



# **pgRouting Manual**

*Release 2.1.0 (b38118a master)*

**pgRouting Contributors**

July 22, 2017



<b>1</b>	<b>General</b>	<b>3</b>
1.1	Introduction . . . . .	3
1.2	Installation . . . . .	4
1.3	Build Guide . . . . .	6
1.4	Support . . . . .	8
<b>2</b>	<b>Tutorial</b>	<b>11</b>
2.1	Tutorial . . . . .	11
2.2	Sample Data . . . . .	20
2.3	User’s Recipes List . . . . .	21
2.4	How to contribute. . . . .	30
<b>3</b>	<b>Version</b>	<b>33</b>
3.1	pgr_version . . . . .	33
<b>4</b>	<b>Data Types</b>	<b>35</b>
4.1	pgRouting Data Types . . . . .	35
<b>5</b>	<b>Topology functions</b>	<b>39</b>
5.1	Topology Functions . . . . .	39
<b>6</b>	<b>Routing functions</b>	<b>63</b>
6.1	Routing Functions . . . . .	63
<b>7</b>	<b>Pre processing or post processing helping functions</b>	<b>119</b>
<b>8</b>	<b>Experimental and Proposed functions</b>	<b>121</b>
<b>9</b>	<b>Experimental functions: Proposed by Steve Woodbridge</b>	<b>123</b>
9.1	Convenience Functions . . . . .	123
9.2	Distance Matrix Functions . . . . .	123
<b>10</b>	<b>Experimental functions: by GSoC</b>	<b>133</b>
10.1	pgr_vrpOneDepot . . . . .	133
10.2	pgr_vrppdtw - Pickup and Delivery problem . . . . .	133
<b>11</b>	<b>Proposed functions: Proposed by Zia Mohammed</b>	<b>135</b>
11.1	pgr_labelGraph . . . . .	135
11.2	Developers’s Functions . . . . .	137
<b>12</b>	<b>Discontinued Functions</b>	<b>149</b>
12.1	Discontinued Functions . . . . .	149

<b>13 Developer</b>	<b>151</b>
13.1 Developer's Guide . . . . .	151
13.2 Release Notes . . . . .	155

pgRouting extends the [PostGIS](http://postgis.net)<sup>1</sup>/[PostgreSQL](http://postgresql.org)<sup>2</sup> geospatial database to provide geospatial routing and other network analysis functionality.

This is the manual for pgRouting 2.1.0 (b38118a master).



The pgRouting Manual is licensed under a [Creative Commons Attribution-Share Alike 3.0 License](http://creativecommons.org/licenses/by-sa/3.0/)<sup>3</sup>. 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 <http://pgrouting.org>. For other licenses used in pgRouting see the [License](#) page.

---

<sup>1</sup><http://postgis.net>

<sup>2</sup><http://postgresql.org>

<sup>3</sup><http://creativecommons.org/licenses/by-sa/3.0/>



## 1.1 Introduction

pgRouting is an extension of PostGIS<sup>1</sup> and PostgreSQL<sup>2</sup> geospatial database and adds routing and other network analysis functionality. A predecessor of pgRouting – pgDijkstra, written by Sylvain Pasche from Camptocamp<sup>3</sup>, was later extended by Orkney<sup>4</sup> and renamed to pgRouting. The project is now supported and maintained by Georepublic<sup>5</sup>, iMaptools<sup>6</sup> and a broad user community.

pgRouting is an OSGeo Labs<sup>7</sup> project of the OSGeo Foundation<sup>8</sup> and included on OSGeo Live<sup>9</sup>.

### 1.1.1 License

The following licenses can be found in pgRouting:

License	
GNU General Public License, version 2	Most features of pgRouting are available under <a href="#">GNU General Public License, version 2</a> <sup>10</sup> .
Boost Software License - Version 1.0	Some Boost extensions are available under <a href="#">Boost Software License - Version 1.0</a> <sup>11</sup> .
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> <sup>12</sup> .

