seq, path_seq, start_vid, end_vid, node, original_id, edge, cost, agg_cost
FROM (SELECT * FROM pgr_dijkstraNear(
  $$SELECT * FROM lineGraph_edges$$,
  (SELECT array_agg(id) FROM vertex_map where original_id = 6),
  (SELECT array_agg(id) FROM vertex_map where original_id = 3))) dn
JOIN vertex_map AS v1 ON (node = v1.id);
seq | path_seq | start_vid | end_vid | node | original_id | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
3 | 3 | -1 | 3 | -3 | 7 | 4 | 1 | 1
1 | 1 | -1 | 3 | -1 | 6 | 1 | 1 | 0
4 | 4 | -1 | 3 | 3 | 3 | -1 | 0 | 2
2 | 2 | -1 | 3 | 7 | 7 | 8 | 0 | 1
(4 rows)

```

The edge to be altered is WHERE cost = 0 AND seq != 1 AND edge != -1 from the previous query:

```

SELECT edge FROM pgr_dijkstraNear(
  $$SELECT * FROM lineGraph_edges$$,
  (SELECT array_agg(id) FROM vertex_map where original_id = 6),
  (SELECT array_agg(id) FROM vertex_map where original_id = 3))
WHERE cost = 0 AND seq != 1 AND edge != -1;
edge
-----
8
(1 row)

```

[Adding a value to the restriction¶](#)

Updating the cost to the edge:

```

UPDATE lineGraph_edges
SET cost = 100
WHERE id IN (
  SELECT edge FROM pgr_dijkstraNear(
  $$SELECT * FROM lineGraph_edges$$,
  (SELECT array_agg(id) FROM vertex_map where original_id = 6),
  (SELECT array_agg(id) FROM vertex_map where original_id = 3))
WHERE cost = 0 AND seq != 1 AND edge != -1);
UPDATE 1

```

Example:

Routing from \6\ to \3\

Now the route does not use edge 8 and does a U turn on a leaf vertex.

```

WITH
results AS (
  SELECT * FROM pgr_dijkstraNear(
  $$SELECT * FROM lineGraph_edges$$,
  (SELECT array_agg(id) FROM vertex_map where original_id = 6),
  (SELECT array_agg(id) FROM vertex_map where original_id = 3))
  SELECT seq, path_seq, start_vid, end_vid, node, original_id, edge, cost, agg_cost
FROM results
LEFT JOIN vertex_map AS v1 ON (node = v1.id) ORDER BY seq;
seq | path_seq | start_vid | end_vid | node | original_id | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | -1 | 3 | -1 | 6 | 1 | 1 | 0
2 | 2 | -1 | 3 | 7 | 7 | 10 | 0 | 1
3 | 3 | -1 | 3 | -5 | 7 | 6 | 1 | 1
4 | 4 | -1 | 3 | 8 | 8 | 28 | 0 | 2
5 | 5 | -1 | 3 | -11 | 8 | 27 | 1 | 2
6 | 6 | -1 | 3 | -8 | 7 | 20 | 0 | 3
7 | 7 | -1 | 3 | -3 | 7 | 4 | 1 | 3
8 | 8 | -1 | 3 | 3 | 3 | -1 | 0 | 4
(8 rows)

```

[Simplifying leaf vertices¶](#)

In this example, there is no additional cost for traversing a leaf vertex.

[Using the vertex map give the leaf verices their original value.¶](#)

On the source column

```
WITH
u_turns AS (
SELECT e.id AS eid, v1.original_id
FROM linegraph_edges as e
JOIN vertex_map AS v1 ON (source = v1.id)
AND v1.original_id IN (3, 6, 8, 11))
UPDATE lineGraph_edges
SET source = original_id
FROM u_turns
WHERE id = eid;
UPDATE 8
```

On the target column

```
WITH
u_turns AS (
SELECT e.id AS eid, v1.original_id
FROM linegraph_edges as e
JOIN vertex_map AS v1 ON (target = v1.id)
AND v1.original_id IN (3, 6, 8, 11))
UPDATE lineGraph_edges
SET target = original_id
FROM u_turns
WHERE id = eid;
UPDATE 8
```

[Removing self loops on leaf nodes¶](#)

The self loops of the leaf nodes are

```
SELECT * FROM linegraph_edges
WHERE source = target
ORDER BY id;
id | source | target | cost | edge
-----+-----+-----+-----+-----
2 | 6 | 6 | 0 | 0
24 | 3 | 3 | 0 | 0
26 | 11 | 11 | 0 | 0
28 | 8 | 8 | 0 | 0
(4 rows)
```

Which can be removed

```
DELETE FROM linegraph_edges
WHERE source = target;
DELETE 4
```

Example:

Routing from \6\ to \3\

Routing can be done now using the original vertices id using [pgr\\_dijkstra](#)

```
WITH
results AS (
SELECT * FROM pgr_dijkstra(
$$SELECT * FROM lineGraph_edges$$, 6, 3))
SELECT seq, path_seq, node, original_id, edge, cost, agg_cost
FROM results
LEFT JOIN vertex_map AS v1 ON (node = v1.id) ORDER BY seq;
seq | path_seq | node | original_id | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 6 | 1 | 1 | 0
2 | 2 | 7 | 7 | 9 | 0 | 1
3 | 3 | -4 | 7 | 5 | 1 | 1
4 | 4 | 11 | 11 | 25 | 1 | 2
5 | 5 | -7 | 7 | 16 | 0 | 3
6 | 6 | -3 | 7 | 4 | 1 | 3
7 | 7 | 3 | 3 | -1 | 0 | 4
(7 rows)
```

[Complete routing graph¶](#)

[Add edges from the original graph¶](#)

Add all the edges that are not involved in the line graph process to the new table

```
SELECT id, source, target, cost, reverse_cost
INTO new_graph from edges
WHERE id NOT IN (4, 7, 8, 10);
SELECT 14
```

Some administrative tasks to get new identifiers for the edges

```
CREATE SEQUENCE new_graph_id_seq;
CREATE SEQUENCE
ALTER TABLE new_graph ALTER COLUMN id SET DEFAULT nextval('new_graph_id_seq');
ALTER TABLE
ALTER TABLE new_graph ALTER COLUMN id SET NOT NULL;
ALTER TABLE
ALTER SEQUENCE new_graph_id_seq OWNED BY new_graph.id;
ALTER SEQUENCE
SELECT setval('new_graph_id_seq', (SELECT max(id) FROM new_graph));
setval
-----
18
(1 row)
```

[Add the newly calculated edges¶](#)

```
INSERT INTO new_graph (source, target, cost, reverse_cost)
SELECT source, target, cost, -1 FROM lineGraph_edges;
INSERT 0 24
```

[Using the routing graph¶](#)

When using this method for routing with soft restrictions there will be uturns

Example:

Routing from \6\ to \3\

```
WITH
results AS (
SELECT * FROM pgr_dijkstra(
```

```
$$SELECT * FROM new_graph$$, 6, 3))
SELECT seq, path_seq, node, original_id, edge, cost, agg_cost
FROM results
LEFT JOIN vertex_map AS v1 ON (node = v1.id) ORDER BY seq;
seq | path_seq | node | original_id | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 6 | 35 | 1 | 0
2 | 2 | 7 | 7 | 20 | 0 | 1
3 | 3 | -4 | 7 | 41 | 1 | 1
4 | 4 | 11 | 11 | 37 | 1 | 2
5 | 5 | -7 | 7 | 27 | 0 | 3
6 | 6 | -3 | 7 | 40 | 1 | 3
7 | 7 | 3 | 3 | -1 | 0 | 4
(7 rows)
```

Example:

Routing from \5\ to \1\

```
WITH
results AS (
  SELECT * FROM pgr_dijkstra(
    $$SELECT * FROM new_graph$$, 5, 1)
  SELECT seq, path_seq, node, original_id, edge, cost, agg_cost
  FROM results
  LEFT JOIN vertex_map AS v1 ON (node = v1.id) ORDER BY seq;
  seq | path_seq | node | original_id | edge | cost | agg_cost
  -----+-----+-----+-----+-----+-----+-----
  1 | 1 | 5 | 5 | 1 | 1 | 0
  2 | 2 | 6 | 6 | 35 | 1 | 1
  3 | 3 | 7 | 7 | 20 | 0 | 2
  4 | 4 | -4 | 7 | 41 | 1 | 2
  5 | 5 | 11 | 11 | 37 | 1 | 3
  6 | 6 | -7 | 7 | 27 | 0 | 4
  7 | 7 | -3 | 7 | 40 | 1 | 4
  8 | 8 | 3 | 3 | 6 | 1 | 5
  9 | 9 | 1 | 1 | -1 | 0 | 6
  (9 rows)
```

- See Also
- [https://en.wikipedia.org/wiki/Line\\_graph](https://en.wikipedia.org/wiki/Line_graph)
  - [https://en.wikipedia.org/wiki/Complete\\_graph](https://en.wikipedia.org/wiki/Complete_graph)

Indices and tables

- [Index](#)
- [Search Page](#)

Introduction

This family of functions is used for transforming a given input graph\((G(V,E))\) into a new graph\((G'(V',E'))\).

See Also

Indices and tables

- [Index](#)
- [Search Page](#)

Ordering - Family of functions

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting
- [pgr\\_cuthillMckeeOrdering - Experimental](#) - Return reverse Cuthill-McKee ordering of an undirected graph.
- [pgr\\_topologicalSort - Experimental](#) - Linear ordering of the vertices for directed acyclic graph.

pgr\_cuthillMckeeOrdering - Experimental

pgr\_cuthillMckeeOrdering — Returns the reverse Cuthill-McKee ordering of an undirected graphs



Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.4.0
  - New experimental function.

Description

In numerical linear algebra, the Cuthill-McKee algorithm (CM), named after Elizabeth Cuthill and James McKee, is an algorithm to permute a sparse matrix that has a symmetric sparsity pattern into a band matrix form with a small bandwidth.

The vertices are basically assigned a breadth-first search order, except that at each step, the adjacent vertices are placed in the queue in order of increasing degree.

The main Characteristics are:

- The implementation is for **undirected** graphs.
- The bandwidth minimization problems are considered NP-complete problems.
- The running time complexity is:  $O(m \log(m) \sqrt{V})$ 
  - where  $\sqrt{V}$  is the number of vertices,
  - $\sqrt{m}$  is the maximum degree of the vertices in the graph.

Boost Graph Inside

Signatures

pgr\_cuthillMcKeeOrdering([Edges SQL](#))

Returns set of (seq, node)

OR EMPTY SET

Example:

Graph ordering of pgRouting [Sample Data](#)

```
SELECT * FROM pgr_cuthillMcKeeOrdering(
  'SELECT id, source, target, cost, reverse_cost FROM edges'
);
seq | node
-----+-----
 1 | 13
 2 | 14
 3 |  2
 4 |  4
 5 |  1
 6 |  9
 7 |  3
 8 |  8
 9 |  5
10 |  7
11 | 12
12 |  6
13 | 11
14 | 17
15 | 10
16 | 16
17 | 15
(17 rows)
```

Parameters

Parameter Type	Description
<a href="#">Edges SQL</a> TEXT	<a href="#">Edges SQL</a> as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, node)

Column	Type	Description
seq	BIGINT	Sequence of the order starting from 1.
node	BIGINT	New ordering in reverse order.

See Also

- Sample Data
- Boost: Cuthill-McKee Ordering
- Wikipedia: Cuthill-McKee Ordering

Indices and tables

- Index
- Search Page

pg\_topologicalSort - Experimental

pg\_topologicalSort — Linear ordering of the vertices for directed acyclic graphs (DAG).

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.0.0
  - New experimental function.

Description

The topological sort algorithm creates a linear ordering of the vertices such that if  $edge((u,v))$  appears in the graph, then  $v$  comes before  $u$  in the ordering.

The main characteristics are:

- Process is valid for directed acyclic graphs only. otherwise it will throw warnings.
- For optimization purposes, if there are more than one answer, the function will return one of them.
- The returned values are ordered in topological order:
- Running time:  $O(V + E)$

Boost Graph Inside

Signatures

Summary

`pgr_topologicalSort`([Edges SQL](#))  
Returns set of (seq, sorted\_v)  
OR EMPTY SET

Example:

Topologically sorting the graph

```
SELECT * FROM pgr_topologicalsort(
  $$SELECT id, source, target, cost
  FROM edges WHERE cost >= 0
  UNION
  SELECT id, target, source, reverse_cost
  FROM edges WHERE cost < 0$$);
seq | sorted_v
```

1	1
2	5
3	2
4	4
5	3
6	13
7	14
8	15
9	10
10	6
11	7
12	8
13	9
14	11
15	16
16	12
17	17
(17 rows)	

Parameters

Parameter	Type	Description
-----------	------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>• When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, sorted\_v)

Column	Type	Description
seq	INTEGER	Sequential value starting from $1$
sorted_v	BIGINT	Linear topological ordering of the vertices

Additional examples

Example:

Topologically sorting the one way segments

```
SELECT * FROM pgr_topologicalsort(
  $$SELECT id, source, target, cost, -1 AS reverse_cost
  FROM edges WHERE cost >= 0
  UNION
  SELECT id, source, target, -1, reverse_cost
  FROM edges WHERE cost < 0$$);
seq | sorted_v
```

1	5
2	2
3	4
4	13
5	14
6	1
7	3
8	15
9	10
10	6
11	7
12	8
13	9
14	11
15	12
16	16
17	17

(17 rows)

Example:

Graph is not a DAG

```
SELECT * FROM pgr_topologicalsort(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$);
ERROR:  Graph is not DAG
CONTEXT:  SQL function "pgr_topologicalsort" statement 1
```

See Also

- [Sample Data](#)
- [Boost: topological sort](#)
- [https://en.wikipedia.org/wiki/Topological\\_sorting](https://en.wikipedia.org/wiki/Topological_sorting)

Indices and tables

- [Index](#)
- [Search Page](#)

See Also

Indices and tables

- [Index](#)
- [Search Page](#)

categories

Vehicle Routing Functions - Category

- Pickup and delivery problem
  - [pgr\\_pickDeliver - Experimental](#) - Pickup & Delivery using a Cost Matrix
  - [pgr\\_pickDeliverEuclidean - Experimental](#) - Pickup & Delivery with Euclidean distances
- Distribution problem
  - [pgr\\_vrpOneDepot - Experimental](#) - From a single depot, distributes orders

Shortest Path Category

- [pgr\\_bellmanFord - Experimental](#)
- [pgr\\_dagShortestPath - Experimental](#)
- [pgr\\_edwardMoore - Experimental](#)

pgr\_bellmanFord - Experimental

pgr\_bellmanFord — Shortest path using Bellman-Ford algorithm.

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.

- May lack documentation.
- Documentation if any might need to be rewritten.
- Documentation examples might need to be automatically generated.
- Might need a lot of feedback from the community.
- Might depend on a proposed function of pgRouting
- Might depend on a deprecated function of pgRouting

Availability

- Version 3.2.0
  - New experimental signature:
    - pgr\_bellmanFord(Combinations)
- Version 3.0.0
  - New experimental function.

Description

Bellman-Ford's algorithm, is named after Richard Bellman and Lester Ford, who first published it in 1958 and 1956, respectively. It is a graph search algorithm that computes shortest paths from a starting vertex (start\_vid) to an ending vertex (end\_vid) in a graph where some of the edge weights may be negative. Though it is more versatile, it is slower than Dijkstra's algorithm. This implementation can be used with a directed graph and an undirected graph.

The main characteristics are:

- Process is valid for edges with both positive and negative edge weights.
- Values are returned when there is a path.
  - When the start vertex and the end vertex are the same, there is no path. The agg\_cost would be 0.
  - When the start vertex and the end vertex are different, and there exists a path between them without having a negative cycle. The agg\_cost would be some finite value denoting the shortest distance between them.
  - When the start vertex and the end vertex are different, and there exists a path between them, but it contains a negative cycle. In such case, agg\_cost for those vertices keep on decreasing furthermore, Hence agg\_cost can't be defined for them.
  - When the start vertex and the end vertex are different, and there is no path. The agg\_cost is infinity.
- For optimization purposes, any duplicated value in the start\_vids or end\_vids are ignored.
- The returned values are ordered:
  - start\_vid ascending
  - end\_vid ascending
- Running time:  $O(|start\_vids| * (V * E))$

 Boost Graph Inside

Signatures

Summary

pgr\_bellmanFord(Edges SQL, start\_vid, end\_vid, [directed])  
pgr\_bellmanFord(Edges SQL, start\_vid, end\_vids, [directed])  
pgr\_bellmanFord(Edges SQL, start\_vids, end\_vid, [directed])  
pgr\_bellmanFord(Edges SQL, start\_vids, end\_vids, [directed])  
pgr\_bellmanFord(Edges SQL, Combinations SQL, [directed])  
Returns set of (seq, path\_seq, [start\_vid], [end\_vid], node, edge, cost, agg\_cost)  
OR EMPTY SET

One to One

pgr\_bellmanFord(Edges SQL, start\_vid, end\_vid, [directed])  
Returns set of (seq, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex 6 to vertex 10 on a directed graph

```
SELECT * FROM pgr_bellmanFord(
'SELECT id, source, target, cost, reverse_cost FROM edges',
6, 10, true);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 4 | 1 | 0
2 | 2 | 7 | 8 | 1 | 1
3 | 3 | 11 | 9 | 1 | 2
4 | 4 | 16 | 16 | 1 | 3
5 | 5 | 15 | 3 | 1 | 4
6 | 6 | 10 | -1 | 0 | 5
(6 rows)
```

One to Many

pgr\_bellmanFord(Edges SQL, start\_vid, end\_vids, [directed])  
Returns set of (seq, path\_seq, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex 6 to vertices (10, 17) on a directed graph

```
SELECT * FROM pgr_bellmanFord(
'SELECT id, source, target, cost, reverse_cost FROM edges',
6, ARRAY[10, 17]);
seq | path_seq | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 10 | 6 | 4 | 1 | 0
2 | 2 | 10 | 7 | 8 | 1 | 1
3 | 3 | 10 | 11 | 9 | 1 | 2
4 | 4 | 10 | 16 | 16 | 1 | 3
5 | 5 | 10 | 15 | 3 | 1 | 4
6 | 6 | 10 | 10 | -1 | 0 | 5
7 | 1 | 17 | 6 | 4 | 1 | 0
8 | 2 | 17 | 7 | 8 | 1 | 1
```

9	3	17	11	11	1	2
10	4	17	12	13	1	3
11	5	17	17	-1	0	4

(11 rows)

Many to One

pgr\_bellmanFord([Edges SQL](#), start vids, end vid, [directed])  
Returns set of (seq, path\_seq, start\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \(\{6, 1\}\) to vertex \(\{17\}\) on a **directed** graph

```
SELECT * FROM pgr_bellmanFord(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[6, 1], 17);
seq | path_seq | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 1 | 6 | 1 | 0
2 | 2 | 1 | 3 | 7 | 1 | 1
3 | 3 | 1 | 7 | 8 | 1 | 2
4 | 4 | 1 | 11 | 11 | 1 | 3
5 | 5 | 1 | 12 | 13 | 1 | 4
6 | 6 | 1 | 17 | -1 | 0 | 5
7 | 1 | 6 | 6 | 4 | 1 | 0
8 | 2 | 6 | 7 | 8 | 1 | 1
9 | 3 | 6 | 11 | 11 | 1 | 2
10 | 4 | 6 | 12 | 13 | 1 | 3
11 | 5 | 6 | 17 | -1 | 0 | 4
(11 rows)
```

Many to Many

pgr\_bellmanFord([Edges SQL](#), start vids, end vids, [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \(\{6, 1\}\) to vertices \(\{10, 17\}\) on an **undirected** graph

```
SELECT * FROM pgr_bellmanFord(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[6, 1], ARRAY[10, 17],
directed => false);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 10 | 6 | 1 | 0
2 | 2 | 1 | 10 | 3 | 7 | 1 | 1
3 | 3 | 1 | 10 | 7 | 4 | 1 | 2
4 | 4 | 1 | 10 | 6 | 2 | 1 | 3
5 | 5 | 1 | 10 | 10 | -1 | 0 | 4
6 | 1 | 1 | 17 | 1 | 6 | 1 | 0
7 | 2 | 1 | 17 | 3 | 7 | 1 | 1
8 | 3 | 1 | 17 | 7 | 8 | 1 | 2
9 | 4 | 1 | 17 | 11 | 11 | 1 | 3
10 | 5 | 1 | 17 | 12 | 13 | 1 | 4
11 | 6 | 1 | 17 | 17 | -1 | 0 | 5
12 | 1 | 6 | 10 | 6 | 2 | 1 | 0
13 | 2 | 6 | 10 | 10 | -1 | 0 | 1
14 | 1 | 6 | 17 | 6 | 4 | 1 | 0
15 | 2 | 6 | 17 | 7 | 8 | 1 | 1
16 | 3 | 6 | 17 | 11 | 11 | 1 | 2
17 | 4 | 6 | 17 | 12 | 13 | 1 | 3
18 | 5 | 6 | 17 | 17 | -1 | 0 | 4
(18 rows)
```

Combinations

pgr\_bellmanFord([Edges SQL](#), [Combinations SQL](#), [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table on an **undirected** graph.

The combinations table:

```
SELECT source, target FROM combinations;
source | target
-----+-----
5 | 6
5 | 10
6 | 5
6 | 15
6 | 14
(5 rows)
```

The query:

```
SELECT * FROM pgr_bellmanFord(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT source, target FROM combinations',
false);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 5 | 6 | 5 | 1 | 1 | 0
2 | 2 | 5 | 6 | 6 | -1 | 0 | 1
3 | 1 | 5 | 10 | 5 | 1 | 1 | 0
4 | 2 | 5 | 10 | 6 | 2 | 1 | 1
5 | 3 | 5 | 10 | 10 | -1 | 0 | 2
6 | 1 | 6 | 5 | 6 | 1 | 1 | 0
7 | 2 | 6 | 5 | 5 | -1 | 0 | 1
8 | 1 | 6 | 15 | 6 | 2 | 1 | 0
9 | 2 | 6 | 15 | 10 | 3 | 1 | 1
10 | 3 | 6 | 15 | 15 | -1 | 0 | 2
(10 rows)
```

Parameters

Column	Type	Description
--------	------	-------------

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
<b>start vid</b>	BIGINT	Identifier of the starting vertex of the path.
<b>start vids</b>	ARRAY[BIGINT]	Array of identifiers of starting vertices.
<b>end vid</b>	BIGINT	Identifier of the ending vertex of the path.
<b>end vids</b>	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters¶

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"> <li>When true the graph is considered <i>Directed</i></li> <li>When false the graph is considered as <i>Undirected</i>.</li> </ul>

Inner Queries¶

Edges SQL¶

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL¶

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns¶

Returns set of (seq, path\_seq [, start\_vid] [, end\_vid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex. Returned when multiple starting vetrices are in the query. <ul style="list-style-type: none"> <li><a href="#">Many to One</a></li> <li><a href="#">Many to Many</a></li> </ul>

Column	Type	Description
		Identifier of the ending vertex. Returned when multiple ending vertices are in the query.
end_vid	BIGINT	<ul style="list-style-type: none"> <li>• <a href="#">One to Many</a></li> <li>• <a href="#">Many to Many</a></li> </ul>
node	BIGINT	Identifier of the node in the path from start_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go from node to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Additional Examples1

Example 1:

Demonstration of repeated values are ignored, and result is sorted.

```
SELECT * FROM pgr_bellmanFord(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[7, 10, 15, 10, 10, 15], ARRAY[10, 7, 10, 15]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 7 | 10 | 7 | 8 | 1 | 0
2 | 2 | 7 | 10 | 11 | 9 | 1 | 1
3 | 3 | 7 | 10 | 16 | 16 | 1 | 2
4 | 4 | 7 | 10 | 15 | 3 | 1 | 3
5 | 5 | 7 | 10 | 10 | -1 | 0 | 4
6 | 1 | 7 | 15 | 7 | 8 | 1 | 0
7 | 2 | 7 | 15 | 11 | 9 | 1 | 1
8 | 3 | 7 | 15 | 16 | 16 | 1 | 2
9 | 4 | 7 | 15 | 15 | -1 | 0 | 3
10 | 1 | 10 | 7 | 10 | 5 | 1 | 0
11 | 2 | 10 | 7 | 11 | 8 | 1 | 1
12 | 3 | 10 | 7 | 7 | -1 | 0 | 2
13 | 1 | 10 | 15 | 10 | 5 | 1 | 0
14 | 2 | 10 | 15 | 11 | 9 | 1 | 1
15 | 3 | 10 | 15 | 16 | 16 | 1 | 2
16 | 4 | 10 | 15 | 15 | -1 | 0 | 3
17 | 1 | 15 | 7 | 15 | 3 | 1 | 0
18 | 2 | 15 | 7 | 10 | 2 | 1 | 1
19 | 3 | 15 | 7 | 6 | 4 | 1 | 2
20 | 4 | 15 | 7 | 7 | -1 | 0 | 3
21 | 1 | 15 | 10 | 15 | 3 | 1 | 0
22 | 2 | 15 | 10 | 10 | -1 | 0 | 1
(22 rows)
```

Example 2:

Making start vids the same as end vids.

```
SELECT * FROM pgr_bellmanFord(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[7, 10, 15], ARRAY[7, 10, 15]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 7 | 10 | 7 | 8 | 1 | 0
2 | 2 | 7 | 10 | 11 | 9 | 1 | 1
3 | 3 | 7 | 10 | 16 | 16 | 1 | 2
4 | 4 | 7 | 10 | 15 | 3 | 1 | 3
5 | 5 | 7 | 10 | 10 | -1 | 0 | 4
6 | 1 | 7 | 15 | 7 | 8 | 1 | 0
7 | 2 | 7 | 15 | 11 | 9 | 1 | 1
8 | 3 | 7 | 15 | 16 | 16 | 1 | 2
9 | 4 | 7 | 15 | 15 | -1 | 0 | 3
10 | 1 | 10 | 7 | 10 | 5 | 1 | 0
11 | 2 | 10 | 7 | 11 | 8 | 1 | 1
12 | 3 | 10 | 7 | 7 | -1 | 0 | 2
13 | 1 | 10 | 15 | 10 | 5 | 1 | 0
14 | 2 | 10 | 15 | 11 | 9 | 1 | 1
15 | 3 | 10 | 15 | 16 | 16 | 1 | 2
16 | 4 | 10 | 15 | 15 | -1 | 0 | 3
17 | 1 | 15 | 7 | 15 | 3 | 1 | 0
18 | 2 | 15 | 7 | 10 | 2 | 1 | 1
19 | 3 | 15 | 7 | 6 | 4 | 1 | 2
20 | 4 | 15 | 7 | 7 | -1 | 0 | 3
21 | 1 | 15 | 10 | 15 | 3 | 1 | 0
22 | 2 | 15 | 10 | 10 | -1 | 0 | 1
(22 rows)
```

Example 3:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_bellmanFord(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)');
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 7 | 6 | 4 | 1 | 0
2 | 2 | 6 | 7 | 7 | -1 | 0 | 1
3 | 1 | 6 | 10 | 6 | 4 | 1 | 0
4 | 2 | 6 | 10 | 7 | 8 | 1 | 1
5 | 3 | 6 | 10 | 11 | 9 | 1 | 2
6 | 4 | 6 | 10 | 16 | 16 | 1 | 3
7 | 5 | 6 | 10 | 15 | 3 | 1 | 4
8 | 6 | 6 | 10 | 10 | -1 | 0 | 5
9 | 1 | 12 | 10 | 12 | 13 | 1 | 0
10 | 2 | 12 | 10 | 17 | 15 | 1 | 1
11 | 3 | 12 | 10 | 16 | 16 | 1 | 2
12 | 4 | 12 | 10 | 15 | 3 | 1 | 3
13 | 5 | 12 | 10 | 10 | -1 | 0 | 4
(13 rows)
```

See Also1

- [Sample Data](#)
- [Boost: Bellman Ford](#)



- [https://en.wikipedia.org/wiki/Bellman%E2%80%93Ford\\_algorithm](https://en.wikipedia.org/wiki/Bellman%E2%80%93Ford_algorithm)

Boost Graph Inside

Indices and tables

- [Index](#)
- [Search Page](#)

## pgr\_dagShortestPath - Experimental¶

pgr\_dagShortestPath — Returns the shortest path for weighted directed acyclic graphs(DAG). In particular, the DAG shortest paths algorithm implemented by Boost.Graph.

□ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.2.0
  - New experimental function.
    - pgr\_dagShortestPath(Combinations)
- Version 3.0.0
  - New experimental function.

## Description¶

Shortest Path for Directed Acyclic Graph(DAG) is a graph search algorithm that solves the shortest path problem for weighted directed acyclic graph, producing a shortest path from a starting vertex (start\_vid) to an ending vertex (end\_vid).

This implementation can only be used with **adirected** graph with no cycles i.e. directed acyclic graph.

The algorithm relies on topological sorting the dag to impose a linear ordering on the vertices, and thus is more efficient for DAG's than either the Dijkstra or Bellman-Ford algorithm.

The main characteristics are:

- Process is valid for weighted directed acyclic graphs only. otherwise it will throw warnings.
- Values are returned when there is a path.
  - When the starting vertex and ending vertex are the same, there is no path.
    - The *agg\_cost* the non included values (*v*, *v*) is 0
  - When the starting vertex and ending vertex are the different and there is no path:
    - The *agg\_cost* the non included values (*u*, *v*) is \(\infty\)
- For optimization purposes, any duplicated value in the *start\_vids* or *end\_vids* are ignored.
- The returned values are ordered:
  - *start\_vid* ascending
  - *end\_vid* ascending
- Running time:  $\mathcal{O}(|\text{start\_vids}| * (V + E))$

Boost Graph Inside

## Signatures¶

Summary

pgr\_dagShortestPath([Edges SQL](#), start\_vid, end\_vid)  
 pgr\_dagShortestPath([Edges SQL](#), start\_vid, end\_vids)  
 pgr\_dagShortestPath([Edges SQL](#), start\_vids, end\_vid)  
 pgr\_dagShortestPath([Edges SQL](#), start\_vids, end\_vids)  
 pgr\_dagShortestPath([Edges SQL](#), Combinations SQL)  
 Returns set of (seq, path\_seq, node, edge, cost, agg\_cost)  
 OR EMPTY SET

One to One¶

`pgr_dagShortestPath(Edges SQL, start vid, end vid)`  
Returns set of (seq, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \5\ to vertex \11\ on a **directed** graph

```
SELECT * FROM pgr_dagShortestPath(
'SELECT id, source, target, cost FROM edges',
5, 11);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 5 | 1 | 1 | 0
2 | 2 | 6 | 4 | 1 | 1
3 | 3 | 7 | 8 | 1 | 2
4 | 4 | 11 | -1 | 0 | 3
(4 rows)
```

One to Many

`pgr_dagShortestPath(Edges SQL, start vid, end vids)`  
Returns set of (seq, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \5\ to vertices \7, 11\

```
SELECT * FROM pgr_dagShortestPath(
'SELECT id, source, target, cost FROM edges',
5, ARRAY[7, 11]);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 5 | 1 | 1 | 0
2 | 2 | 6 | 4 | 1 | 1
3 | 3 | 7 | -1 | 0 | 2
4 | 4 | 1 | 5 | 1 | 1 | 0
5 | 5 | 2 | 6 | 4 | 1 | 1
6 | 6 | 3 | 7 | 8 | 1 | 2
7 | 7 | 4 | 11 | -1 | 0 | 3
(7 rows)
```

Many to One

`pgr_dagShortestPath(Edges SQL, start vids, end vid)`  
Returns set of (seq, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \5, 10\ to vertex \11\

```
SELECT * FROM pgr_dagShortestPath(
'SELECT id, source, target, cost FROM edges',
ARRAY[5, 10], 11);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 5 | 1 | 1 | 0
2 | 2 | 6 | 4 | 1 | 1
3 | 3 | 7 | 8 | 1 | 2
4 | 4 | 11 | -1 | 0 | 3
5 | 5 | 1 | 10 | 5 | 1 | 0
6 | 6 | 2 | 11 | -1 | 0 | 1
(6 rows)
```

Many to Many

`pgr_dagShortestPath(Edges SQL, start vids, end vids)`  
Returns set of (seq, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \5, 15\ to vertices \11, 17\ on an **undirected** graph

```
SELECT * FROM pgr_dagShortestPath(
'SELECT id, source, target, cost FROM edges',
ARRAY[5, 15], ARRAY[11, 17]);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 5 | 1 | 1 | 0
2 | 2 | 6 | 4 | 1 | 1
3 | 3 | 7 | 8 | 1 | 2
4 | 4 | 11 | -1 | 0 | 3
5 | 5 | 1 | 5 | 1 | 1 | 0
6 | 6 | 2 | 6 | 4 | 1 | 1
7 | 7 | 3 | 7 | 8 | 1 | 2
8 | 8 | 4 | 11 | 9 | 1 | 3
9 | 9 | 5 | 16 | 15 | 1 | 4
10 | 10 | 6 | 17 | -1 | 0 | 5
11 | 11 | 1 | 15 | 16 | 1 | 0
12 | 12 | 2 | 16 | 15 | 1 | 1
13 | 13 | 3 | 17 | -1 | 0 | 2
(13 rows)
```

Combinations

`pgr_dagShortestPath(Edges SQL, Combinations SQL)`  
Returns set of (seq, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table on an **undirected** graph

The combinations table:

```
SELECT source, target FROM combinations;
source | target
-----+-----
5 | 6
5 | 10
6 | 5
6 | 15
6 | 14
(5 rows)
```

The query:

```
SELECT * FROM pgr_dagShortestPath(
'SELECT id, source, target, cost FROM edges',
'SELECT source, target FROM combinations');
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 5 | 1 | 1 | 0
2 | 2 | 6 | -1 | 0 | 1
(2 rows)
```

Parameters

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start vid	BIGINT	Identifier of the starting vertex of the path.
start vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end vid	BIGINT	Identifier of the ending vertex of the path.
end vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Return columns

Returns set of (seq, path\_seq [, start\_vid] [, end\_vid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_seq	INTEGER	Relative position in the path. Has value1 for the beginning of a path.

Column	Type	Description
start_vid	BIGINT	Identifier of the starting vertex. Returned when multiple starting vertices are in the query.
		<ul style="list-style-type: none"> <li>• <a href="#">Many to One</a></li> <li>• <a href="#">Many to Many</a></li> </ul>
end_vid	BIGINT	Identifier of the ending vertex. Returned when multiple ending vertices are in the query.
		<ul style="list-style-type: none"> <li>• <a href="#">One to Many</a></li> <li>• <a href="#">Many to Many</a></li> </ul>
node	BIGINT	Identifier of the node in the path from start_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go from node to the next node in the path sequence. -1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

Additional Examples¶

Example 1:

Demonstration of repeated values are ignored, and result is sorted.

```
SELECT * FROM pgr_dagShortestPath(
'SELECT id, source, target, cost FROM edges',
ARRAY[5, 10, 5, 10, 10, 5], ARRAY[11, 17, 17, 11]);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 5 | 1 | 1 | 0
2 | 2 | 6 | 4 | 1 | 1
3 | 3 | 7 | 8 | 1 | 2
4 | 4 | 4 | 11 | -1 | 0 | 3
5 | 1 | 5 | 1 | 1 | 0
6 | 2 | 6 | 4 | 1 | 1
7 | 3 | 7 | 8 | 1 | 2
8 | 4 | 4 | 11 | 9 | 1 | 3
9 | 5 | 16 | 15 | 1 | 4
10 | 6 | 17 | -1 | 0 | 5
11 | 1 | 10 | 5 | 1 | 0
12 | 2 | 11 | -1 | 0 | 1
13 | 1 | 10 | 5 | 1 | 0
14 | 2 | 11 | 9 | 1 | 1
15 | 3 | 16 | 15 | 1 | 2
16 | 4 | 17 | -1 | 0 | 3
(16 rows)
```

Example 2:

Making start\_vids the same as end\_vids

```
SELECT * FROM pgr_dagShortestPath(
'SELECT id, source, target, cost FROM edges',
ARRAY[5, 10, 11], ARRAY[5, 10, 11]);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 5 | 1 | 1 | 0
2 | 2 | 6 | 4 | 1 | 1
3 | 3 | 7 | 8 | 1 | 2
4 | 4 | 4 | 11 | -1 | 0 | 3
5 | 1 | 10 | 5 | 1 | 0
6 | 2 | 11 | -1 | 0 | 1
(6 rows)
```

Example 3:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_dagShortestPath(
'SELECT id, source, target, cost FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)');
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 4 | 1 | 0
2 | 2 | 7 | -1 | 0 | 1
(2 rows)
```

See Also¶

- [Sample Data](#)
- [Boost: DAG shortest paths](#)
- [https://en.wikipedia.org/wiki/Topological\\_sorting](https://en.wikipedia.org/wiki/Topological_sorting)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_edwardMoore - Experimental¶

pgr\_edwardMoore — Returns the shortest path using Edward-Moore algorithm.

☐ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.2.0
  - New experimental signature:
    - pgr\_edwardMoore(Combinations)
- Version 3.0.0
  - New experimental function.

Description

Edward Moore's Algorithm is an improvement of the Bellman-Ford Algorithm. It can compute the shortest paths from a single source vertex to all other vertices in a weighted directed graph. The main difference between Edward Moore's Algorithm and Bellman Ford's Algorithm lies in the run time.

The worst-case running time of the algorithm is  $\mathcal{O}(|V| * |E|)$  similar to the time complexity of Bellman-Ford algorithm. However, experiments suggest that this algorithm has an average running time complexity of  $\mathcal{O}(|E|)$  for random graphs. This is significantly faster in terms of computation speed.

Thus, the algorithm is at-best, significantly faster than Bellman-Ford algorithm and is at-worst, as good as Bellman-Ford algorithm

The main characteristics are:

- Values are returned when there is a path.
  - When the starting vertex and ending vertex are the same, there is no path:
    - The *agg\_cost* the non included values (*v*, *v*) is  $\backslash(0)$
  - When the starting vertex and ending vertex are the different and there is no path:
    - The *agg\_cost* the non included values (*u*, *v*) is  $\backslash(\infty)$
- For optimization purposes, any duplicated value in the *start vids* or *end vids* are ignored.
- The returned values are ordered:
  - *start\_vid* ascending
  - *end\_vid* ascending
- Running time:
  - Worst case:  $\mathcal{O}(|V| * |E|)$
  - Average case:  $\mathcal{O}(|E|)$

Boost Graph Inside

Signatures

Summary

pgr\_edwardMoore([Edges SQL](#), start vid, end vid, [directed])  
pgr\_edwardMoore([Edges SQL](#), start vid, end vids, [directed])  
pgr\_edwardMoore([Edges SQL](#), start vids, end vid, [directed])  
pgr\_edwardMoore([Edges SQL](#), start vids, end vids, [directed])  
pgr\_edwardMoore([Edges SQL](#), [Combinations SQL](#), [directed])  
Returns set of (seq, path\_seq, [start\_vid], [end\_vid], node, edge, cost, agg\_cost)  
OR EMPTY SET

One to One

pgr\_edwardMoore([Edges SQL](#), start vid, end vid, [directed])  
Returns set of (seq, path\_seq, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex  $\backslash(6)$  to vertex  $\backslash(10)$  on a **directed** graph

```
SELECT * FROM pgr_edwardMoore(
'SELECT id, source, target, cost, reverse_cost FROM edges',
6, 10, true);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 4 | 1 | 0
2 | 2 | 7 | 8 | 1 | 1
3 | 3 | 11 | 9 | 1 | 2
4 | 4 | 16 | 16 | 1 | 3
5 | 5 | 15 | 3 | 1 | 4
6 | 6 | 10 | -1 | 0 | 5
(6 rows)
```

One to Many

pgr\_edwardMoore([Edges SQL](#), start vid, end vids, [directed])  
Returns set of (seq, path\_seq, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertex \({6}\) to vertices \({10, 17}\) on a **directed** graph

```
SELECT * FROM pgr_edwardMoore(
'SELECT id, source, target, cost, reverse_cost FROM edges',
6, ARRAY[10, 17]);
seq | path_seq | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 10 | 6 | 4 | 1 | 0
2 | 2 | 10 | 7 | 8 | 1 | 1
3 | 3 | 10 | 11 | 9 | 1 | 2
4 | 4 | 10 | 16 | 16 | 1 | 3
5 | 5 | 10 | 15 | 3 | 1 | 4
6 | 6 | 10 | 10 | -1 | 0 | 5
7 | 1 | 17 | 6 | 4 | 1 | 0
8 | 2 | 17 | 7 | 8 | 1 | 1
9 | 3 | 17 | 11 | 11 | 1 | 2
10 | 4 | 17 | 12 | 13 | 1 | 3
11 | 5 | 17 | 17 | -1 | 0 | 4
(11 rows)
```

Many to One

pgr\_edwardMoore([Edges SQL](#), start vids, end vid, [directed])  
Returns set of (seq, path\_seq, start\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \({6, 1}\) to vertex \({17}\) on a **directed** graph

```
SELECT * FROM pgr_edwardMoore(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[6, 1], 17);
seq | path_seq | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | 6 | 1 | 0
2 | 2 | 1 | 3 | 7 | 1 | 1
3 | 3 | 1 | 7 | 8 | 1 | 2
4 | 4 | 1 | 11 | 11 | 1 | 3
5 | 5 | 1 | 12 | 13 | 1 | 4
6 | 6 | 1 | 17 | -1 | 0 | 5
7 | 1 | 6 | 6 | 4 | 1 | 0
8 | 2 | 6 | 7 | 8 | 1 | 1
9 | 3 | 6 | 11 | 11 | 1 | 2
10 | 4 | 6 | 12 | 13 | 1 | 3
11 | 5 | 6 | 17 | -1 | 0 | 4
(11 rows)
```

Many to Many

pgr\_edwardMoore([Edges SQL](#), start vids, end vids, [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

From vertices \({6, 1}\) to vertices \({10, 17}\) on an **undirected** graph

```
SELECT * FROM pgr_edwardMoore(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[6, 1], ARRAY[10, 17],
directed => false);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 10 | 6 | 1 | 0
2 | 2 | 1 | 10 | 3 | 7 | 1 | 1
3 | 3 | 1 | 10 | 7 | 4 | 1 | 2
4 | 4 | 1 | 10 | 6 | 2 | 1 | 3
5 | 5 | 1 | 10 | 10 | -1 | 0 | 4
6 | 1 | 1 | 17 | 6 | 1 | 0
7 | 2 | 1 | 17 | 3 | 7 | 1 | 1
8 | 3 | 1 | 17 | 7 | 8 | 1 | 2
9 | 4 | 1 | 17 | 11 | 11 | 1 | 3
10 | 5 | 1 | 17 | 12 | 13 | 1 | 4
11 | 6 | 1 | 17 | 17 | -1 | 0 | 5
12 | 1 | 6 | 10 | 6 | 2 | 1 | 0
13 | 2 | 6 | 10 | 10 | -1 | 0 | 1
14 | 1 | 6 | 17 | 6 | 4 | 1 | 0
15 | 2 | 6 | 17 | 7 | 8 | 1 | 1
16 | 3 | 6 | 17 | 11 | 11 | 1 | 2
17 | 4 | 6 | 17 | 12 | 13 | 1 | 3
18 | 5 | 6 | 17 | 17 | -1 | 0 | 4
(18 rows)
```

Combinations

pgr\_edwardMoore([Edges SQL](#), [Combinations SQL](#), [directed])  
Returns set of (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Using a combinations table on an **undirected** graph.

The combinations table:

```
SELECT source, target FROM combinations;
source | target
-----+-----
5 | 6
5 | 10
6 | 5
6 | 15
6 | 14
(5 rows)
```

The query:

```
SELECT * FROM pgr_edwardMoore(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT source, target FROM combinations',
false);
```

seq	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	5	6	5	1	1	0
2	2	5	6	6	-1	0	1
3	1	5	10	5	1	1	0
4	2	5	10	6	2	1	1
5	3	5	10	10	-1	0	2
6	1	6	5	6	1	1	0
7	2	6	5	5	-1	0	1
8	1	6	15	6	2	1	0
9	2	6	15	10	3	1	1
10	3	6	15	15	-1	0	2

(10 rows)

Parameters¶

Column	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below
<a href="#">Combinations SQL</a>	TEXT	<a href="#">Combinations SQL</a> as described below
start_vid	BIGINT	Identifier of the starting vertex of the path.
start_vids	ARRAY[BIGINT]	Array of identifiers of starting vertices.
end_vid	BIGINT	Identifier of the ending vertex of the path.
end_vids	ARRAY[BIGINT]	Array of identifiers of ending vertices.

Optional parameters¶

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries¶

Edges SQL¶

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Combinations SQL¶

Parameter	Type	Description
source	ANY-INTEGER	Identifier of the departure vertex.
target	ANY-INTEGER	Identifier of the arrival vertex.

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

Result columns¶

Returns set of (seq, path\_seq [, start\_vid] [, end\_vid], node, edge, cost, agg\_cost)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
path_seq	INTEGER	Relative position in the path. Has value 1 for the beginning of a path.
start_vid	BIGINT	Identifier of the starting vertex. Returned when multiple starting vertices are in the query. <ul style="list-style-type: none"> <li><a href="#">Many to One</a></li> <li><a href="#">Many to Many</a></li> </ul>
end_vid	BIGINT	Identifier of the ending vertex. Returned when multiple ending vertices are in the query. <ul style="list-style-type: none"> <li><a href="#">One to Many</a></li> <li><a href="#">Many to Many</a></li> </ul>
node	BIGINT	Identifier of the node in the path from start_vid to end_vid.
edge	BIGINT	Identifier of the edge used to go from node to the next node in the path sequence. -1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_vid to node.

#### Additional Examples ¶

##### Example 1:

Demonstration of repeated values are ignored, and result is sorted.

```
SELECT * FROM pgr_edwardMoore(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[7, 10, 15, 10, 10, 15], ARRAY[10, 7, 10, 15]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1	1	7	10	7	8	1	0
2	2	7	10	11	9	1	1
3	3	7	10	16	16	1	2
4	4	7	10	15	3	1	3
5	5	7	10	10	-1	0	4
6	1	7	15	7	8	1	0
7	2	7	15	11	9	1	1
8	3	7	15	16	16	1	2
9	4	7	15	15	-1	0	3
10	1	10	7	10	5	1	0
11	2	10	7	11	8	1	1
12	3	10	7	7	-1	0	2
13	1	10	15	10	5	1	0
14	2	10	15	11	9	1	1
15	3	10	15	16	16	1	2
16	4	10	15	15	-1	0	3
17	1	15	7	15	16	1	0
18	2	15	7	16	9	1	1
19	3	15	7	11	8	1	2
20	4	15	7	7	-1	0	3
21	1	15	10	15	3	1	0
22	2	15	10	10	-1	0	1

(22 rows)

##### Example 2:

Making start vids the same as end vids.

```
SELECT * FROM pgr_edwardMoore(
'SELECT id, source, target, cost, reverse_cost FROM edges',
ARRAY[7, 10, 15], ARRAY[7, 10, 15]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1	1	7	10	7	8	1	0
2	2	7	10	11	9	1	1
3	3	7	10	16	16	1	2
4	4	7	10	15	3	1	3
5	5	7	10	10	-1	0	4
6	1	7	15	7	8	1	0
7	2	7	15	11	9	1	1
8	3	7	15	16	16	1	2
9	4	7	15	15	-1	0	3
10	1	10	7	10	5	1	0
11	2	10	7	11	8	1	1
12	3	10	7	7	-1	0	2
13	1	10	15	10	5	1	0
14	2	10	15	11	9	1	1
15	3	10	15	16	16	1	2
16	4	10	15	15	-1	0	3
17	1	15	7	15	16	1	0
18	2	15	7	16	9	1	1
19	3	15	7	11	8	1	2
20	4	15	7	7	-1	0	3
21	1	15	10	15	3	1	0
22	2	15	10	10	-1	0	1

(22 rows)

##### Example 3:

Manually assigned vertex combinations.

```
SELECT * FROM pgr_edwardMoore(
'SELECT id, source, target, cost, reverse_cost FROM edges',
'SELECT * FROM (VALUES (6, 10), (6, 7), (12, 10)) AS combinations (source, target)');
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1	1	6	7	6	4	1	0
2	2	6	7	7	-1	0	1
3	1	6	10	6	4	1	0
4	2	6	10	7	8	1	1
5	3	6	10	11	9	1	2
6	4	6	10	16	16	1	3
7	5	6	10	15	3	1	4



8	6	6	10	10	-1	0	5
9	1	12	10	12	13	1	0
10	2	12	10	17	15	1	1
11	3	12	10	16	16	1	2
12	4	12	10	15	3	1	3
13	5	12	10	10	-1	0	4

(13 rows)

See Also

- [Sample Data](#)
- [https://en.wikipedia.org/wiki/Shortest\\_Path\\_Faster\\_Algorithm](https://en.wikipedia.org/wiki/Shortest_Path_Faster_Algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

Planar Family

- [pgr\\_isPlanar - Experimental](#)

pgr\_isPlanar - Experimental

pgr\_isPlanar — Returns a boolean depending upon the planarity of the graph.