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

### 1.1.2 Contributors

#### Individuals (in alphabetical order)

Akio Takubo, Anton Patrushev, Ashraf Hossain, Christian Gonzalez, Daniel Kastl, Dave Potts, David Techer, Ema Miyawaki, Florian Thurkow, Frederic Junod, Gerald Fenoy, Jay Mahadeokar, Jinfu

<sup>1</sup><http://postgis.net>

<sup>2</sup><http://postgresql.org>

<sup>3</sup><http://camptocamp.com>

<sup>4</sup><http://www.orkney.co.jp>

<sup>5</sup><http://georepublic.info>

<sup>6</sup><http://imaptools.com/>

<sup>7</sup>[http://wiki.osgeo.org/wiki/OSGeo\\_Labs](http://wiki.osgeo.org/wiki/OSGeo_Labs)

<sup>8</sup><http://osgeo.org>

<sup>9</sup><http://live.osgeo.org/>

<sup>10</sup><http://www.gnu.org/licenses/gpl-2.0.html>

<sup>11</sup>[http://www.boost.org/LICENSE\\_1\\_0.txt](http://www.boost.org/LICENSE_1_0.txt)

<sup>12</sup><http://creativecommons.org/licenses/by-sa/3.0/>

Leng, Kai Behncke, Kishore Kumar, Ko Nagase, Mario Basa, Martin Wiesenhaan, Razequl Islam, Stephen Woodbridge, Sylvain Housseman, Sylvain Pasche, Virginia Vergara

### 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
- iMaptools
- Orkney
- Paragon Corporation

### 1.1.3 Inside



Fig. 1.1: Boost Graph Inside

### 1.1.4 More Information

- The latest software, documentation and news items are available at the pgRouting web site <http://pgrouting.org>.
- PostgreSQL database server at the PostgreSQL main site <http://www.postgresql.org>.
- PostGIS extension at the PostGIS project web site <http://postgis.net>.
- Boost C++ source libraries at <http://www.boost.org>.
- Computational Geometry Algorithms Library (CGAL) at <http://www.cgal.org>.

## 1.2 Installation

This is a basic guide to download and install pgRouting.

---

**Note:** additional notes can be found in [Installation Notes](#)<sup>14</sup>

---

### 1.2.1 Download

Binary packages are provided for the current version on the following platforms:

---

<sup>14</sup><https://github.com/pgRouting/pgrouting/wiki/Notes-on-Download%2C-Installation-and-building-pgRouting>



## Windows

Winnie Bot Builds:

- Winnie PostgreSQL 9.2-9.5 32-bit/64-bit<sup>15</sup>

Production Builds:

- Production builds are part of the Spatial Extensions/PostGIS Bundle available via Application StackBuilder
- Can also get PostGIS Bundle from<sup>16</sup>

## Ubuntu/Debian

Ubuntu packages are available in Launchpad repositories:

- *stable* <https://launchpad.net/~georepublic/+archive/pgrouting>
- *unstable* <https://launchpad.net/~georepublic/+archive/pgrouting-unstable>

```
# Add pgRouting launchpad repository ("stable" or "unstable")
sudo add-apt-repository ppa:georepublic/pgrouting[-unstable]
sudo apt-get update

# Install pgRouting packages
sudo apt-get install postgresql-9.1-pgrouting
```

Use UbuntuGIS-unstable PPA<sup>17</sup> to install PostGIS 2.0.

## RHEL/CentOS/Fedora

- Fedora RPM's: <https://admin.fedoraproject.org/pkgdb/acls/name/pgRouting>

## FreeBSD

pgRouting can be installed via ports:

```
cd /usr/ports/databases/pgRouting
make install clean
```

## OS X

- Homebrew

```
brew install pgrouting
```

## Source Package

You can find all the pgRouting Releases:

<https://github.com/pgRouting/pgrouting/releases>

See *Build Guide* to build the binaries from the source.