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.2.0
  - New experimental function.

Description

A graph is planar if it can be drawn in two-dimensional space with no two of its edges crossing. Such a drawing of a planar graph is called a plane drawing. Every planar graph also admits a straight-line drawing, which is a plane drawing where each edge is represented by a line segment. When a graph has  $K_5$  or  $K_{3,3}$  as subgraph then the graph is not planar.

The main characteristics are:

- This implementation use the Boyer-Myrvold Planarity Testing.
- It will return a boolean value depending upon the planarity of the graph.
- Applicable only for **undirected** graphs.
- The algorithm does not considers traversal costs in the calculations.
- Running time:  $O(|V|)$

Boost Graph Inside

Signatures

Summary

```
pgr_isPlanar(Edges SQL)
RETURNS BOOLEAN

SELECT * FROM pgr_isPlanar(
'SELECT id, source, target, cost, reverse_cost
FROM edges'
);
pgr_isplanar
-----
t
(1 row)
```

Parameters

Parameter Type Description

Edges SQL TEXT Edges SQL as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns a boolean (pgr\_isplanar)

Column	Type	Description
pgr_isplanar	BOOLEAN	<ul style="list-style-type: none"><li>true when the graph is planar.</li><li>false when the graph is not planar.</li></ul>

Additional Examples

The following edges will make the subgraph with vertices {10, 15, 11, 16, 13} a(K\_1) graph.

```
INSERT INTO edges (source, target, cost, reverse_cost) VALUES
(10, 16, 1, 1), (10, 13, 1, 1),
(15, 11, 1, 1), (15, 13, 1, 1),
(11, 13, 1, 1), (16, 13, 1, 1);
INSERT 0 6
```

The new graph is not planar because it has a(K\_5) subgraph. Edges in blue represent(K\_5) subgraph.



```
SELECT * FROM pgr_isPlanar(
'SELECT id, source, target, cost, reverse_cost
FROM edges');
pgr_isplanar
-----
f
(1 row)
```

See Also

- Sample Data
- Boost: Boyer Myrvold

Indices and tables

- Index
- Search Page

Miscellaneous Algorithms

- [pgr\\_lengauerTarjanDominatorTree - Experimental](#)
- [pgr\\_stoerWagner - Experimental](#)
- [pgr\\_transitiveClosure - Experimental](#)
- [pgr\\_hawickCircuits - Experimental](#)

**pgr\_lengauerTarjanDominatorTree - Experimental**

pgr\_lengauerTarjanDominatorTree — Returns the immediate dominator of all vertices.

Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.2.0
  - New experimental function.

Description

The algorithm calculates the *immediate dominator* of each vertex called **idom**, once **idom** of each vertex is calculated then by making every **idom** of each vertex as its parent, the dominator tree can be built.

The main Characteristics are:

- The algorithm works in directed graph only.
- The returned values are not ordered.
- The algorithm returns *idom* of each vertex.
- If the *root vertex* not present in the graph then it returns empty set.
- Running time:  $\mathcal{O}((V+E)\log(V+E))$

Boost Graph Inside

Signatures

Summary

pgr\_lengauerTarjanDominatorTree([Edges SQL](#), **root vertex**)  
Returns set of (seq, vertex\_id, idom)  
OR EMPTY SET

Example:

The dominator tree with root vertex \5\

```
SELECT * FROM pgr_lengauertarjandominatortree(
  $$SELECT id,source,target,cost,reverse_cost FROM edges$$,
  5) ORDER BY vertex_id;
seq | vertex_id | idom
```

1	1	2
9	2	0
2	3	3
10	4	0
17	5	0
4	6	17
3	7	4
7	8	3
11	9	7
5	10	16
6	11	3
8	12	3
12	13	0
13	14	0
16	15	15
15	16	3
14	17	3

(17 rows)

Parameters

Column	Type	Description
Edges SQL	TEXT	SQL query as described above.
root vertex	BIGINT	Identifier of the starting vertex.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Returns set of (seq, vertex\_id, idom)

Column	Type	Description
seq	INTEGER	Sequential value starting from 1.
vertex_id	BIGINT	Identifier of vertex .
idom	BIGINT	Immediate dominator of vertex.

Additional Examples

Example:

Dominator tree of another component.

```
SELECT * FROM pgr_lengauertarjandominatortree(
  $$SELECT id,source,target,cost,reverse_cost FROM edges$$,
  13) ORDER BY vertex_id;
 seq | vertex_id | idom
```

```
-----+-----+-----
1 | 1 | 0
9 | 2 | 0
2 | 3 | 0
10 | 4 | 0
17 | 5 | 0
4 | 6 | 0
3 | 7 | 0
7 | 8 | 0
11 | 9 | 0
5 | 10 | 0
6 | 11 | 0
8 | 12 | 0
12 | 13 | 0
13 | 14 | 12
16 | 15 | 0
15 | 16 | 0
14 | 17 | 0
(17 rows)
```

See Also

- Sample Data
- Boost: Lengauer-Tarjan dominator
- Wikipedia: dominator tree

Indices and tables

- Index
- Search Page

☐ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.0
  - New experimental function.

Description

In graph theory, the Stoer–Wagner algorithm is a recursive algorithm to solve the minimum cut problem in undirected weighted graphs with non-negative weights. The essential idea of this algorithm is to shrink the graph by merging the most intensive vertices, until the graph only contains two combined vertex sets. At each phase, the algorithm finds the minimum s-t cut for two vertices s and t chosen as its will. Then the algorithm shrinks the edge between s and t to search for non s-t cuts. The minimum cut found in all phases will be the minimum weighted cut of the graph.

A cut is a partition of the vertices of a graph into two disjoint subsets. A minimum cut is a cut for which the size or weight of the cut is not larger than the size of any other cut. For an unweighted graph, the minimum cut would simply be the cut with the least edges. For a weighted graph, the sum of all edges' weight on the cut determines whether it is a minimum cut.

The main characteristics are:

- Process is done only on edges with positive costs.
- It's implementation is only on **undirected** graph.
- Sum of the weights of all edges between the two sets is mincut.
  - A **mincut** is a cut having the least weight.
- Values are returned when graph is connected.
  - When there is no edge in graph then EMPTY SET is return.
  - When the graph is unconnected then EMPTY SET is return.
- Sometimes a graph has multiple min-cuts, but all have the same weight. The this function determines exactly one of the min-cuts as well as its weight.
- Running time:  $\mathcal{O}(V^2E + V^2\log V)$ .

Boost Graph Inside

Signatures

pgr\_stoerWagner([Edges SQL](#))  
Returns set of (seq, edge, cost, mincut)  
OR EMPTY SET

Example:

```
min cut of the main subgraph

SELECT * FROM pgr_stoerWagner(
'SELECT id, source, target, cost, reverse_cost
FROM edges WHERE id < 17');
seq | edge | cost | mincut
-----+-----+-----+-----
1 | 6 | 1 | 1
(1 row)
```

Parameters

Parameter	Type	Description
<a href="#">Edges SQL</a>	TEXT	<a href="#">Edges SQL</a> as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
--------	------	---------	-------------

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns¶

Returns set of (seq, edge, cost, mincut)

Column	Type	Description
seq	INT	Sequential value starting from 1.
edge	BIGINT	Edges which divides the set of vertices into two.
cost	FLOAT	Cost to traverse of edge.
mincut	FLOAT	Min-cut weight of a undirected graph.

Additional Example:¶

Example:

min cut of an edge

```
SELECT * FROM pgr_stoerWagner(
  'SELECT id, source, target, cost, reverse_cost
   FROM edges WHERE id = 18');
seq | edge | cost | mincut
-----+-----+-----+-----
1 | 18 | 1 | 1
(1 row)
```

Example:

Using [pgr\\_connectedComponents](#)

```
SELECT * FROM pgr_stoerWagner(
  $$
  SELECT id, source, target, cost, reverse_cost FROM edges
 WHERE source IN (
   SELECT node FROM pgr_connectedComponents(
     'SELECT id, source, target, cost, reverse_cost FROM edges '
     WHERE component = 2)
  $$
);
seq | edge | cost | mincut
-----+-----+-----+-----
1 | 17 | 1 | 1
(1 row)
```

See Also¶

- [Sample Data](#)
- [Boost: Stoer Wagner min cut](#)
- [https://en.wikipedia.org/wiki/Stoer%E2%80%93Wagner\\_algorithm](https://en.wikipedia.org/wiki/Stoer%E2%80%93Wagner_algorithm)

Indices and tables

- [Index](#)
- [Search Page](#)

pgr\_transitiveClosure - Experimental¶

pgr\_transitiveClosure — Transitive closure graph of a directed graph.

☐ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.0.0
  - New experimental function.

Description

Transforms the input directed graph into the transitive closure of the graph.

The main characteristics are:

- Process is valid for directed graphs.
  - The transitive closure of an undirected graph produces a cluster graph
  - Reachability between vertices on an undirected graph happens when they belong to the same connected component. (see [pgr\\_connectedComponents](#))
- The returned values are not ordered
- The returned graph is compressed
- Running time:  $\backslash(O(|V||E|))$

 Boost Graph Inside

Signatures

Summary

The `pgr_transitiveClosure` function has the following signature:

`pgr_transitiveClosure`([Edges SQL](#))  
Returns set of (seq, vid, target\_array)

Example:

Rechability of a subgraph

```
SELECT * FROM pgr_transitiveclosure(
'SELECT id, source, target, cost, reverse_cost
FROM edges WHERE id IN (2, 3, 5, 11, 12, 13, 15)'
ORDER BY vid;
seq | vid | target_array
-----+-----+-----
1 | 6 | {}
6 | 8 | {12,17,16}
2 | 10 | {12,17,16,11,6}
4 | 11 | {12,17,16}
5 | 12 | {17,16}
3 | 15 | {12,17,16,10,11,6}
8 | 16 | {17,16}
7 | 17 | {17,16}
(8 rows)
```

Parameters

Parameter	Type	Description
-----------	------	-------------

[Edges SQL](#) TEXT [Edges SQL](#) as described below.

Inner Queries

Edges SQL

Column	Type	Default	Description
id	ANY-INTEGER		Identifier of the edge.
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)

Column	Type	Default	Description
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"> <li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li> </ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns¶

Returns set of (seq, vid, target\_array)

Column	Type	Description
seq	INTEGER	Sequential value starting from \((1\)
vid	BIGINT	Identifier of the source of the edges
		Identifiers of the targets of the edges
target_array	BIGINT	<ul style="list-style-type: none"> <li>Identifiers of the vertices that are reachable from vertex v.</li> </ul>

See Also¶

- Sample Data
- Boost: transitive closure
- [https://en.wikipedia.org/wiki/Transitive\\_closure](https://en.wikipedia.org/wiki/Transitive_closure)

Indices and tables

- Index
- Search Page

pgr\_hawickCircuits - Experimental¶

pgr\_hawickCircuits — Returns the list of circuits using hawick circuits algorithm.

⚠ Experimental

Warning

Possible server crash

- These functions might create a server crash

Warning

Experimental functions

- They are not officially of the current release.
- They likely will not be officially be part of the next release:
  - The functions might not make use of ANY-INTEGER and ANY-NUMERICAL
  - Name might change.
  - Signature might change.
  - Functionality might change.
  - pgTap tests might be missing.
  - Might need c/c++ coding.
  - May lack documentation.
  - Documentation if any might need to be rewritten.
  - Documentation examples might need to be automatically generated.
  - Might need a lot of feedback from the community.
  - Might depend on a proposed function of pgRouting
  - Might depend on a deprecated function of pgRouting

Availability

- Version 3.4.0
  - New experimental function.

Description¶

Hawick Circuit algorithm, is published in 2008 by Ken Hawick and Health A. James. This algorithm solves the problem of detecting and enumerating circuits in graphs. It is capable of circuit enumeration in graphs with directed-arcs, multiple-arcs and self-arcs with a memory efficient and high-performance im-plementation. It is an extension of Johnson's Algorithm of finding all the elementary circuits of a directed graph.

There are 2 variations defined in the Boost Graph Library. Here, we have implemented only 2nd as it serves the most suitable and practical usecase. In this variation we get the circuits after filtering out the circuits caused by parallel edges. Parallel edge circuits have more use cases when you want to count the no. of circuits.Maybe in future, we will also implement this variation.



The main Characteristics are:

- The algorithm implementation works only for directed graph
- It is a variation of Johnson's algorithm for circuit enumeration.
- The algorithm outputs the distinct circuits present in the graph.
- Time Complexity:  $\mathcal{O}((V + E) (c + 1))$ 
  - where  $\mathcal{O}(|E|)$  is the number of edges in the graph,
  - $\mathcal{O}(|V|)$  is the number of vertices in the graph.
  - $\mathcal{O}(|c|)$  is the number of circuits in the graph.

Boost Graph Inside

Signatures1

pgr\_hawickCircuits([Edges SQL](#))  
Returns set of (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)  
OR EMPTY SET

Example:

Circuits present in the pgRouting[Sample Data](#)

```
SELECT * FROM pgr_hawickCircuits(
  'SELECT id, source, target, cost, reverse_cost FROM edges'
);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 0 | 1 | 1 | 6 | 1 | 0
2 | 1 | 1 | 1 | 1 | 3 | 6 | 1
3 | 1 | 2 | 1 | 1 | 1 | -1 | 0
4 | 2 | 0 | 3 | 3 | 3 | 7 | 1
5 | 2 | 1 | 3 | 3 | 7 | 7 | 1
6 | 2 | 2 | 3 | 3 | 3 | -1 | 0
7 | 3 | 0 | 7 | 7 | 7 | 4 | 1
8 | 3 | 1 | 7 | 7 | 6 | 4 | 1
9 | 3 | 2 | 7 | 7 | 7 | -1 | 0
10 | 4 | 0 | 7 | 7 | 7 | 8 | 1
11 | 4 | 1 | 7 | 7 | 11 | 8 | 1
12 | 4 | 2 | 7 | 7 | 7 | -1 | 0
13 | 5 | 0 | 7 | 7 | 7 | 8 | 1
14 | 5 | 1 | 7 | 7 | 11 | 11 | 1
15 | 5 | 2 | 7 | 7 | 12 | 13 | 1
16 | 5 | 3 | 7 | 7 | 17 | 15 | 1
17 | 5 | 4 | 7 | 7 | 16 | 16 | 1
18 | 5 | 5 | 7 | 7 | 15 | 3 | 1
19 | 5 | 6 | 7 | 7 | 10 | 2 | 1
20 | 5 | 7 | 7 | 7 | 6 | 4 | 1
21 | 5 | 8 | 7 | 7 | 7 | -1 | 0
22 | 6 | 0 | 7 | 7 | 7 | 8 | 1
23 | 6 | 1 | 7 | 7 | 11 | 9 | 1
24 | 6 | 2 | 7 | 7 | 16 | 16 | 1
25 | 6 | 3 | 7 | 7 | 15 | 3 | 1
26 | 6 | 4 | 7 | 7 | 10 | 2 | 1
27 | 6 | 5 | 7 | 7 | 6 | 4 | 1
28 | 6 | 6 | 7 | 7 | 7 | -1 | 0
29 | 7 | 0 | 7 | 7 | 7 | 10 | 1
30 | 7 | 1 | 7 | 7 | 8 | 10 | 1
31 | 7 | 2 | 7 | 7 | 7 | -1 | 0
32 | 8 | 0 | 7 | 7 | 7 | 10 | 1
33 | 8 | 1 | 7 | 7 | 8 | 12 | 1
34 | 8 | 2 | 7 | 7 | 12 | 13 | 1
35 | 8 | 3 | 7 | 7 | 17 | 15 | 1
36 | 8 | 4 | 7 | 7 | 16 | 9 | 1
37 | 8 | 5 | 7 | 7 | 11 | 8 | 1
38 | 8 | 6 | 7 | 7 | 7 | -1 | 0
39 | 9 | 0 | 7 | 7 | 7 | 10 | 1
40 | 9 | 1 | 7 | 7 | 8 | 12 | 1
41 | 9 | 2 | 7 | 7 | 12 | 13 | 1
42 | 9 | 3 | 7 | 7 | 17 | 15 | 1
43 | 9 | 4 | 7 | 7 | 16 | 16 | 1
44 | 9 | 5 | 7 | 7 | 15 | 3 | 1
45 | 9 | 6 | 7 | 7 | 10 | 2 | 1
46 | 9 | 7 | 7 | 7 | 6 | 4 | 1
47 | 9 | 8 | 7 | 7 | 7 | -1 | 0
48 | 10 | 0 | 7 | 7 | 7 | 10 | 1
49 | 10 | 1 | 7 | 7 | 8 | 12 | 1
50 | 10 | 2 | 7 | 7 | 12 | 13 | 1
51 | 10 | 3 | 7 | 7 | 17 | 15 | 1
52 | 10 | 4 | 7 | 7 | 16 | 16 | 1
53 | 10 | 5 | 7 | 7 | 15 | 3 | 1
54 | 10 | 6 | 7 | 7 | 10 | 5 | 1
55 | 10 | 7 | 7 | 7 | 11 | 8 | 1
56 | 10 | 8 | 7 | 7 | 7 | -1 | 0
57 | 11 | 0 | 6 | 6 | 6 | 1 | 1
58 | 11 | 1 | 6 | 6 | 5 | 1 | 1
59 | 11 | 2 | 6 | 6 | -1 | 0
60 | 12 | 0 | 10 | 10 | 10 | 5 | 1
61 | 12 | 1 | 10 | 10 | 11 | 11 | 1
62 | 12 | 2 | 10 | 10 | 12 | 13 | 1
63 | 12 | 3 | 10 | 10 | 17 | 15 | 1
64 | 12 | 4 | 10 | 10 | 16 | 16 | 1
65 | 12 | 5 | 10 | 10 | 15 | 3 | 1
66 | 12 | 6 | 10 | 10 | 10 | -1 | 0
67 | 13 | 0 | 10 | 10 | 10 | 5 | 1
68 | 13 | 1 | 10 | 10 | 11 | 9 | 1
69 | 13 | 2 | 10 | 10 | 16 | 16 | 1
70 | 13 | 3 | 10 | 10 | 15 | 3 | 1
71 | 13 | 4 | 10 | 10 | 10 | -1 | 0
72 | 14 | 0 | 11 | 11 | 11 | 11 | 1
73 | 14 | 1 | 11 | 11 | 12 | 13 | 1
74 | 14 | 2 | 11 | 11 | 17 | 15 | 1
75 | 14 | 3 | 11 | 11 | 16 | 9 | 1
76 | 14 | 4 | 11 | 11 | 11 | -1 | 0
77 | 15 | 0 | 11 | 11 | 11 | 9 | 1
78 | 15 | 1 | 11 | 11 | 16 | 9 | 1
79 | 15 | 2 | 11 | 11 | 11 | -1 | 0
80 | 16 | 0 | 8 | 8 | 8 | 14 | 1
81 | 16 | 1 | 8 | 8 | 9 | 14 | 1
82 | 16 | 2 | 8 | 8 | 8 | -1 | 0
83 | 17 | 0 | 2 | 2 | 2 | 17 | 1
84 | 17 | 1 | 2 | 2 | 4 | 17 | 1
85 | 17 | 2 | 2 | 2 | 2 | -1 | 0
86 | 18 | 0 | 13 | 13 | 13 | 18 | 1
87 | 18 | 1 | 13 | 13 | 14 | 18 | 1
88 | 18 | 2 | 13 | 13 | 13 | -1 | 0
89 | 19 | 0 | 17 | 17 | 17 | 15 | 1
90 | 19 | 1 | 17 | 17 | 16 | 15 | 1
91 | 19 | 2 | 17 | 17 | 17 | -1 | 0
92 | 20 | 0 | 16 | 16 | 16 | 16 | 1
93 | 20 | 1 | 16 | 16 | 15 | 16 | 1
94 | 20 | 2 | 16 | 16 | 16 | -1 | 0
```

(94 rows)

Parameters

Parameter	Type	Default	Description
Edges SQL	TEXT		Edges SQL as described below.

Optional parameters

Column	Type	Default	Description
directed	BOOLEAN	true	<ul style="list-style-type: none"><li>When true the graph is considered <i>Directed</i></li><li>When false the graph is considered as <i>Undirected</i>.</li></ul>

Inner Queries

Edges SQL

Column	Type	Default	Description
source	ANY-INTEGER		Identifier of the first end point vertex of the edge.
target	ANY-INTEGER		Identifier of the second end point vertex of the edge.
cost	ANY-NUMERICAL		Weight of the edge (source, target)
reverse_cost	ANY-NUMERICAL	-1	Weight of the edge (target, source) <ul style="list-style-type: none"><li>When negative: edge (target, source) does not exist, therefore it's not part of the graph.</li></ul>

Where:

ANY-INTEGER:

SMALLINT, INTEGER, BIGINT

ANY-NUMERICAL:

SMALLINT, INTEGER, BIGINT, REAL, FLOAT

Result columns

Column	Type	Description
seq	INTEGER	Sequential value starting from 1
path_id	INTEGER	Id of the circuit starting from 1
path_seq	INTEGER	Relative position in the path. Has value 0 for beginning of the path
start_vid	BIGINT	Identifier of the starting vertex of the circuit.
end_vid	BIGINT	Identifier of the ending vertex of the circuit.
node	BIGINT	Identifier of the node in the path from a vid to next vid.
edge	BIGINT	Identifier of the edge used to go from node to the next node in the path sequence.-1 for the last node of the path.
cost	FLOAT	Cost to traverse from node using edge to the next node in the path sequence.
agg_cost	FLOAT	Aggregate cost from start_v to node.

See Also

- Sample Data
- Boost: Hawick Circuit Algorithm

Indices and tables

- Index
- Search Page

See Also

Indices and tables

- Index
- Search Page

## Release Notes¶

### pgRouting 3.8.0 Release Notes¶

To see all issues & pull requests closed by this release see the[Git closed milestone for 3.8.0](#)

Promotion to official function of pgRouting.

Metric

- [#2760](#): Promoted to official pgr\_degree in version 3.8
  - Error messages adjustment.
  - New signature with only Edges SQL.
  - Function promoted to official.

Utilities

- [#2772](#): Promoted to official pgr\_extractVertices in version 3.8
  - Error messages adjustment.
  - Function promoted to official.
- [#2774](#): Promoted to official pgr\_findCloseEdges in version 3.8
  - Error messages adjustment.
  - partial option is removed.
  - Function promoted to official.
- [#2873](#): Promoted to official pgr\_separateCrossing in version 3.8
  - Function promoted to official.
  - Proposed function.
- [#2874](#): Promoted to official pgr\_separateTouching in version 3.8
  - Function promoted to official.
  - Proposed function.

Proposed functions

Contraction

- [#2790](#): pgr\_contractionDeadEnd new contraction function
- [#2791](#): pgr\_contractionLinear new contraction function
- [#2536](#): Support for contraction hierarchies (pgr\_contractionHierarchies)

Utilities

- [#2848](#): Create pgr\_separateCrossing new utility function
- [#2849](#): Create of pgr\_separateTouching new utility function

Official functions changes

- [#2786](#): pgr\_contraction(edges) new signature
  - New signature:
    - Previously compulsory parameter **Contraction order** is now optional with name methods.
    - New name and order of optional parameters.
  - Deprecated signature pgr\_contraction(text,bigint[],integer,bigint[],boolean)

C/C++ code enhancements

- [#2802](#): Code reorganization on pgr\_contraction
- Other enhancements: [#2869](#)

SQL code enhancements

- [#2850](#): Rewrite pgr\_nodeNetwork

Deprecation of SQL functions

- [#2749](#): Deprecate pgr\_AlphaShape in 3.8
- [#2750](#): Deprecate pgr\_CreateTopology in 3.8
- [#2753](#): Deprecate pgr\_analyzeGraph in 3.8
- [#2754](#): Deprecate pgr\_analyzeOneWay in 3.8
- [#2826](#): Deprecate pgr\_createVerticesTable in 3.8
- [#2847](#): Deprecate pgr\_nodeNetwork in 3.8

In the deprecated functions:

- Migration section is created.
- The use of the functions is removed in the documentation.