<sup>15</sup>[http://postgis.net/windows\\_downloads](http://postgis.net/windows_downloads)

<sup>16</sup><http://download.osgeo.org/postgis/windows/>

<sup>17</sup><https://launchpad.net/~ubuntuugis/+archive/ubuntuugis-unstable>

## Using Git

Git protocol (read-only):

```
git clone git://github.com/pgRouting/pgrouting.git
```

HTTPS protocol (read-only):

```
git clone https://github.com/pgRouting/pgrouting.git
```

See *Build Guide* to build the binaries from the source.

### 1.2.2 Installing in the database

pgRouting is an extension.

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

### 1.2.3 Upgrading the database

To upgrade pgRouting to version 2.1.0 use the following command:

```
ALTER EXTENSION pgrouting UPDATE TO "2.1.0";
```

## 1.3 Build Guide

### 1.3.1 Dependencies

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

- C and C++ compilers
- Postgresql version  $\geq 9.1$
- PostGIS version  $\geq 2.0$
- The Boost Graph Library (BGL). Version  $\geq 1.55$
- CMake  $\geq 2.8.8$
- CGAL  $\geq 4.2$
- (optional, for Documentation) Sphinx  $\geq 1.1$
- (optional, for Documentation as PDF) Latex  $\geq$  [TBD]

### 1.3.2 Configuration

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

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

Create the build directory

```
$ mkdir build
```

To configure:

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

## Configurable variables

The documentation configurable variables are:

**WITH\_DOC** BOOL=OFF – Turn on/off building the documentation

**BUILD\_HTML** BOOL=ON – If WITH\_DOC=ON, turn on/off building HTML

**BUILD\_LATEX** BOOL=OFF – If WITH\_DOC=ON, turn on/off building PDF

**BUILD\_MAN** BOOL=OFF – If WITH\_DOC=ON, turn on/off building MAN pages

Configuring with documentation

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

**Note:** Most of the effort of the documentation has being on the html files.

### 1.3.3 Building

Using make to build the code and the docuemntation

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

```
$ make          # build the code but not the documentation
$ make doc      # build only the documentation
$ make all doc  # build both the code and the documentation
```

### 1.3.4 Installation and reinstallation

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

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

#### For MinGW on Windows

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

#### For Linux

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

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

## 1.4 Support

pgRouting community support is available through [website](#)<sup>18</sup>, [documentation](#)<sup>19</sup>, tutorials, mailing lists and others. If you're looking for *commercial support*, find below a list of companies providing pgRouting development and consulting services.

### 1.4.1 Reporting Problems

Bugs are reported and managed in an [issue tracker](#)<sup>20</sup>. 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](#)<sup>21</sup> 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;
```

### 1.4.2 Mailing List and GIS StackExchange

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

- User mailing list: <http://lists.osgeo.org/mailman/listinfo/pgrouting-users>
- Developer mailing list: <http://lists.osgeo.org/mailman/listinfo/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](#)<sup>22</sup> and tag the question with `pgrouting`. Find all questions tagged with `pgrouting` under <http://gis.stackexchange.com/questions/tagged/pgrouting> or subscribe to the [pgRouting questions feed](#)<sup>23</sup>.

### 1.4.3 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:

---

<sup>18</sup><http://www.pgrouting.org>

<sup>19</sup><http://docs.pgrouting.org>

<sup>20</sup><https://github.com/pgrouting/pgrouting/issues>

<sup>21</sup><https://github.com/pgRouting/pgrouting/issues/new>

<sup>22</sup><http://gis.stackexchange.com/>

<sup>23</sup><http://gis.stackexchange.com/feeds/tag?tagnames=pgrouting&sort=newest>

<b>Company</b>	<b>Offices in</b>	<b>Website</b>
Georepublic	Germany, Japan	<a href="http://georepublic.info">http://georepublic.info</a>
iMaptools	United States	<a href="http://imaptools.com">http://imaptools.com</a>
Orkney Inc.	Japan	<a href="http://www.orkney.co.jp">http://www.orkney.co.jp</a>
Camptocamp	Switzerland, France	<a href="http://www.camptocamp.com">http://www.camptocamp.com</a>



*Tutorial*

- *Getting started*
- *Routing Topology* for an overview of a topology for routing algorithms.
- *Graph Analytics* for an overview of the analysis of a graph.
- *Dictionary of columns & Custom Query* that is used in the routing algorithms.
- *Performance Tips* to improve your performance.
- *Sample Data* that is used in the examples of this manual.
- *User's Recipes List*

## 2.1 Tutorial

*Getting started*

- How to create a database to use for our project
- How to load some data
- How to build a topology
- How to check your graph for errors
- How to compute a route
- How to use other tools to view your graph and route
- How to create a web app

*Advanced Topics*

- *Routing Topology* for an overview of a topology for routing algorithms.
- *Graph Analytics* for an overview of the analysis of a graph.
- *Dictionary of columns & Custom Query* that is used in the routing algorithms.
- *Performance Tips* to improve your performance.

### 2.1.1 Getting Started

This is a simple guide to walk you through the steps of getting started with pgRouting. In this guide we will cover:

- How to create a database to use for our project
- How to load some data

- How to build a topology
- How to check your graph for errors
- How to compute a route
- How to use other tools to view your graph and route
- How to create a web app

## How to create a database to use for our project

The first thing we need to do is create a database and load pgrouting in the database. Typically you will create a database for each project. Once you have a database to work in, you can load your data and build your application in that database. This makes it easy to move your project later if you want to say a production server.

For Postgresql 9.1 and later versions

```
createdb mydatabase
psql mydatabase -c "create extension postgis"
psql mydatabase -c "create extension pgrouting"
```

For older versions of postgresql

```
createdb -T template1 template_postgis
psql template_postgis -c "create language plpgsql"
psql template_postgis -f /usr/share/postgresql/9.0/contrib/postgis-1.5/postgis.sql
psql template_postgis -f /usr/share/postgresql/9.0/contrib/postgis-1.5/spatial_ref_sys.sql
psql template_postgis -f /usr/share/postgresql/9.0/contrib/postgis_comments.sql

createdb -T template_postgis template_pgrouting
psql template_pgrouting -f /usr/share/postgresql/9.0/contrib/pgrouting-2.0/pgrouting.sql

createdb -T template_pgrouting mydatabase
```

## How to load some data

How you load your data will depend in what form it comes it. There are various OpenSource tools that can help you, like:

### shp2pgsql

- this is the postgresql shapefile loader

### ogr2ogr

- this is a vector data conversion utility

### osm2pgsql

- this is a tool for loading OSM data into postgresql

So these tools and probably others will allow you to read vector data and and can load that data into your database as a table of some kind. At this point you need to know a little about your data structure and content. One easy way to browse you data table is with pgAdmin3 or phpPgAdmin.

## How to build a topology

Next we need to build a topology for our street data. What this means is that for any given edge in your street data the ends of that edge will be connected to a unique node and to other edges that are also connected to that same unique node. Once all the edges are connected to nodes we have a graph that can be used for routing with pgrouting. We provide a tools the will help with this:



```
select pgr_createTopology('myroads', 0.000001);
```

See *pgr\_createTopology* for more information.

## How to check your graph for errors

There are lots of possible sources for errors in a graph. The data that you started with may not have been designed with routing in mind. A graph as some very specific requirements. One it that it is *NODED*, this means that except for some very specific use cases, each road segments starts and ends at a node and that in general it does not cross another road segment that it should be connected to.

There can be other errors like the direction of a one-way street being entered in the wrong direction. We do not have tools to search for all possible errors but we have some basic tools that might help.