### All releases¶

#### Release Notes¶

To see the full list of changes check the list of [Git commits](#) on Github.

Mayors

- [pgRouting 3](#)
- [pgRouting 2](#)

- [pgRouting 1](#)

## [pgRouting 3.8](#)

### Minors 3.x

- [pgRouting 3.8](#)
- [pgRouting 3.7](#)
- [pgRouting 3.6](#)
- [pgRouting 3.5](#)
- [pgRouting 3.4](#)
- [pgRouting 3.3](#)
- [pgRouting 3.2](#)
- [pgRouting 3.1](#)
- [pgRouting 3.0](#)

## [pgRouting 3.8](#)

### Contents

- [pgRouting 3.8.0 Release Notes](#)

## [pgRouting 3.8.0 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.8.0](#)

Promotion to official function of pgRouting.

### Metric

- [#2760](#): Promoted to official pgr\_degree in version 3.8
  - Error messages adjustment.
  - New signature with only Edges SQL.
  - Function promoted to official.

### Utilities

- [#2772](#): Promoted to official pgr\_extractVertices in version 3.8
  - Error messages adjustment.
  - Function promoted to official.
- [#2774](#): Promoted to official pgr\_findCloseEdges in version 3.8
  - Error messages adjustment.
  - partial option is removed.
  - Function promoted to official.
- [#2873](#): Promoted to official pgr\_separateCrossing in version 3.8
  - Function promoted to official.
  - Proposed function.
- [#2874](#): Promoted to official pgr\_separateTouching in version 3.8
  - Function promoted to official.
  - Proposed function.

### Proposed functions

### Contraction

- [#2790](#): pgr\_contractionDeadEnd new contraction function
- [#2791](#): pgr\_contractionLinear new contraction function
- [#2536](#): Support for contraction hierarchies (pgr\_contractionHierarchies)

### Utilities

- [#2848](#): Create pgr\_separateCrossing new utility function
- [#2849](#): Create of pgr\_separateTouching new utility function

### Official functions changes

- [#2786](#): pgr\_contraction(edges) new signature
  - New signature:
    - Previously compulsory parameter **Contraction order** is now optional with name methods.
    - New name and order of optional parameters.
  - Deprecated signature pgr\_contraction(text,bigint[],integer,bigint[],boolean)

### C/C++ code enhancements

- [#2802](#): Code reorganization on pgr\_contraction
- Other enhancements: [#2869](#)

### SQL code enhancements

- [#2850](#): Rewrite pgr\_nodeNetwork

### Deprecation of SQL functions

- [#2749](#): Deprecate pgr\_AlphaShape in 3.8
- [#2750](#): Deprecate pgr\_CreateTopology in 3.8

- [#2753](#): Deprecate pgr\_analyzeGraph in 3.8
- [#2754](#): Deprecate pgr\_analyzeOneWay in 3.8
- [#2826](#): Deprecate pgr\_createVerticesTable in 3.8
- [#2847](#): Deprecate pgr\_nodeNetwork in 3.8

In the deprecated functions:

- Migration section is created.
- The use of the functions is removed in the documentation.

#### [pgRouting 3.7](#)

Contents

- [pgRouting 3.7.3 Release Notes](#)
- [pgRouting 3.7.2 Release Notes](#)
- [pgRouting 3.7.1 Release Notes](#)
- [pgRouting 3.7.0 Release Notes](#)

#### [pgRouting 3.7.3 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.7.3](#)

- [#2731](#) Build Failure on Ubuntu 22

#### [pgRouting 3.7.2 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.7.2](#)

Build

- [#2713](#) cmake missing some policies and min version
  - Using OLD policies: CMP0148, CMP0144, CMP0167
  - Minimum cmake version 3.12

Bug fixes

- [#2707](#) Build failure in pgRouting 3.7.1 on Alpine
- [#2706](#) winnie crashing on pgr\_betweennessCentrality

#### [pgRouting 3.7.1 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.7.1](#)

Bug fixes

- [#2680](#) fails to compile under mingw64 gcc 13.2
- [#2689](#) When point is a vertex, the withPoints family do not return results.

C/C++ code enhancemet

- TRSP family

#### [pgRouting 3.7.0 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.7.0](#)

Support

- [#2656](#) Stop support of PostgreSQL 12 on pgrouting v3.7
  - Stopping support of PostgreSQL 12
  - CI does not test for PostgreSQL 12

New experimental functions

- Metrics
  - pgr\_betweennessCentrality

Official functions changes

- [#2605](#) Standardize spanning tree functions output
  - Functions:
    - pgr\_kruskalDD
    - pgr\_kruskalDFS
    - pgr\_kruskalBFS
    - pgr\_primDD
    - pgr\_primDFS
    - pgr\_primBFS
  - Standardizing output columns to (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
    - Added pred result columns.

Experimental promoted to proposed.

- [#2635](#) pgr\_LineGraph ignores directed flag and use negative values for identifiers.
  - pgr\_lineGraph
    - Function promoted to proposed.
    - Works for directed and undirected graphs.

## Code enhancement

- [#2599](#) Driving distance cleanup
- [#2607](#) Read postgresql data on C++
- [#2614](#) Clang tidy does not work

## [pgRouting 3.6.1](#)

### Contents

- [pgRouting 3.6.3 Release Notes](#)
- [pgRouting 3.6.2 Release Notes](#)
- [pgRouting 3.6.1 Release Notes](#)
- [pgRouting 3.6.0 Release Notes](#)

## [pgRouting 3.6.3 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.6.3](#)

### Build

- Explicit minimum requirements:
  - postgres 11.0.0
  - postgis 3.0.0
- g++ 13+ is supported

### Code fixes

- Fix warnings from cpplint.
- Fix warnings from clang 18.

### CI tests

- Add a clang tidy test on changed files.
- Update test not done on versions: 3.0.1, 3.0.2, 3.0.3, 3.0.4, 3.1.0, 3.1.1, 3.1.2

### Documentation

- Results of documentation queries adujsted to boost 1.83.0 version:
  - `pgr_edgeDisjointPaths`
  - `pgr_stoerWagner`

### pgtap tests

- bug fixes

## [pgRouting 3.6.2 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.6.2](#)

### Upgrade fix

- The upgrade was failing for same minor

### Code fixes

- Fix warnings from cpplint

### Others

- Adjust NEWS generator
  - Name change to *NEWS.md* for better visualization on GitHub

## [pgRouting 3.6.1 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.6.1](#)

- [#2588](#) pgrouting 3.6.0 fails to build on OSX

## [pgRouting 3.6.0 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.6.0](#)

### Official functions changes

- [#2516](#) Standardize output `pgr_aStar`
  - Standardize output columns to (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)
    - `pgr_aStar(One to One)` added `start_vid` and `end_vid` columns.
    - `pgr_aStar(One to Many)` added `end_vid` column.
    - `pgr_aStar(Many to One)` added `start_vid` column.
- [#2523](#) Standardize output `pgr_bdAstar`
  - Standardize output columns to (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)
    - `pgr_bdAstar(One to One)` added `start_vid` and `end_vid` columns.
    - `pgr_bdAstar(One to Many)` added `end_vid` column.
    - `pgr_bdAstar(Many to One)` added `start_vid` column.
- [#2547](#) Standardize output and modifying signature `pgr_KSP`
  - Standardizing output columns to (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)
  - `pgr_ksp(One to One)`
    - Added `start_vid` and `end_vid` result columns.

- New proposed signatures:
  - `pgr_ksp(One to Many)`
  - `pgr_ksp(Many to One)`
  - `pgr_ksp(Many to Many)`
  - `pgr_ksp(Combinations)`
- [#2548](#) Standardize output `pgr_drivingDistance`
  - Standardizing output columns to (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
    - `pgr_drivingDistance(Single vertex)`
      - Added depth and start\_vid result columns.
    - `pgr_drivingDistance(Multiple vertices)`
      - Result column name change: from `v` to `start_vid`.
      - Added depth and `pred` result columns.

#### Proposed functions changes

- [#2544](#) Standardize output and modifying signature `pgr_withPointsDD`
  - Signature change: `driving_side` parameter changed from named optional to unnamed compulsory **driving side**.
    - `pgr_withPointsDD(Single vertex)`
    - `pgr_withPointsDD(Multiple vertices)`
  - Standardizing output columns to (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
    - `pgr_withPointsDD(Single vertex)`
      - Added depth, `pred` and `start_vid` column.
    - `pgr_withPointsDD(Multiple vertices)`
      - Added depth, `pred` columns.
  - When `details` is false:
    - Only points that are visited are removed, that is, points reached within the distance are included
  - Deprecated signatures
    - `pgr_withpointsdd(text,text,bigint,double precision,boolean,character,boolean)`
    - `pgr_withpointsdd(text,text,array,double precision,boolean,character,boolean,boolean)`
- [#2546](#) Standardize output and modifying signature `pgr_withPointsKSP`
  - Standardizing output columns to (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)
  - `pgr_withPointsKSP(One to One)`
    - Signature change: `driving_side` parameter changed from named optional to unnamed compulsory **driving side**.
    - Added `start_vid` and `end_vid` result columns.
  - New proposed signatures:
    - `pgr_withPointsKSP(One to Many)`
    - `pgr_withPointsKSP(Many to One)`
    - `pgr_withPointsKSP(Many to Many)`
    - `pgr_withPointsKSP(Combinations)`
  - Deprecated signature
    - `pgr_withpointsksp(text,text,bigint,bigint,integer,boolean,boolean,char,boolean)``

#### C/C++ code enhancements

- [#2504](#) To C++ pg data get, fetch and check.
  - Stopping support for compilation with MSVC.
- [#2505](#) Using namespace.
- [#2512](#) [Dijkstra] Removing duplicate code on Dijkstra.
- [#2517](#) Astar code simplification.
- [#2521](#) Dijkstra code simplification.
- [#2522](#) bdAstar code simplification.

#### Documentation

- [#2490](#) Automatic page history links.
- `..rubric::` Standardize SQL
- [#2555](#) Standardize deprecated messages
- On new internal function: do not use named parameters and default parameters.

#### [pgRouting 3.5f](#)

#### Contents

- [pgRouting 3.5.1 Release Notes](#)
- [pgRouting 3.5.0 Release Notes](#)

#### [pgRouting 3.5.1 Release Notesf](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.5.1](#)

#### Documentation fixes

Changes on the documentation to the following:

- `pgr_degree`
- `pgr_dijkstra`
- `pgr_ksp`
- Automatic page history links
  - using `bootstrap_version` 2 because 3+ does not do dropdowns

Issue fixes

- [#2565](#) `pgr_lengauerTarjanDominatorTree` triggers an assertion

SQL enhancements

- [#2561](#) Not use wildcards on SQL

pgtap tests

- [#2559](#) pgtap test using sampledata

Build fixes

- Fix winnie build

Code fixes

- Fix clang warnings
  - Grouping headers of postgres readers

[pgRouting 3.5.0 Release Notes¶](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.5.0](#)

Official functions changes

- Dijkstra
  - Standardize output columns to (seq, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)
    - `pgr_dijkstra`(One to One) added `start_vid` and `end_vid` columns.
    - `pgr_dijkstra`(One to Many) added `end_vid` column.
    - `pgr_dijkstra`(Many to One) added `start_vid` column.

[pgRouting 3.4¶](#)

Contents

- [pgRouting 3.4.2 Release Notes](#)
- [pgRouting 3.4.1 Release Notes](#)
- [pgRouting 3.4.0 Release Notes](#)

[pgRouting 3.4.2 Release Notes¶](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.4.2](#)

Issue fixes

- [#2394](#): `pgr_bdAstar` accumulates heuristic cost in visited node cost.
- [#2427](#): `pgr_createVerticesTable` & `pgr_createTopology`, variable should be of type Record.

[pgRouting 3.4.1 Release Notes¶](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.4.1](#)

Issue fixes

- [#2401](#): pgRouting 3.4.0 do not build docs when sphinx is too low or missing
- [#2398](#): v3.4.0 does not upgrade from 3.3.3

[pgRouting 3.4.0 Release Notes¶](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.4.0](#)

Issue fixes

- [#1891](#): `pgr_ksp` doesn't give all correct shortest path

New proposed functions.

- With points
  - `pgr_withPointsVia`(One Via)
- Turn Restrictions
  - Via with turn restrictions
    - `pgr_trspVia`(One Via)
    - `pgr_trspVia_withPoints`(One Via)
  - `pgr_trsp`
    - `pgr_trsp`(One to One)
    - `pgr_trsp`(One to Many)
    - `pgr_trsp`(Many to One)
    - `pgr_trsp`(Many to Many)
    - `pgr_trsp`(Combinations)
  - `pgr_trsp_withPoints`



- `pgr_trsp_withPoints(One to One)`
- `pgr_trsp_withPoints(One to Many)`
- `pgr_trsp_withPoints(Many to One)`
- `pgr_trsp_withPoints(Many to Many)`
- `pgr_trsp_withPoints(Combinations)`
- Topology
  - `pgr_degree`
- Utilities
  - `pgr_findCloseEdges(One point)`
  - `pgr_findCloseEdges(Many points)`

#### New experimental functions

- Ordering
  - `pgr_cuthillMcKeeOrdering`
- Unclassified
  - `pgr_hawickCircuits`

#### Official functions changes

- Flow functions
  - `pgr_maxCardinalityMatch(text)`
    - Deprecating: `pgr_maxCardinalityMatch(text,boolean)`

#### Deprecated Functions

- Turn Restrictions
  - `pgr_trsp(text,integer,integer,boolean,boolean,text)`
  - `pgr_trsp(text,integer,float8,integer,float8,boolean,boolean,text)`
  - `pgr_trspViaVertices(text,array,boolean,boolean,text)`
  - `pgr_trspViaEdges(text,integer[],float[],boolean,boolean,text)`

#### [pgRouting 3.3.5](#)

##### Contents

- [pgRouting 3.3.5 Release Notes](#)
- [pgRouting 3.3.4 Release Notes](#)
- [pgRouting 3.3.3 Release Notes](#)
- [pgRouting 3.3.2 Release Notes](#)
- [pgRouting 3.3.1 Release Notes](#)
- [pgRouting 3.3.0 Release Notes](#)

#### [pgRouting 3.3.5 Release Notes](#)

- [#2401](#): pgRouting 3.4.0 do not build docs when sphinx is too low or missing

#### [pgRouting 3.3.4 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.3.4](#)

##### Issue fixes

- [#2400](#): pgRouting 3.3.3 does not build in focal

#### [pgRouting 3.3.3 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.3.3](#)

##### Issue fixes

- [#1891](#): `pgr_ksp` doesn't give all correct shortest path

#### Official functions changes

- Flow functions
  - `pgr_maxCardinalityMatch(text,boolean)`
    - Ignoring optional boolean parameter, as the algorithm works only for undirected graphs.

#### [pgRouting 3.3.2 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.3.2](#)

- Revised documentation
  - Simplifying table names and table columns, for example:
    - `edges` instead of `edge_table`
      - Removing unused columns `category_id` and `reverse_category_id`.
    - `combinations` instead of `combinations_table`
      - Using PostGIS standard for geometry column.
        - `geom` instead of `the_geom`
  - Avoiding usage of functions that modify indexes, columns etc on tables.
    - Using `pgr_extractVertices` to create a routing topology

- Restructure of the pgRouting concepts page.

#### Issue fixes

- [#2276](#): edgeDisjointPaths issues with start\_vid and combinations
- [#2312](#): pgr\_extractVertices error when target is not BIGINT
- [#2357](#): Apply clang-tidy performance-\*

#### [pgRouting 3.3.1 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.3.1](#) on Github.

#### Issue fixes

- [#2216](#): Warnings when using clang
- [#2266](#): Error processing restrictions

#### [pgRouting 3.3.0 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.3.0](#) on Github.

#### Issue fixes

- [#2057](#): trspViaEdges columns in different order
- [#2087](#): pgr\_extractVertices to proposed
- [#2201](#): pgr\_depthFirstSearch to proposed
- [#2202](#): pgr\_sequentialVertexColoring to proposed
- [#2203](#): pgr\_dijkstraNear and pgr\_dijkstraNearCost to proposed

#### New experimental functions

- Coloring
  - pgr\_edgeColoring

#### Experimental promoted to Proposed

- Dijkstra
  - pgr\_dijkstraNear
    - pgr\_dijkstraNear(Combinations)
    - pgr\_dijkstraNear(Many to Many)
    - pgr\_dijkstraNear(Many to One)
    - pgr\_dijkstraNear(One to Many)
  - pgr\_dijkstraNearCost
    - pgr\_dijkstraNearCost(Combinations)
    - pgr\_dijkstraNearCost(Many to Many)
    - pgr\_dijkstraNearCost(Many to One)
    - pgr\_dijkstraNearCost(One to Many)
- Coloring
  - pgr\_sequentialVertexColoring
- Topology
  - pgr\_extractVertices
- Traversal
  - pgr\_depthFirstSearch(Multiple vertices)
  - pgr\_depthFirstSearch(Single vertex)

#### [pgRouting 3.2](#)

#### Contents

- [pgRouting 3.2.2 Release Notes](#)
- [pgRouting 3.2.1 Release Notes](#)
- [pgRouting 3.2.0 Release Notes](#)

#### [pgRouting 3.2.2 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.2.2](#) on Github.

#### Issue fixes

- [#2093](#): Compilation on Visual Studio
- [#2189](#): Build error on RHEL 7

#### [pgRouting 3.2.1 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.2.1](#) on Github.

#### Issue fixes

- [#1883](#): pgr\_TSPeuclidean crashes connection on Windows
  - The solution is to use Boost::graph::metric\_tsp\_approx
  - To not break user's code the optional parameters related to the TSP Annaeling are ignored
  - The function with the annaeling optional parameters is deprecated

#### [pgRouting 3.2.0 Release Notes](#)

To see all issues & pull requests closed by this release see the[Git closed milestone for 3.2.0](#) on Github.

#### Build

- [#1850](#): Change Boost min version to 1.56
  - Removing support for Boost v1.53, v1.54 & v1.55

#### New experimental functions

- `pgr_bellmanFord(Combinations)`
- `pgr_binaryBreadthFirstSearch(Combinations)`
- `pgr_bipartite`
- `pgr_dagShortestPath(Combinations)`
- `pgr_depthFirstSearch`
- `Dijkstra Near`
  - `pgr_dijkstraNear`
    - `pgr_dijkstraNear(One to Many)`
    - `pgr_dijkstraNear(Many to One)`
    - `pgr_dijkstraNear(Many to Many)`
    - `pgr_dijkstraNear(Combinations)`
  - `pgr_dijkstraNearCost`
    - `pgr_dijkstraNearCost(One to Many)`
    - `pgr_dijkstraNearCost(Many to One)`
    - `pgr_dijkstraNearCost(Many to Many)`
    - `pgr_dijkstraNearCost(Combinations)`
- `pgr_edwardMoore(Combinations)`
- `pgr_isPlanar`
- `pgr_lengauerTarjanDominatorTree`
- `pgr_makeConnected`
- `Flow`
  - `pgr_maxFlowMinCost(Combinations)`
  - `pgr_maxFlowMinCost_Cost(Combinations)`
- `pgr_sequentialVertexColoring`

#### New proposed functions.

- `Astar`
  - `pgr_aStar(Combinations)`
  - `pgr_aStarCost(Combinations)`
- `Bidirectional Astar`
  - `pgr_bdAstar(Combinations)`
  - `pgr_bdAstarCost(Combinations)`
- `Bidirectional Dijkstra`
  - `pgr_bdDijkstra(Combinations)`
  - `pgr_bdDijkstraCost(Combinations)`
- `Flow`
  - `pgr_boykovKolmogorov(Combinations)`
  - `pgr_edgeDisjointPaths(Combinations)`
  - `pgr_edmondsKarp(Combinations)`
  - `pgr_maxFlow(Combinations)`
  - `pgr_pushRelabel(Combinations)`
- `pgr_withPoints(Combinations)`
- `pgr_withPointsCost(Combinations)`

#### [pgRouting 3.15](#)

#### Contents

- [pgRouting 3.1.4 Release Notes](#)
- [pgRouting 3.1.3 Release Notes](#)
- [pgRouting 3.1.2 Release Notes](#)
- [pgRouting 3.1.1 Release Notes](#)
- [pgRouting 3.1.0 Release Notes](#)

#### [pgRouting 3.1.4 Release Notes](#)

To see all issues & pull requests closed by this release see the[Git closed milestone for 3.1.4](#) on Github.

#### Issues fixes

- [#2189](#): Build error on RHEL 7

#### [pgRouting 3.1.3 Release Notes](#)

To see all issues & pull requests closed by this release see the[Git closed milestone for 3.1.3](#) on Github.

Issues fixes

- [#1825](#): Boost versions are not honored
- [#1849](#): Boost 1.75.0 geometry "point\_xy.hpp" build error on macOS environment
- [#1861](#): vrp functions crash server

#### [pgRouting 3.1.2 Release Notes¶](#)

To see all issues & pull requests closed by this release see the[Git closed milestone for 3.1.2](#) on Github.

Issues fixes

- [#1304](#): FreeBSD 12 64-bit crashes on pgr\_vrOneDepot tests Experimental Function
- [#1356](#): tools/testers/pg\_prove\_tests.sh fails when PostgreSQL port is not passed
- [#1725](#): Server crash on pgr\_pickDeliver and pgr\_vrpOneDepot on openbsd
- [#1760](#): TSP server crash on ubuntu 20.04 #1760
- [#1770](#): Remove warnings when using clang compiler

#### [pgRouting 3.1.1 Release Notes¶](#)

To see all issues & pull requests closed by this release see the[Git closed milestone for 3.1.1](#) on Github.

Issues fixes

- [#1733](#): pgr\_bdAstar fails when source or target vertex does not exist in the graph
- [#1647](#): Linear Contraction contracts self loops
- [#1640](#): pgr\_withPoints fails when points\_sql is empty
- [#1616](#): Path evaluation on C++ not updated before the results go back to C
- [#1300](#): pgr\_chinesePostman crash on test data

#### [pgRouting 3.1.0 Release Notes¶](#)

To see all issues & pull requests closed by this release see the[Git closed milestone for 3.1.0](#) on Github.

New proposed functions.

- pgr\_dijkstra(combinations)
- pgr\_dijkstraCost(combinations)

Build changes

- Minimal requirement for Sphinx: version 1.8

#### [pgRouting 3.0¶](#)

Contents

- [pgRouting 3.0.6 Release Notes](#)
- [pgRouting 3.0.5 Release Notes](#)
- [pgRouting 3.0.4 Release Notes](#)
- [pgRouting 3.0.3 Release Notes](#)
- [pgRouting 3.0.2 Release Notes](#)
- [pgRouting 3.0.1 Release Notes](#)
- [pgRouting 3.0.0 Release Notes](#)

#### [pgRouting 3.0.6 Release Notes¶](#)

To see all issues & pull requests closed by this release see the[Git closed milestone for 3.0.6](#) on Github.

Issues fixes

- [#2189](#): Build error on RHEL 7

#### [pgRouting 3.0.5 Release Notes¶](#)

To see all issues & pull requests closed by this release see the[Git closed milestone for 3.0.5](#) on Github.

Backport issue fixes

- [#1825](#): Boost versions are not honored
- [#1849](#): Boost 1.75.0 geometry "point\_xy.hpp" build error on macOS environment
- [#1861](#): vrp functions crash server

#### [pgRouting 3.0.4 Release Notes¶](#)

To see all issues & pull requests closed by this release see the[Git closed milestone for 3.0.4](#) on Github.

Backport issue fixes

- [#1304](#): FreeBSD 12 64-bit crashes on pgr\_vrOneDepot tests Experimental Function
- [#1356](#): tools/testers/pg\_prove\_tests.sh fails when PostgreSQL port is not passed
- [#1725](#): Server crash on pgr\_pickDeliver and pgr\_vrpOneDepot on openbsd
- [#1760](#): TSP server crash on ubuntu 20.04 #1760
- [#1770](#): Remove warnings when using clang compiler

#### [pgRouting 3.0.3 Release Notes¶](#)

Backport issue fixes

- [#1733](#): pgr\_bdAstar fails when source or target vertex does not exist in the graph

- [#1647](#): Linear Contraction contracts self loops
- [#1640](#): pgr\_withPoints fails when points\_sql is empty
- [#1616](#): Path evaluation on C++ not updated before the results go back to C
- [#1300](#): pgr\_chinesePostman crash on test data

#### [pgRouting 3.0.2 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.0.2](#) on Github.

Issues fixes

- [#1378](#): Visual Studio build failing

#### [pgRouting 3.0.1 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.0.1](#) on Github.

Issues fixes

- [#232](#): Honor client cancel requests in C /C++ code

#### [pgRouting 3.0.0 Release Notes](#)

To see all issues & pull requests closed by this release see the [Git closed milestone for 3.0.0](#) on Github.