```
select pgr_analyzegraph('myroads', 0.000001);
select pgr_analyzeoneway('myroads', s_in_rules, s_out_rules,
                             t_in_rules, t_out_rules
                             direction)
```

See *Graph Analytics* for more information.

If your data needs to be *NODED*, we have a tool that can help for that also.

See *pgr\_nodeNetwork* for more information.

## How to compute a route

Once you have all the prep work done above, computing a route is fairly easy. We have a lot of different algorithms but they can work with your prepared road network. The general form of a route query is:

```
select pgr_<algorithm>(<SQL for edges>, start, end, <additional options>)
```

As you can see this is fairly straight forward and you can look at the specific algorithms for the details on how to use them. What you get as a result from these queries will be a set of records of type *pgr\_costResult[]* or *pgr\_geomResult[]*. These results have information like edge id and/or the node id along with the cost or geometry for the step in the path from *start* to *end*. Using the ids you can join these results back to your edge table to get more information about each step in the path.

- See also *pgr\_costResult[]* and *pgr\_geomResult[]*.

## 2.1.2 Routing Topology

**Author** Stephen Woodbridge <woodbri@swoodbridge.com<sup>1</sup>>

**Copyright** Stephen Woodbridge. The source code is released under the MIT-X license.

### Overview

Typically when GIS files are loaded into the data database for use with pgRouting they do not have topology information associated with them. To create a useful topology the data needs to be “noded”. This means that where two or more roads form an intersection there it needs to be a node at the intersection and all the road segments need to be broken at the intersection, assuming that you can navigate from any of these segments to any other segment via that intersection.

You can use the *graph analysis functions* to help you see where you might have topology problems in your data. If you need to node your data, we also have a function *pgr\_nodeNetwork()* that might work for you. This function splits ALL crossing segments and nodes them. There are some cases where this might NOT be the right thing to do.

<sup>1</sup>woodbri@swoodbridge.com

For example, when you have an overpass and underpass intersection, you do not want these noded, but pgr\_nodeNetwork does not know that is the case and will node them which is not good because then the router will be able to turn off the overpass onto the underpass like it was a flat 2D intersection. To deal with this problem some data sets use z-levels at these types of intersections and other data might not node these intersection which would be ok.

For those cases where topology needs to be added the following functions may be useful. One way to prep the data for pgRouting is to add the following columns to your table and then populate them as appropriate. This example makes a lot of assumption like that you original data tables already has certain columns in it like `one_way`, `fcc`, and possibly others and that they contain specific data values. This is only to give you an idea of what you can do with your data.

```
ALTER TABLE edge_table
  ADD COLUMN source integer,
  ADD COLUMN target integer,
  ADD COLUMN cost_len double precision,
  ADD COLUMN cost_time double precision,
  ADD COLUMN rcost_len double precision,
  ADD COLUMN rcost_time double precision,
  ADD COLUMN x1 double precision,
  ADD COLUMN y1 double precision,
  ADD COLUMN x2 double precision,
  ADD COLUMN y2 double precision,
  ADD COLUMN to_cost double precision,
  ADD COLUMN rule text,
  ADD COLUMN isolated integer;

SELECT pgr_createTopology('edge_table', 0.000001, 'the_geom', 'id');
```

The function `pgr_createTopology()` will create the `vertices_tmp` table and populate the source and target columns. The following example populated the remaining columns. In this example, the `fcc` column contains feature class code and the CASE statements converts it to an average speed.

```
UPDATE edge_table SET x1 = st_x(st_startpoint(the_geom)),
  y1 = st_y(st_startpoint(the_geom)),
  x2 = st_x(st_endpoint(the_geom)),
  y2 = st_y(st_endpoint(the_geom)),
  cost_len = st_length_spheroid(the_geom, 'SPHEROID["WGS84",6378137,298.25728]'),
  rcost_len = st_length_spheroid(the_geom, 'SPHEROID["WGS84",6378137,298.25728]'),
  len_km = st_length_spheroid(the_geom, 'SPHEROID["WGS84",6378137,298.25728]')/1000.0,
  len_miles = st_length_spheroid(the_geom, 'SPHEROID["WGS84",6378137,298.25728]')
  / 1000.0 * 0.6213712,
  speed_mph = CASE WHEN fcc='A10' THEN 65
    WHEN fcc='A15' THEN 65
    WHEN fcc='A20' THEN 55
    WHEN fcc='A25' THEN 55
    WHEN fcc='A30' THEN 45
    WHEN fcc='A35' THEN 45
    WHEN fcc='A40' THEN 35
    WHEN fcc='A45' THEN 35
    WHEN fcc='A50' THEN 25
    WHEN fcc='A60' THEN 25
    WHEN fcc='A61' THEN 25
    WHEN fcc='A62' THEN 25
    WHEN fcc='A64' THEN 25
    WHEN fcc='A70' THEN 15
    WHEN fcc='A69' THEN 10
    ELSE null END,
  speed_kmh = CASE WHEN fcc='A10' THEN 104
    WHEN fcc='A15' THEN 104
    WHEN fcc='A20' THEN 88
    WHEN fcc='A25' THEN 88
    WHEN fcc='A30' THEN 72
```

```

        WHEN fcc='A35' THEN 72
        WHEN fcc='A40' THEN 56
        WHEN fcc='A45' THEN 56
        WHEN fcc='A50' THEN 40
        WHEN fcc='A60' THEN 50
        WHEN fcc='A61' THEN 40
        WHEN fcc='A62' THEN 40
        WHEN fcc='A64' THEN 40
        WHEN fcc='A70' THEN 25
        WHEN fcc='A69' THEN 15
        ELSE null END;

-- UPDATE the cost infomation based on oneway streets

UPDATE edge_table SET
  cost_time = CASE
    WHEN one_way='TF' THEN 10000.0
    ELSE cost_len/1000.0/speed_kmh::numeric*3600.0
  END,
  rcost_time = CASE
    WHEN one_way='FT' THEN 10000.0
    ELSE cost_len/1000.0/speed_kmh::numeric*3600.0
  END;

-- clean up the database because we have updated a lot of records

VACUUM ANALYZE VERBOSE edge_table;

```

Now your database should be ready to use any (most?) of the pgRouting algorithms.

## See Also

- *pgr\_createTopology*
- *pgr\_nodeNetwork*
- *pgr\_pointToId*

## 2.1.3 Graph Analytics

**Author** Stephen Woodbridge <woodbri@swoodbridge.com<sup>2</sup>>

**Copyright** Stephen Woodbridge. The source code is released under the MIT-X license.

### Overview

It is common to find problems with graphs that have not been constructed fully noded or in graphs with z-levels at intersection that have been entered incorrectly. An other problem is one way streets that have been entered in the wrong direction. We can not detect errors with respect to “ground” truth, but we can look for inconsistencies and some anomalies in a graph and report them for additional inspections.

We do not current have any visualization tools for these problems, but I have used mapserver to render the graph and highlight potential problem areas. Someone familiar with graphviz might contribute tools for generating images with that.

<sup>2</sup>woodbri@swoodbridge.com

## Analyze a Graph

With *pgr\_analyzeGraph* the graph can be checked for errors. For example for table “mytab” that has “mytab\_vertices\_pgr” as the vertices table:

```
SELECT pgr_analyzeGraph('mytab', 0.000002);
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: 158
NOTICE: Dead ends: 20028
NOTICE: Potential gaps found near dead ends: 527
NOTICE: Intersections detected: 2560
NOTICE: Ring geometries: 0
pgr_analyzeGraph
-----
OK
(1 row)
```

In the vertices table “mytab\_vertices\_pgr”:

- Deadends are identified by cnt=1
- Potential gap problems are identified with chk=1.