Fixed Issues

- [#1153](#): Renamed pgr\_eucledianTSP to pgr\_TSPeuclidean
- [#1188](#): Removed CGAL dependency
- [#1002](#): Fixed contraction issues:
  - [#1004](#): Contracts when forbidden vertices do not belong to graph
  - [#1005](#): Intermideate results eliminated
  - [#1006](#): No loss of information

New Functions

- Kruskal family
  - pgr\_kruskal
  - pgr\_kruskalBFS
  - pgr\_kruskalDD
  - pgr\_kruskalDFS
- Prim family
  - pgr\_prim
  - pgr\_primDD
  - pgr\_primDFS
  - pgr\_primBFS

Proposed moved to official on pgRouting

- aStar Family
  - pgr\_aStar(One to Many)
  - pgr\_aStar(Many to One)
  - pgr\_aStar(Many to Many)
  - pgr\_aStarCost(One to One)
  - pgr\_aStarCost(One to Many)
  - pgr\_aStarCost(Many to One)
  - pgr\_aStarCost(Many to Many)
  - pgr\_aStarCostMatrix
- bdAstar Family
  - pgr\_bdAstar(One to Many)
  - pgr\_bdAstar(Many to One)
  - pgr\_bdAstar(Many to Many)
  - pgr\_bdAstarCost(One to One)
  - pgr\_bdAstarCost(One to Many)
  - pgr\_bdAstarCost(Many to One)
  - pgr\_bdAstarCost(Many to Many)
  - pgr\_bdAstarCostMatrix
- bdDijkstra Family
  - pgr\_bdDijkstra(One to Many)
  - pgr\_bdDijkstra(Many to One)
  - pgr\_bdDijkstra(Many to Many)
  - pgr\_bdDijkstraCost(One to One)
  - pgr\_bdDijkstraCost(One to Many)
  - pgr\_bdDijkstraCost(Many to One)
  - pgr\_bdDijkstraCost(Many to Many)

- pgr\_bdDijkstraCostMatrix
- Flow Family
  - pgr\_pushRelabel(One to One)
  - pgr\_pushRelabel(One to Many)
  - pgr\_pushRelabel(Many to One)
  - pgr\_pushRelabel(Many to Many)
  - pgr\_edmondsKarp(One to One)
  - pgr\_edmondsKarp(One to Many)
  - pgr\_edmondsKarp(Many to One)
  - pgr\_edmondsKarp(Many to Many)
  - pgr\_boykovKolmogorov (One to One)
  - pgr\_boykovKolmogorov (One to Many)
  - pgr\_boykovKolmogorov (Many to One)
  - pgr\_boykovKolmogorov (Many to Many)
  - pgr\_maxCardinalityMatching
  - pgr\_maxFlow
  - pgr\_edgeDisjointPaths(One to One)
  - pgr\_edgeDisjointPaths(One to Many)
  - pgr\_edgeDisjointPaths(Many to One)
  - pgr\_edgeDisjointPaths(Many to Many)
- Components family
  - pgr\_connectedComponents
  - pgr\_strongComponents
  - pgr\_biconnectedComponents
  - pgr\_articulationPoints
  - pgr\_bridges
- Contraction:
  - Removed unnecessary column seq
  - Bug Fixes

#### New experimental functions

- pgr\_maxFlowMinCost
- pgr\_maxFlowMinCost\_Cost
- pgr\_extractVertices
- pgr\_turnRestrictedPath
- pgr\_stoerWagner
- pgr\_dagShortestpath
- pgr\_topologicalSort
- pgr\_transitiveClosure
- VRP category
  - pgr\_pickDeliverEuclidean
  - pgr\_pickDeliver
- Chinese Postman family
  - pgr\_chinesePostman
  - pgr\_chinesePostmanCost
- Breadth First Search family
  - pgr\_breadthFirstSearch
  - pgr\_binaryBreadthFirstSearch
- Bellman Ford family
  - pgr\_bellmanFord
  - pgr\_edwardMoore

#### Moved to legacy

- Experimental functions
  - pgr\_labelGraph - Use the components family of functions instead.
  - Max flow - functions were renamed on v2.5.0
    - pgr\_maxFlowPushRelabel
    - pgr\_maxFlowBoykovKolmogorov
    - pgr\_maxFlowEdmondsKarp
    - pgr\_maximumcardinalitymatching
  - VRP
    - pgr\_gsoc\_vrppdw

- TSP old signatures
- pgr\_pointsAsPolygon
- pgr\_alphaShape old signature

## [pgRouting 2.6](#)

### Minors 2.x

- [pgRouting 2.6](#)
- [pgRouting 2.5](#)
- [pgRouting 2.4](#)
- [pgRouting 2.3](#)
- [pgRouting 2.2](#)
- [pgRouting 2.1](#)
- [pgRouting 2.0](#)

## [pgRouting 2.6.3](#)

### Contents

- [pgRouting 2.6.3 Release Notes](#)
- [pgRouting 2.6.2 Release Notes](#)
- [pgRouting 2.6.1 Release Notes](#)
- [pgRouting 2.6.0 Release Notes](#)

## [pgRouting 2.6.3 Release Notes](#)

To see the issues closed by this release see the [Git closed milestone for 2.6.3](#) on Github.

### Bug fixes

- [#1219](#) Implicit cast for via\_path integer to text
- [#1193](#) Fixed pgr\_pointsAsPolygon breaking when comparing strings in WHERE clause
- [#1185](#) Improve FindPostgreSQL.cmake

## [pgRouting 2.6.2 Release Notes](#)

To see the issues closed by this release see the [Git closed milestone for 2.6.2](#) on Github.

### Bug fixes

- [#1152](#) Fixes driving distance when vertex is not part of the graph
- [#1098](#) Fixes windows test
- [#1165](#) Fixes build for python3 and perl5

## [pgRouting 2.6.1 Release Notes](#)

To see the issues closed by this release see the [Git closed milestone for 2.6.1](#) on Github.

- Fixes server crash on several functions.
  - pgr\_floydWarshall
  - pgr\_johnson
  - pgr\_aStar
  - pgr\_bdAstar
  - pgr\_bdDijkstra
  - pgr\_alphashape
  - pgr\_dijkstraCostMatrix
  - pgr\_dijkstra
  - pgr\_dijkstraCost
  - pgr\_drivingDistance
  - pgr\_KSP
  - pgr\_dijkstraVia (proposed)
  - pgr\_boykovKolmogorov (proposed)
  - pgr\_edgeDisjointPaths (proposed)
  - pgr\_edmondsKarp (proposed)
  - pgr\_maxCardinalityMatch (proposed)
  - pgr\_maxFlow (proposed)
  - pgr\_withPoints (proposed)
  - pgr\_withPointsCost (proposed)
  - pgr\_withPointsKSP (proposed)
  - pgr\_withPointsDD (proposed)
  - pgr\_withPointsCostMatrix (proposed)
  - pgr\_contractGraph (experimental)
  - pgr\_pushRelabel (experimental)
  - pgr\_vrpOneDepot (experimental)

- `pgr_gsoc_vrppdtw` (experimental)
- Fixes for deprecated functions where also applied but not tested
- Removed compilation warning for g++8
- Fixed a fallthrough on Astar and bdAstar.

#### [pgRouting 2.6.0 Release Notes¶](#)

To see the issues closed by this release see the [Git closed milestone for 2.6.0](#) on Github.

New experimental functions

- `pgr_lineGraphFull`

Bug fixes

- Fix `pgr_trsp(text, integer, double precision, integer, double precision, boolean, boolean[, text])`
  - without restrictions
    - calls `pgr_dijkstra` when both end points have a fraction IN (0,1)
    - calls `pgr_withPoints` when at least one fraction NOT IN (0,1)
  - with restrictions
    - calls original `trsp` code

Internal code

- Cleaned the internal code of `trsp(text, integer, integer, boolean, boolean [, text])`
  - Removed the use of pointers
  - Internal code can accept BIGINT
- Cleaned the internal code of `withPoints`

#### [pgRouting 2.5¶](#)

Contents

- [pgRouting 2.5.5 Release Notes](#)
- [pgRouting 2.5.4 Release Notes](#)
- [pgRouting 2.5.3 Release Notes](#)
- [pgRouting 2.5.2 Release Notes](#)
- [pgRouting 2.5.1 Release Notes](#)
- [pgRouting 2.5.0 Release Notes](#)

#### [pgRouting 2.5.5 Release Notes¶](#)

To see the issues closed by this release see the [Git closed milestone for 2.5.5](#) on Github.

Bug fixes

- Fixes driving distance when vertex is not part of the graph
- Fixes windows test
- Fixes build for python3 and perl5

#### [pgRouting 2.5.4 Release Notes¶](#)

To see the issues closed by this release see the [Git closed milestone for 2.5.4](#) on Github.

- Fixes server crash on several functions.
  - `pgr_floydWarshall`
  - `pgr_johnson`
  - `pgr_aStar`
  - `pgr_bdAstar`
  - `pgr_bdDijkstra`
  - `pgr_alphashape`
  - `pgr_dijkstraCostMatrix`
  - `pgr_dijkstra`
  - `pgr_dijkstraCost`
  - `pgr_drivingDistance`
  - `pgr_KSP`
  - `pgr_dijkstraVia` (proposed)
  - `pgr_boykovKolmogorov` (proposed)
  - `pgr_edgeDisjointPaths` (proposed)
  - `pgr_edmondsKarp` (proposed)
  - `pgr_maxCardinalityMatch` (proposed)
  - `pgr_maxFlow` (proposed)
  - `pgr_withPoints` (proposed)
  - `pgr_withPointsCost` (proposed)
  - `pgr_withPointsKSP` (proposed)
  - `pgr_withPointsDD` (proposed)
  - `pgr_withPointsCostMatrix` (proposed)



- `pgr_contractGraph` (experimental)
- `pgr_pushRelabel` (experimental)
- `pgr_vrpOneDepot` (experimental)
- `pgr_gsoc_vrppdtw` (experimental)
- Fixes for deprecated functions where also applied but not tested
- Removed compilation warning for g++8
- Fixed a fallthrough on Astar and bdAstar.

#### [pgRouting 2.5.3 Release Notes¶](#)

To see the issues closed by this release see the [Git closed milestone for 2.5.3](#) on Github.

Bug fixes

- Fix for postgresql 11: Removed a compilation error when compiling with PostgreSQL

#### [pgRouting 2.5.2 Release Notes¶](#)

To see the issues closed by this release see the [Git closed milestone for 2.5.2](#) on Github.

Bug fixes

- Fix for postgresql 10.1: Removed a compiler condition

#### [pgRouting 2.5.1 Release Notes¶](#)

To see the issues closed by this release see the [Git closed milestone for 2.5.1](#) on Github.

Bug fixes

- Fixed prerequisite minimum version of: cmake

#### [pgRouting 2.5.0 Release Notes¶](#)

To see the issues closed by this release see the [Git closed issues for 2.5.0](#) on Github.

enhancement:

- `pgr_version` is now on SQL language

Breaking change on:

- `pgr_edgeDisjointPaths`:
  - Added `path_id`, `cost` and `agg_cost` columns on the result
  - Parameter names changed
  - The many version results are the union of the One to One version

New Signatures

- `pgr_bdAstar(One to One)`

New proposed functions.

- `pgr_bdAstar(One to Many)`
- `pgr_bdAstar(Many to One)`
- `pgr_bdAstar(Many to Many)`
- `pgr_bdAstarCost(One to One)`
- `pgr_bdAstarCost(One to Many)`
- `pgr_bdAstarCost(Many to One)`
- `pgr_bdAstarCost(Many to Many)`
- `pgr_bdAstarCostMatrix`
- `pgr_bdDijkstra(One to Many)`
- `pgr_bdDijkstra(Many to One)`
- `pgr_bdDijkstra(Many to Many)`
- `pgr_bdDijkstraCost(One to One)`
- `pgr_bdDijkstraCost(One to Many)`
- `pgr_bdDijkstraCost(Many to One)`
- `pgr_bdDijkstraCost(Many to Many)`
- `pgr_bdDijkstraCostMatrix`
- `pgr_lineGraph`
- `pgr_lineGraphFull`
- `pgr_connectedComponents`
- `pgr_strongComponents`
- `pgr_biconnectedComponents`
- `pgr_articulationPoints`
- `pgr_bridges`

Deprecated signatures

- `pgr_bdastar` - use `pgr_bdAstar` instead

Renamed functions

- `pgr_maxFlowPushRelabel` - use `pgr_pushRelabel` instead
- `pgr_maxFlowEdmondsKarp` - use `pgr_edmondsKarp` instead

- `pgr_maxFlowBoykovKolmogorov` - use `pgr_boykovKolmogorov` instead
- `pgr_maximumCardinalityMatching` - use `pgr_maxCardinalityMatch` instead

#### Deprecated Function

- `pgr_pointToEdgeNode`

#### [pgRouting 2.4¶](#)

##### Contents

- [pgRouting 2.4.2 Release Notes](#)
- [pgRouting 2.4.1 Release Notes](#)
- [pgRouting 2.4.0 Release Notes](#)

#### [pgRouting 2.4.2 Release Notes¶](#)

To see the issues closed by this release see the [Git closed milestone for 2.4.2](#) on Github.

##### Improvement

- Works for postgresSQL 10

##### Bug fixes

- Fixed: Unexpected error column "cname"
- Replace `__linux__` with `__GLIBC__` for glibc-specific headers and functions

#### [pgRouting 2.4.1 Release Notes¶](#)

To see the issues closed by this release see the [Git closed milestone for 2.4.1](#) on Github.

##### Bug fixes

- Fixed compiling error on macOS
- Condition error on `pgr_withPoints`

#### [pgRouting 2.4.0 Release Notes¶](#)

To see the issues closed by this release see the [Git closed issues for 2.4.0](#) on Github.

##### New Functions

- `pgr_bdDijkstra`

##### New proposed signatures:

- `pgr_maxFlow`
- `pgr_aStar(One to Many)`
- `pgr_aStar(Many to One)`
- `pgr_aStar(Many to Many)`
- `pgr_aStarCost(One to One)`
- `pgr_aStarCost(One to Many)`
- `pgr_aStarCost(Many to One)`
- `pgr_aStarCost(Many to Many)`
- `pgr_aStarCostMatrix`

##### Deprecated signatures.

- `pgr_bddijkstra` - use `pgr_bdDijkstra` instead

##### Deprecated Functions

- `pgr_pointsToVids`

##### Bug fixes

- Bug fixes on proposed functions
  - `pgr_withPointsKSP`: fixed ordering
- TRSP original code is used with no changes on the compilation warnings

#### [pgRouting 2.3¶](#)

#### [pgRouting 2.3.2 Release Notes¶](#)

To see the issues closed by this release see the [Git closed issues for 2.3.2](#) on Github.

##### Bug Fixes

- Fixed `pgr_gsoc_vrppdtw` crash when all orders fit on one truck.
- Fixed `pgr_trsp`:
  - Alternate code is not executed when the point is in reality a vertex
  - Fixed ambiguity on seq

#### [pgRouting 2.3.1 Release Notes¶](#)

To see the issues closed by this release see the [Git closed issues for 2.3.1](#) on Github.

##### Bug Fixes

- Leaks on proposed `max_flow` functions
- Regression error on `pgr_trsp`
- Types discrepancy on `pgr_createVerticesTable`

#### [pgRouting 2.3.0 Release Notes¶](#)

To see the issues closed by this release see the[Git closed issues for 2.3.0](#) on Github.

#### New Signatures

- pgr\_TSP
- pgr\_aStar

#### New Functions

- pgr\_eucledianTSP

#### New proposed functions.

- pgr\_dijkstraCostMatrix
- pgr\_withPointsCostMatrix
- pgr\_maxFlowPushRelabel(One to One)
- pgr\_maxFlowPushRelabel(One to Many)
- pgr\_maxFlowPushRelabel(Many to One)
- pgr\_maxFlowPushRelabel(Many to Many)
- pgr\_maxFlowEdmondsKarp(One to One)
- pgr\_maxFlowEdmondsKarp(One to Many)
- pgr\_maxFlowEdmondsKarp(Many to One)
- pgr\_maxFlowEdmondsKarp(Many to Many)
- pgr\_maxFlowBoykovKolmogorov (One to One)
- pgr\_maxFlowBoykovKolmogorov (One to Many)
- pgr\_maxFlowBoykovKolmogorov (Many to One)
- pgr\_maxFlowBoykovKolmogorov (Many to Many)
- pgr\_maximumCardinalityMatching
- pgr\_edgeDisjointPaths(One to One)
- pgr\_edgeDisjointPaths(One to Many)
- pgr\_edgeDisjointPaths(Many to One)
- pgr\_edgeDisjointPaths(Many to Many)
- pgr\_contractGraph

#### Deprecated signatures

- pgr\_tsp - use pgr\_TSP or pgr\_eucledianTSP instead
- pgr\_aStar - use pgr\_aStar instead

#### Deprecated Functions

- pgr\_flip\_edges
- pgr\_vidsToDmatrix
- pgr\_pointsToDMatrix
- pgr\_textToPoints

#### [pgRouting 2.2.4](#)

##### Contents

- [pgRouting 2.2.4 Release Notes](#)
- [pgRouting 2.2.3 Release Notes](#)
- [pgRouting 2.2.2 Release Notes](#)
- [pgRouting 2.2.1 Release Notes](#)
- [pgRouting 2.2.0 Release Notes](#)

#### [pgRouting 2.2.4 Release Notes](#)

To see the issues closed by this release see the[Git closed issues for 2.2.4](#) on Github.

#### Bug Fixes

- Bogus uses of extern "C"
- Build error on Fedora 24 + GCC 6.0
- Regression error pgr\_nodeNetwork

#### [pgRouting 2.2.3 Release Notes](#)

To see the issues closed by this release see the[Git closed issues for 2.2.3](#) on Github.

#### Bug Fixes

- Fixed compatibility issues with PostgreSQL 9.6.

#### [pgRouting 2.2.2 Release Notes](#)

To see the issues closed by this release see the[Git closed issues for 2.2.2](#) on Github.

#### Bug Fixes

- Fixed regression error on pgr\_drivingDistance

#### [pgRouting 2.2.1 Release Notes](#)

To see the issues closed by this release see the[Git closed issues for 2.2.1](#) on Github.

## Bug Fixes

- Server crash fix on pgr\_alphaShape
- Bug fix on With Points family of functions

## [pgRouting 2.2.0 Release Notes](#)

To see the issues closed by this release see the [Git closed issues for 2.2.0](#) on Github.

## Improvements

- pgr\_nodeNetwork
  - Adding a row\_where and outall optional parameters
- Signature fix
  - pgr\_dijkstra – to match what is documented

## New Functions

- pgr\_floydWarshall
- pgr\_Johnson
- pgr\_dijkstraCost(One to One)
- pgr\_dijkstraCost(One to Many)
- pgr\_dijkstraCost(Many to One)
- pgr\_dijkstraCost(Many to Many)

## Proposed Functionality

- pgr\_withPoints(One to One)
- pgr\_withPoints(One to Many)
- pgr\_withPoints(Many to One)
- pgr\_withPoints(Many to Many)
- pgr\_withPointsCost(One to One)
- pgr\_withPointsCost(One to Many)
- pgr\_withPointsCost(Many to One)
- pgr\_withPointsCost(Many to Many)
- pgr\_withPointsDD(single vertex)
- pgr\_withPointsDD(multiple vertices)
- pgr\_withPointsKSP
- pgr\_dijkstraVia

## Deprecated Functions

- pgr\_apspWarshall use pgr\_floydWarshall instead
- pgr\_apspJohnson use pgr\_Johnson instead
- pgr\_kDijkstraCost use pgr\_dijkstraCost instead
- pgr\_kDijkstraPath use pgr\_dijkstra instead

## Renamed and Deprecated Function

- pgr\_makeDistanceMatrix renamed to \_pgr\_makeDistanceMatrix

## [pgRouting 2.1](#)

## Contents

- [pgRouting 2.1.0 Release Notes](#)

## [pgRouting 2.1.0 Release Notes](#)

To see the issues closed by this release see the [Git closed issues for 2.1.0](#) on Github.

## New Signatures

- pgr\_dijkstra(One to Many)
- pgr\_dijkstra(Many to One)
- pgr\_dijkstra(Many to Many)
- pgr\_drivingDistance(multiple vertices)

## Refactored

- pgr\_dijkstra(One to One)
- pgr\_ksp
- pgr\_drivingDistance(single vertex)

## Improvements

- pgr\_alphaShape function now can generate better (multi)polygon with holes and alpha parameter.

## Proposed Functionality

- Proposed functions from Steve Woodbridge, (Classified as Convenience by the author.)
  - pgr\_pointToEdgeNode - convert a point geometry to a vertex\_id based on closest edge.
  - pgr\_flipEdges - flip the edges in an array of geometries so the connect end to end.
  - pgr\_textToPoints - convert a string of x,y;x,y;... locations into point geometries.
  - pgr\_pointsToVids - convert an array of point geometries into vertex ids.

- `pgr_pointsToDMatrix` - Create a distance matrix from an array of points.
- `pgr_vidsToDMatrix` - Create a distance matrix from an array of `vertex_id`.
- `pgr_vidsToDMatrix` - Create a distance matrix from an array of `vertex_id`.
- Added proposed functions from GSoc Projects:
  - `pgr_vrppdtw`
  - `pgr_vrponedepot`

#### Deprecated Functions

- `pgr_getColumnName`
- `pgr_getTableName`
- `pgr_isColumnCndexed`
- `pgr_isColumnInTable`
- `pgr_quote_ident`
- `pgr_versionless`
- `pgr_startPoint`
- `pgr_endPoint`
- `pgr_pointTold`

#### No longer supported

- Removed the 1.x legacy functions

#### Bug Fixes

- Some bug fixes in other functions

#### Refactoring Internal Code

- A C and C++ library for developer was created
  - encapsulates postgresSQL related functions
  - encapsulates Boost.Graph graphs
    - Directed Boost.Graph
    - Undirected Boost.graph.
  - allow any-integer in the id's
  - allow any-numerical on the `cost/reverse_cost` columns
- Instead of generating many libraries: - All functions are encapsulated in one library - The library has the prefix 2-1-0

#### [pgRouting 2.0¶](#)

#### Contents

- [pgRouting 2.0.1 Release Notes](#)
- [pgRouting 2.0.0 Release Notes](#)

#### [pgRouting 2.0.1 Release Notes¶](#)

Minor bug fixes.

#### Bug Fixes

- No track of the bug fixes were kept.

#### [pgRouting 2.0.0 Release Notes¶](#)

To see the issues closed by this release see the [Git closed issues for 2.0.0](#) on Github.

With the release of pgRouting 2.0.0 the library has abandoned backwards compatibility to [pgRouting 1.0](#) releases. The main Goals for this release are:

- Major restructuring of pgRouting.
- Standardization of the function naming
- Preparation of the project for future development.

As a result of this effort:

- pgRouting has a simplified structure
- Significant new functionality has being added
- Documentation has being integrated
- Testing has being integrated
- And made it easier for multiple developers to make contributions.

#### Important Changes

- Graph Analytics - tools for detecting and fixing connection some problems in a graph
- A collection of useful utility functions
- Two new All Pairs Short Path algorithms (`pgr_apspJohnson`, `pgr_apspWarshall`)
- Bi-directional Dijkstra and A-star search algorithms (`pgr_bdAstar`, `pgr_bdDijkstra`)
- One to many nodes search (`pgr_kDijkstra`)
- K alternate paths shortest path (`pgr_ksp`)
- New TSP solver that simplifies the code and the build process (`pgr_tsp`), dropped "Gaul Library" dependency
- Turn Restricted shortest path (`pgr_trsp`) that replaces Shooting Star
- Dropped support for Shooting Star

- Built a test infrastructure that is run before major code changes are checked in
- Tested and fixed most all of the outstanding bugs reported against 1.x that existing in the 2.0-dev code base.
- Improved build process for Windows
- Automated testing on Linux and Windows platforms trigger by every commit
- Modular library design
- Compatibility with PostgreSQL 9.1 or newer
- Compatibility with PostGIS 2.0 or newer
- Installs as PostgreSQL EXTENSION
- Return types re factored and unified
- Support for table SCHEMA in function parameters
- Support for st\_ PostGIS function prefix
- Added pgr\_ prefix to functions and types
- Better documentation: <https://docs.pgrouting.org>
- shooting\_star is discontinued

## [pgRouting 1¶](#)

### [pgRouting 1.0¶](#)

#### Contents

- [Changes for release 1.05](#)
- [Changes for release 1.03](#)
- [Changes for release 1.02](#)
- [Changes for release 1.01](#)
- [Changes for release 1.0](#)
- [Changes for release 1.0.0b](#)
- [Changes for release 1.0.0a](#)
- [Changes for release 0.9.9](#)
- [Changes for release 0.9.8](#)

To see the issues closed by this release see the [Git closed issues for 1.x](#) on Github. The following release notes have been copied from the previous RELEASE\_NOTES file and are kept as a reference.

#### [Changes for release 1.05¶](#)

- Bug fixes

#### [Changes for release 1.03¶](#)

- Much faster topology creation
- Bug fixes

#### [Changes for release 1.02¶](#)

- Shooting\* bug fixes
- Compilation problems solved

#### [Changes for release 1.01¶](#)

- Shooting\* bug fixes

#### [Changes for release 1.0¶](#)

- Core and extra functions are separated
- Cmake build process
- Bug fixes

#### [Changes for release 1.0.0b¶](#)

- Additional SQL file with more simple names for wrapper functions
- Bug fixes

#### [Changes for release 1.0.0a¶](#)

- Shooting\* shortest path algorithm for real road networks
- Several SQL bugs were fixed

#### [Changes for release 0.9.9¶](#)

- PostgreSQL 8.2 support
- Shortest path functions return empty result if they could not find any path

#### [Changes for release 0.9.8¶](#)

- Renumbering scheme was added to shortest path functions
- Directed shortest path functions were added
- routing\_postgis.sql was modified to use dijkstra in TSP search

## [Migration guide¶](#)

Several functions are having changes on the signatures, and/or have been replaced by new functions.

Results can be different because of the changes.

#### Warning

All deprecated functions will be removed on next major version 4.0.0

#### Contents

- [Migration guide](#)
  - [Migration of pgr\\_alphaShape](#)
  - [Migration of pgr\\_nodeNetwork](#)
  - [Migration of pgr\\_createTopology](#)
  - [Migration of pgr\\_createVerticesTable](#)
  - [Migration of pgr\\_analyzeOneWay](#)
  - [Migration of pgr\\_analyzeGraph](#)
  - [Migration of pgr\\_aStar](#)
  - [Migration of pgr\\_bdAstar](#)
  - [Migration of pgr\\_dijkstra](#)
  - [Migration of pgr\\_drivingDistance](#)
  - [Migration of pgr\\_kruskalDD / pgr\\_kruskalBFS / pgr\\_kruskalDFS](#)
  - [Migration of pgr\\_KSP](#)
  - [Migration of pgr\\_maxCardinalityMatch](#)
  - [Migration of pgr\\_primDD / pgr\\_primBFS / pgr\\_primDFS](#)
  - [Migration of pgr\\_withPointsDD](#)
  - [Migration of pgr\\_withPointsKSP](#)
  - [Migration of pgr\\_trsp \(Vertices\)](#)
  - [Migration of pgr\\_trsp \(Edges\)](#)
  - [Migration of pgr\\_trspViaVertices](#)
  - [Migration of pgr\\_trspViaEdges](#)
  - [Migration of restrictions](#)
  - [See Also](#)

#### [Migration of pgr\\_alphaShape¶](#)

Starting from [v3.8.0](#)

**Before Deprecation:** The following was calculated:

- An alphaShape was calculated

**After Deprecation:**

PostGIS has two ways of generating alphaShape.

If you have SFCGAL, which you can install using

CREATE EXTENSION postgis\_sfcgal

- Since PostGIS 3.5+ use [CG\\_AlphaShape](#)
- For PostGIS 3.5+ use the old name `ST_AlphaShape`

Other PostGIS options are \* [ST\\_ConvexHull](#) \* [ST\\_ConcaveHull](#)

#### [Migration of pgr\\_nodeNetwork¶](#)

Starting from [v3.8.0](#)

**Before Deprecation:** A table with `<edges>_nodded` was created. with split edges.

#### Migration

Use [pgr\\_separateTouching](#) and/or use [pgr\\_separateCrossing](#)

#### [Migration of pgr\\_createTopology¶](#)

Starting from [v3.8.0](#)

**Before Deprecation:** The following was calculated:

- A table with `<edges>_vertices_pgr` was created.

**After Deprecation:** The user is responsible to create the complete topology.

#### [Build a routing topology¶](#)

The basic information to use the majority of the pgRouting functions `id`, `source`, `target`, `cost`, [`reverse_cost`] is what in pgRouting is called the routing topology.

`reverse_cost` is optional but strongly recommended to have in order to reduce the size of the database due to the size of the geometry columns. Having said that, in this documentation `reverse_cost` is used in this documentation.

When the data comes with geometries and there is no routing topology, then this step is needed.

All the start and end vertices of the geometries need an identifier that is to be stored in `source` and `target` columns of the table of the data. Likewise, `cost` and `reverse_cost` need to have the value of traversing the edge in both directions.

If the columns do not exist they need to be added to the table in question. (see [ALTER TABLE](#))

The function [pgr\\_extractVertices](#) is used to create a vertices table based on the edge identifier and the geometry of the edge of the graph.

```
SELECT * INTO vertices
FROM pgr_extractVertices('SELECT id, geom FROM edges ORDER BY id');
SELECT 18
```

Finally using the data stored on the vertices tables the `source` and `target` are filled up.

```
/* -- set the source information */
UPDATE edges AS e
SET source = v.id, x1 = x, y1 = y
FROM vertices AS v
WHERE ST_StartPoint(e.geom) = v.geom;
UPDATE 24
/* -- set the target information */
UPDATE edges AS e
SET target = v.id, x2 = x, y2 = y
FROM vertices AS v
WHERE ST_EndPoint(e.geom) = v.geom;
UPDATE 24
```

Migration of pgr\_createVerticesTable¶

Starting from [v3.8.0](#)

**Before Deprecation:** The following was calculated:

- A table with `<edges>_vertices_pgr` was created.

**After Deprecation:** The user is responsible to create the vertices table, indexes, etc. They may use [pgr\\_extractVertices](#) for that purpose.

```
SELECT * INTO vertices
FROM pgr_extractVertices('SELECT id, geom FROM edges ORDER BY id');
SELECT 17
```

Migration of pgr\_analyzeOneWay¶

Starting from [v3.8.0](#)

**Before Deprecation:** The following was calculated:

- Number of potential problems in directionality

WHERE

Directionality problems were calculated based on codes.

Dead ends.

A routing problem can arise when from a vertex there is only a way on or a way out but not both:

Either saving or using directly [pgr\\_extractVertices](#) get the dead ends information and determine if the adjacent edge is one way or not.

In this example [pgr\\_extractVertices](#) has already been applied.

```
WITH
deadends AS (
  SELECT (in_edges || out_edges)[1] as id
  FROM vertices where array_length(in_edges || out_edges, 1) = 1)
SELECT * FROM edges JOIN deadends USING (id)
WHERE cost < 0 OR reverse_cost < 0;
id | source | target | cost | reverse_cost | capacity | reverse_capacity | x1 | y1 | x2 | y2 | geom
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
(0 rows)
```

Bridges.

Another routing problem can arise when there is an edge of an undirected graph whose deletion increases its number of connected components, and the bridge is only one way.

To determine if the bridges are or not one way.

```
SELECT id, cost < 0 OR reverse_cost < 0 AS is_OneWayway
FROM pgr_bridges('SELECT id, source, target, cost, reverse_cost FROM edges')
JOIN edges ON (edge = id);
id | is_oneway
----+-----
1 | f
6 | f
7 | f
14 | f
17 | f
18 | f
(6 rows)
```

Migration of pgr\_analyzeGraph¶

Starting from [v3.8.0](#)

**Before Deprecation:** The following was calculated:

- Number of isolated segments.
- Number of dead ends.
- Number of potential gaps found near dead ends.
- Number of intersections. (between 2 edges)

WHERE

Graph component:

A connected subgraph that is not part of any larger connected subgraph.

Isolated segment:

A graph component with only one segment.

Dead ends:

A vertex that participates in only one edge.

gaps:

Space between two geometries.

Intersection:

Is a topological relationship between two geometries.

Migration.

Components.

Instead of counting only isolated segments, determine all the components of the graph.

Depending of the final application requirements use:



- [pgr\\_connectedComponents](#)
- [pgr\\_strongComponents](#)
- [pgr\\_biconnectedComponents](#)

For example:

```
SELECT *
FROM pgr_connectedComponents(
  'SELECT id, source, target, cost, reverse_cost FROM edges'
);
```

seq	component	node
1	1	1
2	1	3
3	1	5
4	1	6
5	1	7
6	1	8
7	1	9
8	1	10
9	1	11
10	1	12
11	1	15
12	1	16
13	1	17
14	2	2
15	2	4
16	13	13
17	13	14

(17 rows)

Dead ends.

Instead of counting the dead ends, determine all the dead ends of the graph using [pgr\\_degree](#).

For example:

```
SELECT *
FROM pgr_degree($$SELECT id, source, target FROM edges$$)
WHERE degree = 1;
```

node	degree
9	1
5	1
4	1
14	1
13	1
2	1
1	1

(7 rows)

Potential gaps near dead ends.

Instead of counting potential gaps between geometries, determine the geometric gaps in the graph using [pgr\\_findCloseEdges](#).

For example:

```
WITH
deadends AS (
  SELECT id, geom, (in_edges || out_edges)[1] as inhere
  FROM vertices where array_length(in_edges || out_edges, 1) = 1),
results AS (
  SELECT (pgr_findCloseEdges('SELECT id, geom FROM edges WHERE id != ' || inhere , geom, 0.001)).*
  FROM deadends)
SELECT d.id, edge_id, distance, st_AsText(geom) AS point, st_asText(edge) edge
FROM results JOIN deadends d USING (geom);
```

id	edge_id	distance	point	edge
4	14	1.00008890058e-12	POINT(1.999999999999 3.5)	LINESTRING(1.99999999999 3.5,2 3.5)

(1 row)

Topological relationships.

Instead of counting intersections, determine topological relationships between geometries.

Several PostGIS functions can be used: [ST\\_Intersects](#), [ST\\_Crosses](#), [ST\\_Overlaps](#), etc.

For example:

```
SELECT e1.id AS id1, e2.id AS id2
FROM edges e1, edges e2 WHERE e1 < e2 AND st_crosses(e1.geom, e2.geom);
```

id1	id2
13	18

(1 row)

#### [Migration of pgr\\_aStar](#)

Starting from [v3.6.0](#)

Signatures to be migrated:

- `pgr_aStar (One to One)`
- `pgr_aStar (One to Many)`
- `pgr_aStar (Many to One)`

Before Migration

- Output columns were (seq, path\_seq, [start\_vid], [end\_vid], node, edge, cost, agg\_cost)
  - Depending on the overload used, the columns `start_vid` and `end_vid` might be missing:
    - `pgr_aStar (One to One)` does not have `start_vid` and `end_vid`.
    - `pgr_aStar (One to Many)` does not have `start_vid`.
    - `pgr_aStar (Many to One)` does not have `end_vid`.

Migration:

- Be aware of the existence of the additional columns.
- In `pgr_aStar (One to One)`
  - `start_vid` contains the **start vid** parameter value.

- end\_vid contains the **end vid** parameter value.

```
SELECT * FROM pgr_aStar(
  $$SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2 FROM edges$$,
  6, 10);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
2 | 2 | 6 | 10 | 7 | 8 | 1 | 1
3 | 3 | 6 | 10 | 11 | 9 | 1 | 2
4 | 4 | 6 | 10 | 16 | 16 | 1 | 3
5 | 5 | 6 | 10 | 15 | 3 | 1 | 4
6 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(6 rows)
```

- In pgr\_aStar (*One to Many*)

- start\_vid contains the **start vid** parameter value.

```
SELECT * FROM pgr_aStar(
  $$SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2 FROM edges$$,
  6, ARRAY[3, 10]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 3 | 6 | 4 | 1 | 0
2 | 2 | 6 | 3 | 7 | 7 | 1 | 1
3 | 3 | 6 | 3 | 3 | -1 | 0 | 2
4 | 1 | 6 | 10 | 6 | 4 | 1 | 0
5 | 2 | 6 | 10 | 7 | 8 | 1 | 1
6 | 3 | 6 | 10 | 11 | 9 | 1 | 2
7 | 4 | 6 | 10 | 16 | 16 | 1 | 3
8 | 5 | 6 | 10 | 15 | 3 | 1 | 4
9 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(9 rows)
```

- In pgr\_aStar (*Many to One*)

- end\_vid contains the **end vid** parameter value.

```
SELECT * FROM pgr_aStar(
  $$SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2 FROM edges$$,
  ARRAY[3, 6], 10);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 3 | 10 | 3 | 7 | 1 | 0
2 | 2 | 3 | 10 | 7 | 8 | 1 | 1
3 | 3 | 3 | 10 | 11 | 9 | 1 | 2
4 | 4 | 3 | 10 | 16 | 16 | 1 | 3
5 | 5 | 3 | 10 | 15 | 3 | 1 | 4
6 | 6 | 3 | 10 | 10 | -1 | 0 | 5
7 | 1 | 6 | 10 | 6 | 4 | 1 | 0
8 | 2 | 6 | 10 | 7 | 8 | 1 | 1
9 | 3 | 6 | 10 | 11 | 9 | 1 | 2
10 | 4 | 6 | 10 | 16 | 16 | 1 | 3
11 | 5 | 6 | 10 | 15 | 3 | 1 | 4
12 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(12 rows)
```

- If needed filter out the added columns, for example:

```
SELECT seq, path_seq, node, edge, cost, agg_cost FROM pgr_aStar(
  $$SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2 FROM edges$$,
  6, 10);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 4 | 1 | 0
2 | 2 | 7 | 8 | 1 | 1
3 | 3 | 11 | 9 | 1 | 2
4 | 4 | 16 | 16 | 1 | 3
5 | 5 | 15 | 3 | 1 | 4
6 | 6 | 10 | -1 | 0 | 5
(6 rows)
```

- If needed add the new columns, similar to the following example where `pgr_dijkstra` is used, and the function had to be modified to be able to return the new columns:

- In [v3.0](#) the function `my_dijkstra` uses `pgr_dijkstra`.
- Starting from [v3.5](#) the function `my_dijkstra` returns the new additional columns of `pgr_dijkstra`.

#### Migration of pgr\_bdAstar

Starting from [v3.6.0](#)

Signatures to be migrated:

- pgr\_bdAstar (*One to One*)
- pgr\_bdAstar (*One to Many*)
- pgr\_bdAstar (*Many to One*)

Before Migration:

- Output columns were (seq, path\_seq, [start\_vid], [end\_vid], node, edge, cost, agg\_cost)
  - Depending on the overload used, the columns `start_vid` and `end_vid` might be missing:
    - pgr\_bdAstar (*One to One*) does not have `start_vid` and `end_vid`.
    - pgr\_bdAstar (*One to Many*) does not have `start_vid`.
    - pgr\_bdAstar (*Many to One*) does not have `end_vid`.

Migration:

- Be aware of the existence of the additional columns.
- In pgr\_bdAstar (*One to One*)
  - start\_vid contains the **start vid** parameter value.
  - end\_vid contains the **end vid** parameter value.

```
SELECT * FROM pgr_bdAstar(
  $$SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2 FROM edges$$,
  6, 10);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 10 | 6 | 4 | 1 | 0
```

2	2	6	10	7	8	1	1
3	3	6	10	11	9	1	2
4	4	6	10	16	16	1	3
5	5	6	10	15	3	1	4
6	6	6	10	10	-1	0	5

(6 rows)

- In `pgr_bdAstar` (*One to Many*)
  - `start_vid` contains the **start vid** parameter value.

```
SELECT * FROM pgr_bdAstar(
  $$SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2 FROM edges$$,
  6, ARRAY[3, 10]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1	1	6	3	6	4	1	0
2	2	6	3	7	7	1	1
3	3	6	3	3	-1	0	2
4	1	6	10	6	4	1	0
5	2	6	10	7	8	1	1
6	3	6	10	11	9	1	2
7	4	6	10	16	16	1	3
8	5	6	10	15	3	1	4
9	6	6	10	10	-1	0	5

(9 rows)

- In `pgr_bdAstar` (*Many to One*)
  - `end_vid` contains the **end vid** parameter value.

```
SELECT * FROM pgr_bdAstar(
  $$SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2 FROM edges$$,
  ARRAY[3, 6], 10);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1	1	3	10	3	7	1	0
2	2	3	10	7	8	1	1
3	3	3	10	11	9	1	2
4	4	3	10	16	16	1	3
5	5	3	10	15	3	1	4
6	6	3	10	10	-1	0	5
7	1	6	10	6	4	1	0
8	2	6	10	7	8	1	1
9	3	6	10	11	9	1	2
10	4	6	10	16	16	1	3
11	5	6	10	15	3	1	4
12	6	6	10	10	-1	0	5

(12 rows)

- If needed filter out the added columns, for example:

```
SELECT seq, path_seq, node, edge, cost, agg_cost FROM pgr_bdAstar(
  $$SELECT id, source, target, cost, reverse_cost, x1, y1, x2, y2 FROM edges$$,
  6, 10);
seq | path_seq | node | edge | cost | agg_cost
```

1	1	6	4	1	0
2	2	7	8	1	1
3	3	11	9	1	2
4	4	16	16	1	3
5	5	15	3	1	4
6	6	10	-1	0	5

(6 rows)

- If needed add the new columns, similar to the following example where `pgr_dijkstra` is used, and the function had to be modified to be able to return the new columns:
  - In [v3.0](#) the function `my_dijkstra` uses `pgr_dijkstra`.
  - Starting from [v3.5](#) the function `my_dijkstra` returns the new additional columns of `pgr_dijkstra`.

#### Migration of `pgr_dijkstra`

Starting from [v3.5.0](#)

Signatures to be migrated:

- `pgr_dijkstra` (*One to One*)
- `pgr_dijkstra` (*One to Many*)
- `pgr_dijkstra` (*Many to One*)

Before Migration:

- Output columns were (seq, path\_seq, [start\_vid], [end\_vid], node, edge, cost, agg\_cost)
  - Depending on the overload used, the columns `start_vid` and `end_vid` might be missing:
    - `pgr_dijkstra` (*One to One*) does not have `start_vid` and `end_vid`.
    - `pgr_dijkstra` (*One to Many*) does not have `start_vid`.
    - `pgr_dijkstra` (*Many to One*) does not have `end_vid`.

Migration:

- Be aware of the existence of the additional columns.
- In `pgr_dijkstra` (*One to One*)
  - `start_vid` contains the **start vid** parameter value.
  - `end_vid` contains the **end vid** parameter value.

```
SELECT * FROM pgr_dijkstra(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  6, 10);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
```

1	1	6	10	6	4	1	0
2	2	6	10	7	8	1	1
3	3	6	10	11	9	1	2
4	4	6	10	16	16	1	3
5	5	6	10	15	3	1	4
6	6	6	10	10	-1	0	5

(6 rows)

- In `pgr_dijkstra (One to Many)`
  - `start_vid` contains the **start vid** parameter value.

```
SELECT * FROM pgr_dijkstra(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  6, ARRAY[3, 10]);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 3 | 6 | 4 | 1 | 0
2 | 2 | 6 | 3 | 7 | 7 | 1 | 1
3 | 3 | 6 | 3 | 3 | -1 | 0 | 2
4 | 1 | 6 | 10 | 6 | 4 | 1 | 0
5 | 2 | 6 | 10 | 7 | 8 | 1 | 1
6 | 3 | 6 | 10 | 11 | 9 | 1 | 2
7 | 4 | 6 | 10 | 16 | 16 | 1 | 3
8 | 5 | 6 | 10 | 15 | 3 | 1 | 4
9 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(9 rows)
```

- In `pgr_dijkstra (Many to One)`
  - `end_vid` contains the **end vid** parameter value.

```
SELECT * FROM pgr_dijkstra(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  ARRAY[3, 6], 10);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 3 | 10 | 3 | 7 | 1 | 0
2 | 2 | 3 | 10 | 7 | 8 | 1 | 1
3 | 3 | 3 | 10 | 11 | 9 | 1 | 2
4 | 4 | 3 | 10 | 16 | 16 | 1 | 3
5 | 5 | 3 | 10 | 15 | 3 | 1 | 4
6 | 6 | 3 | 10 | 10 | -1 | 0 | 5
7 | 1 | 6 | 10 | 6 | 4 | 1 | 0
8 | 2 | 6 | 10 | 7 | 8 | 1 | 1
9 | 3 | 6 | 10 | 11 | 9 | 1 | 2
10 | 4 | 6 | 10 | 16 | 16 | 1 | 3
11 | 5 | 6 | 10 | 15 | 3 | 1 | 4
12 | 6 | 6 | 10 | 10 | -1 | 0 | 5
(12 rows)
```

- If needed filter out the added columns, for example:

```
SELECT seq, path_seq, node, edge, cost, agg_cost FROM pgr_dijkstra(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  6, 10);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 6 | 4 | 1 | 0
2 | 2 | 7 | 8 | 1 | 1
3 | 3 | 11 | 9 | 1 | 2
4 | 4 | 16 | 16 | 1 | 3
5 | 5 | 15 | 3 | 1 | 4
6 | 6 | 10 | -1 | 0 | 5
(6 rows)
```

- If needed add the new columns, for example:
  - In [v3.0](#) the function `my_dijkstra` uses `pgr_dijkstra`.
  - Starting from [v3.5](#) the function `my_dijkstra` returns the new additional columns of `pgr_dijkstra`.

#### [Migration of `pgr\_drivingDistance`](#)

Starting from [v3.6.0 `pgr\_drivingDistance`](#) result columns are being standardized.

from:

```
(seq, [from_v.] node, edge, cost, agg_cost)
```

to:

```
(seq, depth, start_vid, pred, node, edge, cost, agg_cost)
```

Signatures to be migrated:

- `pgr_drivingDistance(Single vertex)`
- `pgr_drivingDistance(Multiple vertices)`

Before Migration:

Output columns were (seq, [from\_v.] node, edge, cost, agg\_cost)

- `pgr_drivingDistance(Single vertex)`
  - Does not have `start_vid` and `depth` result columns.
- `pgr_drivingDistance(Multiple vertices)`
  - Has `from_v` instead of `start_vid` result column.
  - does not have `depth` result column.

Migration:

- Be aware of the existence and name change of the result columns.

`pgr_drivingDistance(Single vertex)`

Using [this](#) example.

- `start_vid` contains the **start vid** parameter value.
- `depth` contains the depth of the node.
- `pred` contains the predecessor of the node.

```
SELECT * FROM pgr_drivingDistance(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  11, 3, 0);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 11 | 11 | 11 | -1 | 0 | 0
2 | 1 | 11 | 11 | 7 | 8 | 1 | 1
3 | 1 | 11 | 11 | 12 | 11 | 1 | 1
4 | 1 | 11 | 11 | 16 | 9 | 1 | 1
5 | 2 | 11 | 7 | 3 | 7 | 1 | 2
```

6	2	11	7	6	4	1	2
7	2	11	7	8	10	1	2
8	2	11	16	15	16	1	2
9	2	11	16	17	15	1	2
10	3	11	3	1	6	1	3
11	3	11	6	5	1	1	3
12	3	11	8	9	14	1	3
13	3	11	15	10	3	1	3

(13 rows)

If needed filter out the added columns, for example, to return the original columns

```
SELECT seq, node, edge, cost, agg_cost
FROM pgr_drivingDistance(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  11, 3.0);
seq | node | edge | cost | agg_cost
```

1	11	-1	0	0
2	7	8	1	1
3	12	11	1	1
4	16	9	1	1
5	3	7	1	2
6	6	4	1	2
7	8	10	1	2
8	15	16	1	2
9	17	15	1	2
10	1	6	1	3
11	5	1	1	3
12	9	14	1	3
13	10	3	1	3

(13 rows)

`pgr_drivingDistance(Multiple vertices)`

Using [this](#) example.