```
SELECT count(*) as deadends FROM mytab_vertices_pgr WHERE cnt = 1;
deadends
-----
20028
(1 row)

SELECT count(*) as gaps FROM mytab_vertices_pgr WHERE chk = 1;
gaps
-----
527
(1 row)
```

For isolated road segments, for example, a segment where both ends are deadends. you can find these with the following query:

```
SELECT *
FROM mytab a, mytab_vertices_pgr b, mytab_vertices_pgr c
WHERE a.source=b.id AND b.cnt=1 AND a.target=c.id AND c.cnt=1;
```

If you want to visualize these on a graphic image, then you can use something like mapserver to render the edges and the vertices and style based on cnt or if they are isolated, etc. You can also do this with a tool like graphviz, or geoserver or other similar tools.

## Analyze One Way Streets

*pgr\_analyzeOneway* analyzes one way streets in a graph and identifies any flipped segments. Basically if you count the edges coming into a node and the edges exiting a node the number has to be greater than one.

This query will add two columns to the vertices\_tmp table `ein int` and `eout int` and populate it with the appropriate counts. After running this on a graph you can identify nodes with potential problems with the following query.

The rules are defined as an array of text strings that if match the `col` value would be counted as true for the source or target in or out condition.

## Example

Lets assume we have a table “st” of edges and a column “one\_way” that might have values like:

- ‘FT’ - oneway from the source to the target node.
- ‘TF’ - oneway from the target to the source node.
- ‘B’ - two way street.
- ‘’ - empty field, assume twoway.
- <NULL> - NULL field, use two\_way\_if\_null flag.

Then we could form the following query to analyze the oneway streets for errors.

```
SELECT pgr_analyzeOneway('mytab',
    ARRAY['', 'B', 'TF'],
    ARRAY['', 'B', 'FT'],
    ARRAY['', 'B', 'FT'],
    ARRAY['', 'B', 'TF'],
);

-- now we can see the problem nodes
SELECT * FROM mytab_vertices_pgr WHERE ein=0 OR eout=0;

-- and the problem edges connected to those nodes
SELECT gid FROM mytab a, mytab_vertices_pgr b WHERE a.source=b.id AND ein=0 OR eout=0
UNION
SELECT gid FROM mytab a, mytab_vertices_pgr b WHERE a.target=b.id AND ein=0 OR eout=0;
```

Typically these problems are generated by a break in the network, the one way direction set wrong, maybe an error related to z-levels or a network that is not properly noded.

The above tools do not detect all network issues, but they will identify some common problems. There are other problems that are hard to detect because they are more global in nature like multiple disconnected networks. Think of an island with a road network that is not connected to the mainland network because the bridge or ferry routes are missing.

## See Also

- [pgr\\_analyzeGraph](#)
- [pgr\\_analyzeOneway](#)
- [pgr\\_nodeNetwork](#)

## 2.1.4 Dictionary of columns & Custom Query

**path** a sequence of vertices/edges from A to B.

**route** a sequence of paths.

### Description of the edges\_sql query

**edges\_sql** an SQL query, which should return a set of rows with the following columns:

**id** ANY-INTEGGER identifier of the edge.

**source** ANY-INTEGGER identifier of the first end point vertex of the edge.

**target** ANY-INTEGGER identifier of the second end pont vertex of the edge.

**cost** ANY-NUMERICAL weight of the edge (*source*, *target*), if negative: edge (*source*, *target*) does not exist, therefore it's not part of the graph.

**reverse\_cost** ANY-NUMERICAL (optional) weight of the edge (*target*, *source*), if negative: edge (*target*, *source*) does not exist, therefore it's not part of the graph.

Where:

**ANY-INTEGERS** smallint, int, bigint

**ANY-NUMERICALS** smallint, int, bigint, real, float

### Description of the parameters of the signatures

**edges\_sql** TEXT SQL query as described above.

**start\_vid** BIGINT identifier of the starting vertex of the path.

**start\_vids** array[ANY-INTEGERS] array of identifiers of starting vertices.

**end\_vid** BIGINT identifier of the ending vertex of the path.

**end\_vids** array[ANY-INTEGERS] array of identifiers of ending vertices.

**directed** boolean (optional). When *false* the graph is considered as Undirected. Default is *true* which considers the graph as Directed.

### Description of the return values

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

**seq** INTEGER is a sequential value starting from 1.

**route\_seq** INTEGER relative position in the route. Has value 1 for the beginning of a route.

**route\_id** INTEGER id of the route.

**path\_seq** INTEGER relative position in the path. Has value 1 for the beginning of a path.

**path\_id** INTEGER id of the path.

**start\_vid** BIGINT id of the starting vertex. Used when multiple starting vertices are in the query.

**end\_vid** BIGINT id of the ending vertex. Used when multiple ending vertices are in the query.

**node** BIGINT id of the node in the path from start\_vid to end\_v.

**edge** BIGINT id 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 total cost from start\_vid to node.

### Descriptions for version 2.0 signatures

In general, the routing algorithms need an SQL query that contain one or more of the following required columns with the preferred type:

**id** int4

**source** int4

**target** int4

**cost** float8

**reverse\_cost** float8

**x** float8

**y** float8

**x1** float8  
**y1** float8  
**x2** float8  
**y2** float8