- The `from_v` result column name changes to `start_vid`.
- `depth` contains the depth of the `node`.
- `pred` contains the predecessor of the `node`.

```
SELECT *
FROM pgr_drivingDistance(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  ARRAY[11, 16], 3.0, equicost => true);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
```

1	0	11	11	11	-1	0	0
2	1	11	11	7	8	1	1
3	1	11	11	12	11	1	1
4	2	11	7	3	7	1	2
5	2	11	7	6	4	1	2
6	2	11	7	8	10	1	2
7	3	11	3	1	6	1	3
8	3	11	6	5	1	1	3
9	3	11	8	9	14	1	3
10	0	16	16	16	-1	0	0
11	1	16	16	15	16	1	1
12	1	16	16	17	15	1	1
13	2	16	15	10	3	1	2

(13 rows)

If needed filter out and rename columns, for example, to return the original columns:

```
SELECT seq, start_vid AS from_v, node, edge, cost, agg_cost
FROM pgr_drivingDistance(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  ARRAY[11, 16], 3.0, equicost => true);
seq | from_v | node | edge | cost | agg_cost
```

1	11	11	-1	0	0
2	11	7	8	1	1
3	11	12	11	1	1
4	11	3	7	1	2
5	11	6	4	1	2
6	11	8	10	1	2
7	11	1	6	1	3
8	11	5	1	1	3
9	11	9	14	1	3
10	16	16	-1	0	0
11	16	15	16	1	1
12	16	17	15	1	1
13	16	10	3	1	2

(13 rows)

[Migration of pgr\\_kruskalDD / pgr\\_kruskalBFS / pgr\\_kruskalDFS](#)

Starting from [v3.7.0 pgr\\_kruskalDD](#), [pgr\\_kruskalBFS](#) and [pgr\\_kruskalDFS](#) result columns are being standardized.

from:

```
(seq, depth, start_vid, node, edge, cost, agg_cost)
```

to:

```
(seq, depth, start_vid, pred, node, edge, cost, agg_cost)
```

- `pgr_kruskalDD`
  - Single vertex
  - Multiple vertices
- `pgr_kruskalDFS`
  - Single vertex
  - Multiple vertices
- `pgr_kruskalBFS`
  - Single vertex
  - Multiple vertices

Before Migration:

Output columns were (seq, depth, start\_vid, node, edge, cost, agg\_cost)

- Single vertex and Multiple vertices
  - Do not have *pred* result column.

Migration:

- Be aware of the existence of *pred* result columns.
- If needed filter out the added columns

#### Kruskal single vertex

Using `pgr_KruskalDD` as example. Migration is similar to al the affected functions.

Comparing with [this](#) example.

Now column *pred* exists and contains the predecessor of the *node*.

```
SELECT * FROM pgr_kruskalDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
6, 3.5);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 10 | 15 | 3 | 1 | 2
5 | 3 | 6 | 15 | 16 | 16 | 1 | 3
(5 rows)
```

If needed filter out the added columns, for example, to return the original columns

```
SELECT seq, depth, start_vid, node, edge, cost, agg_cost
FROM pgr_kruskalDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
6, 3.5);
seq | depth | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 15 | 3 | 1 | 2
5 | 3 | 6 | 16 | 16 | 1 | 3
(5 rows)
```

#### Kruskal multiple vertices

Using `pgr_KruskalDD` as example. Migration is similar to al the affected functions.

Comparing with [this](#) example.

Now column *pred* exists and contains the predecessor of the *node*.

```
SELECT * FROM pgr_kruskalDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
ARRAY[9, 6], 3.5);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 10 | 15 | 3 | 1 | 2
5 | 3 | 6 | 15 | 16 | 16 | 1 | 3
6 | 0 | 9 | 9 | 9 | -1 | 0 | 0
7 | 1 | 9 | 9 | 8 | 14 | 1 | 1
8 | 2 | 9 | 8 | 7 | 10 | 1 | 2
9 | 3 | 9 | 7 | 3 | 7 | 1 | 3
10 | 2 | 9 | 8 | 12 | 12 | 1 | 2
11 | 3 | 9 | 12 | 11 | 11 | 1 | 3
12 | 3 | 9 | 12 | 17 | 13 | 1 | 3
(12 rows)
```

If needed filter out the added columns, for example, to return the original columns

```
SELECT seq, depth, start_vid, node, edge, cost, agg_cost
FROM pgr_kruskalDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
ARRAY[9, 6], 3.5);
seq | depth | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 15 | 3 | 1 | 2
5 | 3 | 6 | 16 | 16 | 1 | 3
6 | 0 | 9 | 9 | -1 | 0 | 0
7 | 1 | 9 | 8 | 14 | 1 | 1
8 | 2 | 9 | 7 | 10 | 1 | 2
9 | 3 | 9 | 3 | 7 | 1 | 3
10 | 2 | 9 | 12 | 12 | 1 | 2
11 | 3 | 9 | 11 | 11 | 1 | 3
12 | 3 | 9 | 17 | 13 | 1 | 3
(12 rows)
```

#### Migration of `pgr_KSP`

Starting from [v3.6.0 pgr\\_KSP](#) result columns are being standardized.

from:

```
(seq, path_id, path_seq, node, edge, cost, agg_cost)
```

from:

```
(seq, path_id, path_seq, start_vid, end_vid, node, edge, cost, agg_cost)
```

Signatures to be migrated:

- `pgr_KSP` (One to One)

Before Migration:

- Output columns were (seq, path\_id, path\_seq, node, edge, cost, agg\_cost)
  - the columns *start\_vid* and *end\_vid* do not exist.
    - `pgr_KSP` (One to One) does not have *start\_vid* and *end\_vid*.

Migration:

- Be aware of the existence of the additional columns.

`pgr_KSP (One to One)`

Using [this](#) example.

- `start_vid` contains the **start vid** parameter value.
- `end_vid` contains the **end vid** parameter value.

```
SELECT * FROM pgr_KSP(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  6, 17, 2);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 6 | 17 | 6 | 4 | 1 | 0
2 | 1 | 2 | 6 | 17 | 7 | 10 | 1 | 1
3 | 1 | 3 | 6 | 17 | 8 | 12 | 1 | 2
4 | 1 | 4 | 6 | 17 | 12 | 13 | 1 | 3
5 | 1 | 5 | 6 | 17 | 17 | -1 | 0 | 4
6 | 2 | 1 | 6 | 17 | 6 | 4 | 1 | 0
7 | 2 | 2 | 6 | 17 | 7 | 8 | 1 | 1
8 | 2 | 3 | 6 | 17 | 11 | 9 | 1 | 2
9 | 2 | 4 | 6 | 17 | 16 | 15 | 1 | 3
10 | 2 | 5 | 6 | 17 | 17 | -1 | 0 | 4
(10 rows)
```

If needed filter out the added columns, for example, to return the original columns:

```
SELECT seq, path_id, path_seq, node, edge, cost, agg_cost FROM pgr_KSP(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  6, 17, 2);
seq | path_id | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 6 | 4 | 1 | 0
2 | 1 | 2 | 7 | 10 | 1 | 1
3 | 1 | 3 | 8 | 12 | 1 | 2
4 | 1 | 4 | 12 | 13 | 1 | 3
5 | 1 | 5 | 17 | -1 | 0 | 4
6 | 2 | 1 | 6 | 4 | 1 | 0
7 | 2 | 2 | 7 | 8 | 1 | 1
8 | 2 | 3 | 11 | 9 | 1 | 2
9 | 2 | 4 | 16 | 15 | 1 | 3
10 | 2 | 5 | 17 | -1 | 0 | 4
(10 rows)
```

**Migration of `pgr_maxCardinalityMatch`**

`pgr_maxCardinalityMatch` works only for undirected graphs, therefore the `directed` flag has been removed.

Starting from [v3.4.0](#)

Signature to be migrated:

```
pgr_maxCardinalityMatch(Edges SQL, [directed])
RETURNS SETOF (seq, edge, source, target)
```

Migration is needed, because:

- Use `cost` and `reverse_cost` on the inner query
- Results are ordered
- Works for undirected graphs.
- New signature
  - `pgr_maxCardinalityMatch(text)` returns only `edge` column.
  - The optional flag `directed` is removed.

Before migration:

```
SELECT * FROM pgr_maxCardinalityMatch(
  $$SELECT id, source, target, cost AS going, reverse_cost AS coming FROM edges$$,
  directed => true
);
WARNING: pgr_maxCardinalityMatch(text,boolean) deprecated signature on v3.4.0
seq | edge | source | target
-----+-----+-----+-----
1 | 1 | 5 | 6
2 | 5 | 10 | 11
3 | 6 | 1 | 3
4 | 13 | 12 | 17
5 | 14 | 8 | 9
6 | 16 | 15 | 16
7 | 17 | 2 | 4
8 | 18 | 13 | 14
(8 rows)
```

- Columns used are `going` and `coming` to represent the existence of an edge.
- Flag `directed` was used to indicate if it was for **directed** or **undirected** graph.
  - The flag `directed` is ignored.
    - Regardless of it's value it gives the result considering the graph as **undirected**.

Migration:

- Use the columns `cost` and `reverse_cost` to represent the existence of an edge.
- Do not use the flag `directed`.
- In the query returns only `edge` column.

```
SELECT * FROM pgr_maxCardinalityMatch(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$
);
edge
-----
1
5
6
13
14
16
17
18
(8 rows)
```

Migration of [pgr\\_primDD](#) / [pgr\\_primBFS](#) / [pgr\\_primDFS](#)

Starting from [v3.7.0](#) [pgr\\_primDD](#), [pgr\\_primBFS](#) and [pgr\\_primDFS](#) result columns are being standardized.

from:

(seq, depth, start\_vid, node, edge, cost, agg\_cost)

to:

(seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)

- [pgr\\_primDD](#)
  - Single vertex
  - Multiple vertices
- [pgr\\_primDFS](#)
  - Single vertex
  - Multiple vertices
- [pgr\\_primBFS](#)
  - Single vertex
  - Multiple vertices

Before Migration:

Output columns were (seq, depth, start\_vid, node, edge, cost, agg\_cost)

- Single vertex and Multiple vertices
  - Do not have *pred* result column.

Migration:

- Be aware of the existence of *pred* result columns.
- If needed filter out the added columns

Prim single vertex

Using [pgr\\_primDD](#) as example. Migration is similar to al the affected functions.

Comparing with [this](#) example.

Now column *pred* exists and contains the predecessor of the node.

```
SELECT * FROM pgr_primDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
6, 3.5);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 10 | 15 | 3 | 1 | 2
5 | 2 | 6 | 10 | 11 | 5 | 1 | 2
6 | 3 | 6 | 11 | 16 | 9 | 1 | 3
7 | 3 | 6 | 11 | 12 | 11 | 1 | 3
8 | 1 | 6 | 6 | 7 | 4 | 1 | 1
9 | 2 | 6 | 7 | 3 | 7 | 1 | 2
10 | 3 | 6 | 3 | 1 | 6 | 1 | 3
11 | 2 | 6 | 7 | 8 | 10 | 1 | 2
12 | 3 | 6 | 8 | 9 | 14 | 1 | 3
(12 rows)
```

If needed filter out the added columns, for example, to return the original columns

```
SELECT seq, depth, start_vid, node, edge, cost, agg_cost
FROM pgr_primDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
6, 3.5);
seq | depth | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 15 | 3 | 1 | 2
5 | 2 | 6 | 11 | 5 | 1 | 2
6 | 3 | 6 | 16 | 9 | 1 | 3
7 | 3 | 6 | 12 | 11 | 1 | 3
8 | 1 | 6 | 7 | 4 | 1 | 1
9 | 2 | 6 | 3 | 7 | 1 | 2
10 | 3 | 6 | 1 | 6 | 1 | 3
11 | 2 | 6 | 8 | 10 | 1 | 2
12 | 3 | 6 | 9 | 14 | 1 | 3
(12 rows)
```

Prim multiple vertices

Using [pgr\\_primDD](#) as example. Migration is similar to al the affected functions.

Comparing with [this](#) example.

Now column *pred* exists and contains the predecessor of the node.

```
SELECT * FROM pgr_primDD(
'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
ARRAY[9, 6], 3.5);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | 6 | 6 | 6 | -1 | 0 | 0
2 | 1 | 6 | 6 | 5 | 1 | 1 | 1
3 | 1 | 6 | 6 | 10 | 2 | 1 | 1
4 | 2 | 6 | 10 | 15 | 3 | 1 | 2
5 | 2 | 6 | 10 | 11 | 5 | 1 | 2
6 | 3 | 6 | 11 | 16 | 9 | 1 | 3
7 | 3 | 6 | 11 | 12 | 11 | 1 | 3
8 | 1 | 6 | 6 | 7 | 4 | 1 | 1
9 | 2 | 6 | 7 | 3 | 7 | 1 | 2
10 | 3 | 6 | 3 | 1 | 6 | 1 | 3
11 | 2 | 6 | 7 | 8 | 10 | 1 | 2
12 | 3 | 6 | 8 | 9 | 14 | 1 | 3
13 | 0 | 9 | 9 | 9 | -1 | 0 | 0
14 | 1 | 9 | 9 | 8 | 14 | 1 | 1
15 | 2 | 9 | 8 | 7 | 10 | 1 | 2
```



16		3		9		7		6		4		1		3
17		3		9		7		3		7		1		3
(17 rows)														

If needed filter out the added columns, for example, to return the original columns

```
SELECT seq, depth, start_vid, node, edge, cost, agg_cost
FROM pgr_primDD(
  'SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id',
  ARRAY[9, 6], 3.5);
seq | depth | start_vid | node | edge | cost | agg_cost
```

1		0		6		6		-1		0		0
2		1		6		5		1		1		1
3		1		6		10		2		1		1
4		2		6		15		3		1		2
5		2		6		11		5		1		2
6		3		6		16		9		1		3
7		3		6		12		11		1		3
8		1		6		7		4		1		1
9		2		6		3		7		1		2
10		3		6		1		6		1		3
11		2		6		8		10		1		2
12		3		6		9		14		1		3
13		0		9		9		-1		0		0
14		1		9		8		14		1		1
15		2		9		7		10		1		2
16		3		9		6		4		1		3
17		3		9		3		7		1		3
(17 rows)												

Migration of pgr\_withPointsDD¶

Starting from [v3.6.0 pgr\\_withPointsDD - Proposed](#) result columns are being standardized.

from:

```
(seq, [start_vid], node, edge, cost, agg_cost)
```

to:

```
(seq, depth, start_vid, pred, node, edge, cost, agg_cost)
```

And driving\_side parameter changed from named optional to unnamed compulsory **driving\_side** and its validity differ for directed and undirected graphs.

Signatures to be migrated:

- pgr\_withPointsDD (Single vertex)
- pgr\_withPointsDD (Multiple vertices)

Before Migration:

- pgr\_withPointsDD (Single vertex)
  - Output columns were (seq, node, edge, cost, agg\_cost)
  - Does not have start\_vid, pred and depth result columns.
  - driving\_side parameter was named optional now it is compulsory unnamed.
- pgr\_withPointsDD (Multiple vertices)
  - Output columns were (seq, start\_vid, node, edge, cost, agg\_cost)
  - Does not have depth and pred result columns.
  - driving\_side parameter was named optional now it is compulsory unnamed.

Driving side was optional

The default values on this query are:

directed:

```
true
```

driving\_side:

```
'b'
```

details:

```
false
```

```
SELECT * FROM pgr_withPointsDD(
  $$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
  $$SELECT pid, edge_id, fraction, side from pointsOfInterest$$,
  -1, 3.3);
WARNING:  pgr_withpointsdd(text,text,bigint,double precision,boolean,character,boolean) deprecated signature on 3.6.0
seq | node | edge | cost | agg_cost
```

1		-1		-1		0		0
2		5		1		0.4		0.4
3		6		1		0.6		0.6
4		7		4		1		1.6
5		3		7		1		2.6
6		8		10		1		2.6
7		11		8		1		2.6
8		-3		12		0.6		3.2
9		-4		6		0.7		3.3
(9 rows)								

Driving side was named optional

The default values on this query are:

directed:

```
true
```

details:

```
false
```

```
SELECT * FROM pgr_withPointsDD(
  $$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
  $$SELECT pid, edge_id, fraction, side from pointsOfInterest$$,
  -1, 3.3, driving_side => 'r');
WARNING:  pgr_withpointsdd(text,text,bigint,double precision,boolean,character,boolean) deprecated signature on 3.6.0
```

seq	node	edge	cost	agg_cost
1	-1	-1	0	0
2	5	1	0.4	0.4
3	6	1	1	1.4
4	7	4	1	2.4

(4 rows)

On directed graph `b` could be used as **driving side**

The default values on this query are:

details:

false				
SELECT * FROM pgr_withPointsDD( \$\$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id\$\$, \$\$SELECT pid, edge_id, fraction, side from pointsOfInterest\$\$, -1, 3.3, directed => true, driving_side => 'b'); WARNING: pgr_withpointsdd(text,text,bigint,double precision,boolean,character,boolean) deprecated signature on 3.6.0				
seq	node	edge	cost	agg_cost
1	-1	-1	0	0
2	5	1	0.4	0.4
3	6	1	0.6	0.6
4	7	4	1	1.6
5	3	7	1	2.6
6	8	10	1	2.6
7	11	8	1	2.6
8	-3	12	0.6	3.2
9	-4	6	0.7	3.3

(9 rows)

On undirected graph `r` could be used as **driving side**

Also `l` could be used as **driving side**

```
SELECT * FROM pgr_withPointsDD(  
  $$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,  
  $$SELECT pid, edge_id, fraction, side from pointsOfInterest$$,  
  -1, 3.3, 'r', directed => true);  
seq | depth | start_vid | pred | node | edge | cost | agg_cost  
-----  
1 | 0 | -1 | -1 | -1 | -1 | 0 | 0  
2 | 1 | -1 | -1 | 5 | 1 | 0.4 | 0.4  
3 | 2 | -1 | 5 | 6 | 1 | 1 | 1.4  
4 | 3 | -1 | -6 | 7 | 4 | 1 | 2.4  
(4 rows)
```

After Migration:

- Be aware of the existence of the additional result Columns.
- New output columns are (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
- **driving side** parameter is unnamed compulsory, and valid values differ for directed and undirected graphs.
  - Does not have a default value.
  - In directed graph: valid values are [f, R, l, L]
  - In undirected graph: valid values are [b, B]
  - Using an invalid value throws an ERROR.

`pgr_withPointsDD (Single vertex)`

Using [this](#) example.

- (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
- start\_vid contains the **start vid** parameter value.
- depth contains the **depth** from the start\_vid vertex to the node.
- pred contains the predecessor of the node.

To migrate, use an unnamed valid value for **driving side** after the **distance** parameter:

```
SELECT * FROM pgr_withPointsDD(  
  $$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,  
  $$SELECT pid, edge_id, fraction, side from pointsOfInterest$$,  
  -1, 3.3, 'r', directed => true);
```

seq	depth	start_vid	pred	node	edge	cost	agg_cost
1	0	-1	-1	-1	-1	0	0
2	1	-1	-1	5	1	0.4	0.4
3	2	-1	5	6	1	1	1.4
4	3	-1	-6	7	4	1	2.4

(4 rows)

To get results from previous versions:

- filter out the additional columns, for example;
- When details => false to remove the points use WHERE node >= 0 OR cost = 0

SELECT seq, node, edge, cost, agg_cost FROM pgr_withPointsDD( \$\$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id\$\$, \$\$SELECT pid, edge_id, fraction, side from pointsOfInterest\$\$, -1, 3.3, 'r', details => true);				
seq	node	edge	cost	agg_cost
1	-1	-1	0	0
2	5	1	0.4	0.4
3	6	1	1	1.4
4	-6	4	0.7	2.1
5	7	4	0.3	2.4

(5 rows)

`pgr_withPointsDD (Multiple vertices)`

Using [this](#) example.

- (seq, depth, start\_vid, pred, node, edge, cost, agg\_cost)
- depth contains the **depth** from the start\_vid vertex to the node.

- pred contains the predecessor of the node.

```
SELECT * FROM pgr_withPointsDD(
  $$SELECT * FROM edges ORDER BY id$$,
  $$SELECT pid, edge_id, fraction, side from pointsOfInterest$$,
  ARRAY[-1, 16], 3.3, 'T', equicost => true);
seq | depth | start_vid | pred | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 0 | -1 | -1 | -1 | -1 | 0 | 0
2 | 1 | -1 | -1 | 6 | 1 | 0.6 | 0.6
3 | 2 | -1 | 6 | 7 | 4 | 1 | 1.6
4 | 2 | -1 | 6 | 5 | 1 | 1 | 1.6
5 | 3 | -1 | 7 | 3 | 7 | 1 | 2.6
6 | 3 | -1 | 7 | 8 | 10 | 1 | 2.6
7 | 4 | -1 | 8 | -3 | 12 | 0.6 | 3.2
8 | 4 | -1 | 3 | -4 | 6 | 0.7 | 3.3
9 | 0 | 16 | 16 | 16 | -1 | 0 | 0
10 | 1 | 16 | 16 | 11 | 9 | 1 | 1
11 | 1 | 16 | 16 | 15 | 16 | 1 | 1
12 | 1 | 16 | 16 | 17 | 15 | 1 | 1
13 | 2 | 16 | 15 | 10 | 3 | 1 | 2
14 | 2 | 16 | 11 | 12 | 11 | 1 | 2
(14 rows)
```

To get results from previous versions:

- Filter out the additional columns
- When details => false to remove the points use WHERE node >= 0 OR cost = 0

```
SELECT seq, start_vid, node, edge, cost, agg_cost FROM pgr_withPointsDD(
  $$SELECT id, source, target, cost, reverse_cost FROM edges ORDER BY id$$,
  $$SELECT pid, edge_id, fraction, side from pointsOfInterest$$,
  ARRAY[-1, 16], 3.3, 'T', equicost => true) WHERE node >= 0 OR cost = 0;
seq | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | -1 | -1 | -1 | 0 | 0
2 | -1 | 6 | 1 | 0.6 | 0.6
3 | -1 | 7 | 4 | 1 | 1.6
4 | -1 | 5 | 1 | 1 | 1.6
5 | -1 | 3 | 7 | 1 | 2.6
6 | -1 | 8 | 10 | 1 | 2.6
9 | 16 | 16 | -1 | 0 | 0
10 | 16 | 11 | 9 | 1 | 1
11 | 16 | 15 | 16 | 1 | 1
12 | 16 | 17 | 15 | 1 | 1
13 | 16 | 10 | 3 | 1 | 2
14 | 16 | 12 | 11 | 1 | 2
(12 rows)
```

#### Migration of pgr\_withPointsKSP

Starting from [v3.6.0 pgr\\_withPointsKSP - Proposed](#) result columns are being standardized.

from:

```
(seq, path_id, path_seq, node, edge, cost, agg_cost)
```

from:

```
(seq, path_id, path_seq, start_vid, end_vid, node, edge, cost, agg_cost)
```

And driving side parameter changed from named optional to unnamed compulsory **driving side** and its validity differ for directed and undirected graphs.

Signatures to be migrated:

- pgr\_withPointsKSP (*One to One*)

Before Migration:

- Output columns were (seq, path\_seq, [start\_pid], [end\_pid], node, edge, cost, agg\_cost)
  - the columns start\_vid and end\_vid do not exist.

Migration:

- Be aware of the existence of the additional result Columns.
- New output columns are (seq, path\_id, path\_seq, start\_vid, end\_vid, node, edge, cost, agg\_cost)
- **driving side** parameter is unnamed compulsory, and valid values differ for directed and undirected graphs.
  - Does not have a default value.
  - In directed graph: valid values are [r, R, l, L]
  - In undirected graph: valid values are [b, B]
  - Using an invalid value throws an ERROR.

#### pgr\_withPointsKSP (*One to One*)

Using [this](#) example.

- start\_vid contains the **start vid** parameter value.
- end\_vid contains the **end vid** parameter value.

```
SELECT * FROM pgr_withPointsKSP(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT pid, edge_id, fraction, side from pointsOfInterest$$,
  -1, -2, 'T');
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | -1 | -2 | -1 | 1 | 0.6 | 0
2 | 1 | 2 | 1 | -1 | -2 | 6 | 4 | 1 | 0.6
3 | 1 | 3 | 1 | -1 | -2 | 7 | 8 | 1 | 1.6
4 | 1 | 4 | 1 | -1 | -2 | 11 | 11 | 1 | 2.6
5 | 1 | 5 | 1 | -1 | -2 | 12 | 13 | 1 | 3.6
6 | 1 | 6 | 1 | -1 | -2 | 17 | 15 | 0.6 | 4.6
7 | 1 | 7 | 1 | -1 | -2 | -2 | -1 | 0 | 5.2
8 | 2 | 1 | 1 | -1 | -2 | -1 | 1 | 0.6 | 0
9 | 2 | 2 | 1 | -1 | -2 | 6 | 4 | 1 | 0.6
10 | 2 | 3 | 1 | -1 | -2 | 7 | 8 | 1 | 1.6
11 | 2 | 4 | 1 | -1 | -2 | 11 | 9 | 1 | 2.6
12 | 2 | 5 | 1 | -1 | -2 | 16 | 15 | 1.6 | 3.6
13 | 2 | 6 | 1 | -1 | -2 | -2 | -1 | 0 | 5.2
(13 rows)
```

If needed filter out the additional columns, for example, to return the original columns:

```
SELECT seq, path_id, path_seq, node, edge, cost, agg_cost FROM pgr_withPointsKSP(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT pid, edge_id, fraction, side from pointsOfInterest$$,
  -1, -2, 2, 7);
seq | path_id | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 1 | -1 | 1 | 0.6 | 0
2 | 1 | 2 | 6 | 4 | 1 | 0.6 | 0.6
3 | 1 | 3 | 7 | 8 | 1 | 1.6 | 1.6
4 | 1 | 4 | 11 | 11 | 1 | 2.6 | 2.6
5 | 1 | 5 | 12 | 13 | 1 | 3.6 | 3.6
6 | 1 | 6 | 17 | 15 | 0.6 | 4.6
7 | 1 | 7 | -2 | -1 | 0 | 5.2 | 5.2
8 | 2 | 1 | -1 | 1 | 0.6 | 0
9 | 2 | 2 | 6 | 4 | 1 | 0.6 | 0.6
10 | 2 | 3 | 7 | 8 | 1 | 1.6 | 1.6
11 | 2 | 4 | 11 | 9 | 1 | 2.6 | 2.6
12 | 2 | 5 | 16 | 15 | 1.6 | 3.6
13 | 2 | 6 | -2 | -1 | 0 | 5.2 | 5.2
(13 rows)
```

[Migration of pgr\\_trsp \(Vertices\)](#)

Signature:

pgr\_trsp(Edges SQL, source, target, directed boolean, has\_rcost boolean  
[,restrict\_sq text]);

RETURNS SETOF (seq, id1, id2, cost)

Deprecated:

- [v3.4.0](#)
- [Use pgr\\_dijkstra when there are no restrictions.](#)
- [Use pgr\\_trsp when there are restrictions.](#)

See Also

- [pgr\\_dijkstra](#)
- [pgr\\_trsp - Proposed](#)
- [Migration of restrictions](#)

[Use pgr\\_dijkstra when there are no restrictions.](#)

Use [pgr\\_dijkstra](#) instead.

```
SELECT * FROM pgr_dijkstra(
  $$SELECT id, source, target, cost, reverse_cost
  FROM edges WHERE id != 16$$,
  15, 16);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 15 | 16 | 15 | 3 | 1 | 0
2 | 2 | 15 | 16 | 10 | 5 | 1 | 1
3 | 3 | 15 | 16 | 11 | 9 | 1 | 2
4 | 4 | 15 | 16 | 16 | -1 | 0 | 3
(4 rows)
```

To get the original column names:

```
SELECT seq, node::INTEGER AS id1, edge::INTEGER AS id2, cost
FROM pgr_dijkstra(
  $$SELECT id, source, target, cost, reverse_cost
  FROM edges WHERE id != 16$$,
  15, 16);
seq | id1 | id2 | cost
-----+-----+-----+-----
1 | 15 | 3 | 1
2 | 10 | 5 | 1
3 | 11 | 9 | 1
4 | 16 | -1 | 0
(4 rows)
```

- id1 is the node
- id2 is the edge

[Use pgr\\_trsp when there are restrictions.](#)

Use [pgr\\_trsp - Proposed](#) (One to One) instead.

```
SELECT * FROM pgr_trsp(
  $$SELECT id, source, target, cost, reverse_cost
  FROM edges WHERE id != 16$$,
  $$SELECT * FROM new_restrictions$$,
  15, 16);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 15 | 16 | 15 | 3 | 1 | 0
2 | 2 | 15 | 16 | 10 | 5 | 1 | 1
3 | 3 | 15 | 16 | 11 | 11 | 1 | 2
4 | 4 | 15 | 16 | 12 | 13 | 1 | 3
5 | 5 | 15 | 16 | 17 | 15 | 1 | 4
6 | 6 | 15 | 16 | 16 | -1 | 0 | 5
(6 rows)
```

To get the original column names:

```
SELECT seq, node::INTEGER AS id1, edge::INTEGER AS id2, cost
FROM pgr_trsp(
  $$SELECT id, source, target, cost, reverse_cost
  FROM edges WHERE id != 16$$,
  $$SELECT * FROM new_restrictions$$,
  15, 16);
seq | id1 | id2 | cost
-----+-----+-----+-----
1 | 15 | 3 | 1
2 | 10 | 5 | 1
3 | 11 | 11 | 1
4 | 12 | 13 | 1
5 | 17 | 15 | 1
6 | 16 | -1 | 0
(6 rows)
```

- id1 is the node
- id2 is the edge

Migration of pgr\_trsp (Edges)¶

Signature:

```
pgr_trsp(sql text, source_edge integer, source_pos float8,
        target_edge integer, target_pos float8,
        directed boolean, has_rcost boolean
        [,restrict_sql text]);
RETURNS SETOF (seq, id1, id2, cost)
```

Deprecated:

[v3.4.0](#)

- [Use pgr\\_withPoints when there are no restrictions.](#)
- [Use pgr\\_trsp\\_withPoints when there are restrictions.](#)

See Also

- [pgr\\_withPoints - Proposed](#)
- [pgr\\_trsp\\_withPoints - Proposed](#)
- [Migration of restrictions](#)

Use pgr\_withPoints when there are no restrictions.¶

Use [pgr\\_withPoints - Proposed](#) (One to One) instead.

```
SELECT * FROM pgr_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT pid, edge_id, fraction FROM pointsOfInterest WHERE pid IN (4, 3)$$,
  -4, -3,
  details => false);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -4 | 6 | 0.7 | 0
2 | 2 | 3 | 7 | 1 | 0.7
3 | 3 | 7 | 10 | 1 | 1.7
4 | 4 | 8 | 12 | 0.6 | 2.7
5 | 5 | -3 | -1 | 0 | 3.3
(5 rows)
```

To get the original column names:

```
SELECT seq, node::INTEGER AS id1, edge::INTEGER AS id2, cost
FROM pgr_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT * FROM (VALUES (1, 6, 0.3),(2, 12, 0.6)) AS t(pid, edge_id, fraction)$$,
  -1, -2,
  details => false);
seq | id1 | id2 | cost
-----+-----+-----+-----
1 | -1 | 6 | 0.7
2 | 3 | 7 | 1
3 | 7 | 10 | 1
4 | 8 | 12 | 0.6
5 | -2 | -1 | 0
(5 rows)
```

- id1 is the node
- id2 is the edge

Use pgr\_trsp\_withPoints when there are restrictions.¶

Use [pgr\\_trsp\\_withPoints - Proposed](#) instead.

```
SELECT * FROM pgr_trsp_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT * FROM new_restrictions$$,
  $$SELECT pid, edge_id, fraction FROM pointsOfInterest WHERE pid IN (4, 3)$$,
  -4, -3,
  details => false);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -4 | -3 | -4 | 6 | 0.7 | 0
2 | 2 | 2 | -4 | -3 | 3 | 7 | 1 | 0.7
3 | 3 | 3 | -4 | -3 | 7 | 8 | 1 | 1.7
4 | 4 | 4 | -4 | -3 | 11 | 9 | 1 | 2.7
5 | 5 | 5 | -4 | -3 | 16 | 16 | 1 | 3.7
6 | 6 | 6 | -4 | -3 | 15 | 3 | 1 | 4.7
7 | 7 | 7 | -4 | -3 | 10 | 2 | 1 | 5.7
8 | 8 | 8 | -4 | -3 | 6 | 4 | 1 | 6.7
9 | 9 | 9 | -4 | -3 | 7 | 10 | 1 | 7.7
10 | 10 | 10 | -4 | -3 | 8 | 12 | 0.6 | 8.7
11 | 11 | 11 | -4 | -3 | -3 | -1 | 0 | 9.3
(11 rows)
```

To get the original column names:

```
SELECT seq, node::INTEGER AS id1, edge::INTEGER AS id2, cost
FROM pgr_trsp_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT * FROM new_restrictions$$,
  $$SELECT * FROM (VALUES (1, 6, 0.3),(2, 12, 0.6)) AS t(pid, edge_id, fraction)$$,
  -1, -2,
  details => false)
WHERE edge != -1;
seq | id1 | id2 | cost
-----+-----+-----+-----
1 | -1 | 6 | 0.7
2 | 3 | 7 | 1
3 | 7 | 8 | 1
4 | 11 | 9 | 1
5 | 16 | 16 | 1
6 | 15 | 3 | 1
7 | 10 | 2 | 1
8 | 6 | 4 | 1
9 | 7 | 10 | 1
10 | 8 | 12 | 0.6
(10 rows)
```

- id1 is the node
- id2 is the edge

Migration of pgr\_trspViaVertices¶

Signature:

```
pgr_trspViaVertices(sql text, vids integer[],
    directed boolean, has_rcost boolean
    [, turn_restrict_sql text]);
RETURNS SETOF (seq, id1, id2, id3, cost)
```

Deprecated:

- [v3.4.0](#)
- [Use pgr\\_dijkstraVia when there are no restrictions](#)
- [Use pgr\\_trspVia when there are restrictions](#)

See Also

- [pgr\\_dijkstraVia - Proposed](#)
- [pgr\\_trspVia - Proposed](#)
- [Migration of restrictions](#)

[Use pgr\\_dijkstraVia when there are no restrictions¶](#)

Use [pgr\\_dijkstraVia - Proposed](#) instead.

```
SELECT * FROM pgr_dijkstraVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  ARRAY[6, 3, 6]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 6 | 3 | 6 | 3 | 6 | 4 | 1 | 0 | 0
2 | 1 | 1 | 6 | 3 | 7 | 7 | 1 | 1 | 1 | 1
3 | 1 | 1 | 6 | 3 | 3 | 3 | -1 | 0 | 2 | 2
4 | 2 | 1 | 3 | 6 | 3 | 7 | 1 | 0 | 2 | 2
5 | 2 | 1 | 3 | 6 | 7 | 4 | 1 | 1 | 1 | 3
6 | 2 | 1 | 3 | 6 | 6 | -2 | 0 | 2 | 2 | 4
(6 rows)
```

To get the original column names:

```
SELECT row_number() over(ORDER BY seq) AS seq,
  path_id::INTEGER AS id1, node::INTEGER AS id2,
  CASE WHEN edge >= 0 THEN edge::INTEGER ELSE -1 END AS id3, cost::FLOAT
FROM pgr_dijkstraVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  ARRAY[6, 3, 6])
WHERE edge != -1;
seq | id1 | id2 | id3 | cost
-----+-----+-----+-----+-----
1 | 1 | 6 | 3 | 6
2 | 1 | 7 | 7 | 1
3 | 1 | 3 | 3 | -1
4 | 2 | 3 | 7 | 1
5 | 2 | 7 | 4 | 1
6 | 2 | 6 | -1 | 0
(5 rows)
```

- id1 is the path identifier
- id2 is the node
- id3 is the edge

[Use pgr\\_trspVia when there are restrictions¶](#)

Use [pgr\\_trspVia - Proposed](#) instead.

```
SELECT * FROM pgr_trspVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT * FROM new_restrictions$$,
  ARRAY[6, 3, 6]);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 6 | 3 | 6 | 3 | 6 | 4 | 1 | 0 | 0
2 | 1 | 1 | 6 | 3 | 7 | 8 | 1 | 1 | 1 | 1
3 | 1 | 1 | 6 | 3 | 11 | 9 | 1 | 2 | 2 | 2
4 | 1 | 1 | 6 | 3 | 16 | 16 | 1 | 3 | 3 | 3
5 | 1 | 1 | 6 | 3 | 15 | 3 | 1 | 4 | 4 | 4
6 | 1 | 1 | 6 | 3 | 10 | 5 | 1 | 5 | 5 | 5
7 | 1 | 1 | 6 | 3 | 11 | 8 | 1 | 6 | 6 | 6
8 | 1 | 1 | 6 | 3 | 7 | 7 | 1 | 7 | 7 | 7
9 | 1 | 1 | 6 | 3 | 3 | -1 | 0 | 8 | 8 | 8
10 | 2 | 1 | 3 | 6 | 3 | 7 | 1 | 0 | 8 | 8
11 | 2 | 1 | 3 | 6 | 7 | 4 | 1 | 1 | 9 | 9
12 | 2 | 1 | 3 | 6 | 6 | -2 | 0 | 2 | 10 | 10
(12 rows)
```

To get the original column names:

```
SELECT row_number() over(ORDER BY seq) AS seq,
  path_id::INTEGER AS id1, node::INTEGER AS id2,
  CASE WHEN edge >= 0 THEN edge::INTEGER ELSE -1 END AS id3, cost::FLOAT
FROM pgr_trspVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT * FROM new_restrictions$$,
  ARRAY[6, 3, 6])
WHERE edge != -1;
seq | id1 | id2 | id3 | cost
-----+-----+-----+-----+-----
1 | 1 | 6 | 3 | 6
2 | 1 | 7 | 8 | 1
3 | 1 | 11 | 9 | 1
4 | 1 | 16 | 16 | 1
5 | 1 | 15 | 3 | 1
6 | 1 | 10 | 5 | 1
7 | 1 | 11 | 8 | 1
8 | 1 | 7 | 7 | 1
9 | 2 | 3 | 7 | 1
10 | 2 | 7 | 4 | 1
11 | 2 | 6 | -1 | 0
(11 rows)
```

- id1 is the path identifier
- id2 is the node
- id3 is the edge

[Migration of pgr\\_trspViaEdges¶](#)

Signature:

```
pgr_trspViaEdges(sql text, eids integer[], pcts float8[],
    directed boolean, has_rcost boolean
    [, turn_restrict_sql text]);
RETURNS SETOF (seq, id1, id2, id3, cost)
```

Deprecated:

- [v3.4.0](#)
- [Use pgr\\_withPointsVia when there are no restrictions](#)
- [Use pgr\\_trspVia\\_withPoints when there are restrictions](#)

See Also

- [pgr\\_withPointsVia - Proposed](#)
- [pgr\\_trspVia\\_withPoints - Proposed](#)
- [Migration of restrictions](#)

[Use pgr\\_withPointsVia when there are no restrictions](#)

Use [pgr\\_withPointsVia - Proposed](#) instead.

```
SELECT * FROM pgr_withPointsVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT pid, edge_id, fraction FROM pointsOfInterest WHERE pid IN (3, 4, 6)$$,
  ARRAY[-4, -3, -6],
  details => false);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -4 | -3 | 4 | 6 | 0.7 | 0 | 0
2 | 1 | 2 | -4 | -3 | 3 | 7 | 1 | 0.7 | 0.7
3 | 1 | 3 | -4 | -3 | 7 | 10 | 1 | 1.7 | 1.7
4 | 1 | 4 | -4 | -3 | 8 | 12 | 0.6 | 2.7 | 2.7
5 | 1 | 5 | -4 | -3 | -1 | 0 | 3.3 | 3.3
6 | 2 | 1 | -3 | -6 | 3 | 12 | 0.4 | 0 | 3.3
7 | 2 | 2 | -3 | -6 | 12 | 13 | 1 | 0.4 | 3.7
8 | 2 | 3 | -3 | -6 | 17 | 15 | 1 | 1.4 | 4.7
9 | 2 | 4 | -3 | -6 | 16 | 9 | 1 | 2.4 | 5.7
10 | 2 | 5 | -3 | -6 | 11 | 8 | 1 | 3.4 | 6.7
11 | 2 | 6 | -3 | -6 | 7 | 4 | 0.3 | 4.4 | 7.7
12 | 2 | 7 | -3 | -6 | -2 | 0 | 4.7 | 8
(12 rows)
```

To get the original column names:

```
SELECT row_number() over(ORDER BY seq) AS seq,
  path_id::INTEGER AS id1, node::INTEGER AS id2,
  CASE WHEN edge >= 0 THEN edge::INTEGER ELSE -1 END AS id3, cost::FLOAT
FROM pgr_withPointsVia(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT pid, edge_id, fraction FROM pointsOfInterest WHERE pid IN (3, 4, 6)$$,
  ARRAY[-4, -3, -6],
  details => false);
seq | id1 | id2 | id3 | cost
-----+-----+-----+-----+-----
1 | 1 | -4 | 6 | 0.7
2 | 1 | 3 | 7 | 1
3 | 1 | 7 | 10 | 1
4 | 1 | 8 | 12 | 0.6
5 | 1 | -3 | -1 | 0
6 | 2 | -3 | 12 | 0.4
7 | 2 | 12 | 13 | 1
8 | 2 | 17 | 15 | 1
9 | 2 | 16 | 9 | 1
10 | 2 | 11 | 8 | 1
11 | 2 | 7 | 4 | 0.3
12 | 2 | -6 | -1 | 0
(12 rows)
```

- id1 is the path identifier
- id2 is the node
- id3 is the edge

[Use pgr\\_trspVia\\_withPoints when there are restrictions](#)

Use [pgr\\_trspVia\\_withPoints - Proposed](#) instead.

```
SELECT * FROM pgr_trspVia_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT * FROM new_restrictions$$,
  $$SELECT pid, edge_id, fraction FROM pointsOfInterest WHERE pid IN (3, 4, 6)$$,
  ARRAY[-4, -3, -6],
  details => false);
seq | path_id | path_seq | start_vid | end_vid | node | edge | cost | agg_cost | route_agg_cost
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | -4 | -3 | 4 | 6 | 0.7 | 0 | 0
2 | 1 | 2 | -4 | -3 | 3 | 7 | 1 | 0.7 | 0.7
3 | 1 | 3 | -4 | -3 | 7 | 4 | 0.6 | 1.7 | 1.7
4 | 1 | 4 | -4 | -3 | 7 | 10 | 1 | 2.3 | 2.3
5 | 1 | 5 | -4 | -3 | 8 | 12 | 0.6 | 3.3 | 3.3
6 | 1 | 6 | -4 | -3 | -1 | 0 | 3.9 | 3.9
7 | 2 | 1 | -3 | -6 | 3 | 12 | 0.4 | 0 | 3.9
8 | 2 | 2 | -3 | -6 | 12 | 13 | 1 | 0.4 | 4.3
9 | 2 | 3 | -3 | -6 | 17 | 15 | 1 | 1.4 | 5.3
10 | 2 | 4 | -3 | -6 | 16 | 9 | 1 | 2.4 | 6.3
11 | 2 | 5 | -3 | -6 | 11 | 8 | 1 | 3.4 | 7.3
12 | 2 | 6 | -3 | -6 | 7 | 4 | 0.3 | 4.4 | 8.3
13 | 2 | 7 | -3 | -6 | -2 | 0 | 4.7 | 8.6
(13 rows)
```

To get the original column names:

```
SELECT row_number() over(ORDER BY seq) AS seq,
  path_id::INTEGER AS id1, node::INTEGER AS id2,
  CASE WHEN edge >= 0 THEN edge::INTEGER ELSE -1 END AS id3, cost::FLOAT
FROM pgr_trspVia_withPoints(
  $$SELECT id, source, target, cost, reverse_cost FROM edges$$,
  $$SELECT * FROM new_restrictions$$,
  $$SELECT * FROM (VALUES (1, 6, 0.3),(2, 12, 0.6),(3, 4, 0.7)) AS t(pid, edge_id, fraction)$$,
  ARRAY[-1, -2, -3],
  details => false);
seq | id1 | id2 | id3 | cost
-----+-----+-----+-----+-----
1 | 1 | -1 | 6 | 0.7
2 | 1 | 3 | 7 | 1
3 | 1 | 7 | 4 | 0.6
```

4		1		7		10		1
5		1		8		12		0.6
6		1		-2		-1		0
7		2		-2		12		0.4
8		2		12		13		1
9		2		17		15		1
10		2		16		9		1
11		2		11		8		1
12		2		7		4		0.3
13		2		-3		-1		0

(13 rows)

- id1 is the path identifier
- id2 is the node
- id3 is the edge

Migration of restrictions¶

Starting from [v3.4.0](#)

The structure of the restrictions have changed:

Old restrictions structure¶

On the deprecated signatures:

- Column rid is ignored
- via\_path
  - Must be in reverse order.
  - Is of type TEXT.
  - When more than one via edge must be separated with,.
- target\_id
  - Is the last edge of the forbidden path.
  - Is of type INTEGER.
- to\_cost
  - Is of type FLOAT.

Creation of the old restrictions table

```
CREATE TABLE old_restrictions (
  rid BIGINT NOT NULL,
  to_cost FLOAT,
  target_id BIGINT,
  via_path TEXT
);
CREATE TABLE
```

Old restrictions fill up

```
INSERT INTO old_restrictions (rid, to_cost, target_id, via_path) VALUES
(1, 100, 7, '4'),
(1, 100, 11, '8'),
(1, 100, 10, '7'),
(2, 4, 9, '5,3'),
(3, 100, 9, '16');
INSERT 0 5
```

Old restrictions contents¶

```
SELECT * FROM old_restrictions;
rid | to_cost | target_id | via_path
-----+-----+-----+-----
1 | 100 | 7 | 4
1 | 100 | 11 | 8
1 | 100 | 10 | 7
2 | 4 | 9 | 5, 3
3 | 100 | 9 | 16
(5 rows)
```

The restriction with rid = 2 is representing  $(3 \rightarrow 5 \rightarrow 9)$

- $(3 \rightarrow 5)$ 
  - is on column via\_path in reverse order
  - is of type TEXT
- $(9)$ 
  - is on column target\_id
  - is of type INTEGER

New restrictions structure¶

- Column id is ignored
- Column path
  - Is of type ARRAY[ANY-INTEGER].
  - Contains all the edges involved on the restriction.
  - The array has the ordered edges of the restriction.
- Column cost
  - Is of type ANY-NUMERICAL

The creation of the restrictions table

```
CREATE TABLE restrictions (
  id SERIAL PRIMARY KEY,
  path BIGINT[],
  cost FLOAT
);
CREATE TABLE
```



Adding the restrictions

```
INSERT INTO restrictions (path, cost) VALUES
(ARRAY[4, 7], 100),
(ARRAY[8, 11], 100),
(ARRAY[7, 10], 100),
(ARRAY[3, 5, 9], 4),
(ARRAY[9, 16], 100);
INSERT 0 5
```

Restrictions data

```
SELECT * FROM restrictions;
id | path | cost
---+-----+-----
1 | {4,7} | 100
2 | {8,11} | 100
3 | {7,10} | 100
4 | {3,5,9} | 4
5 | {9,16} | 100
(5 rows)
```

The restriction with rid = 2 represents the path \((3 \rightarrow 5 \rightarrow 9)\).

- By inspection the path is clear.

Migration

To transform the old restrictions table to the new restrictions structure,

- Create a new table with the new restrictions structure.
  - In this migration guide new\_restrictions is been used.
- For this migration pgRouting supplies an auxiliary function for reversal of an array \_pgr\_array\_reverse needed for the migration.
  - \_pgr\_array\_reverse:
    - Was created temporally for this migration
    - Is not documented.
    - Will be removed on the next mayor version 4.0.0

```
SELECT rid AS id,
_pgr_array_reverse(
array_prepend(target_id, string_to_array(via_path::text, ' '::BIGINT))) AS path,
to_cost AS cost
INTO new_restrictions
FROM old_restrictions;
SELECT 5
```

The migrated table contents:

```
SELECT * FROM new_restrictions;
id | path | cost
---+-----+-----
1 | {4,7} | 100
1 | {8,11} | 100
1 | {7,10} | 100
2 | {3,5,9} | 4
3 | {16,9} | 100
(5 rows)
```

See Also

- [TRSP - Family of functions](#)
- [withPoints - Category](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Indices and tables

- [Index](#)
- [Search Page](#)

Contents