```
SELECT source, target, cost FROM edge_table;
SELECT id, source, target, cost FROM edge_table;
SELECT id, source, target, cost, x1, y1, x2, y2 ,reverse_cost FROM edge_table
```

When the edge table has a different name to represent the required columns:

```
SELECT src as source, target, cost FROM othertable;
SELECT gid as id, src as source, target, cost FROM othertable;
SELECT gid as id, src as source, target, cost, fromX as x1, fromY as y1, toX as x2, toY as y2 ,ReverseCost as reverse_cost
FROM othertable;
```

The topology functions use the same names for id, source and target columns of the edge table, The following parameters have as default value:

**id** int4 Default id  
**source** int4 Default source  
**target** int4 Default target  
**the\_geom** text Default the\_geom  
**oneway** text Default oneway  
**rows\_where** text Default true to indicate all rows (this is not a column)

The following parameters do not have a default value and when used they have to be inserted in strict order:

**edge\_table** text  
**tolerance** float8  
**s\_in\_rules** text[]  
**s\_out\_rules** text[]  
**t\_in\_rules** text[]  
**t\_out\_rules** text[]

When the columns required have the default names this can be used (pgr\_func is to represent a topology function)

```
pgr_func('edge_table')           -- when tolerance is not required
pgr_func('edge_table',0.001)    -- when tolerance is required
-- s_in_rule, s_out_rule, st_in_rules, t_out_rules are required
SELECT pgr_analyzeOneway('edge_table', ARRAY['', 'B', 'TF'], ARRAY['', 'B', 'FT'],
                        ARRAY['', 'B', 'FT'], ARRAY['', 'B', 'TF'])
```

When the columns required do not have the default names its strongly recommended to use the *named notation*.

```
pgr_func('othertable', id:='gid', source:='src',the_geom:='mygeom')
pgr_func('othertable',0.001,the_geom:='mygeom',id:='gid',source:='src')
SELECT pgr_analyzeOneway('othertable', ARRAY['', 'B', 'TF'], ARRAY['', 'B', 'FT'],
                        ARRAY['', 'B', 'FT'], ARRAY['', 'B', 'TF']
                        source:='src', oneway:='dir')
```

## 2.1.5 Performance Tips

For the Routing functions:

**Note:** To get faster results bound your queries to the area of interest of routing to have, for example, no more than one million rows.

---

### For the topology functions:

When “you know” that you are going to remove a set of edges from the edges table, and without those edges you are going to use a routing function you can do the following:

Analyze the new topology based on the actual topology:

```
pgr_analyzegraph('edge_table',rows_where:='id < 17');
```

Or create a new topology if the change is permanent:

```
pgr_createTopology('edge_table',rows_where:='id < 17');  
pgr_analyzegraph('edge_table',rows_where:='id < 17');
```

Use an SQL that “removes” the edges in the routing function

```
SELECT id, source, target from edge_table WHERE id < 17
```

When “you know” that the route will not go out of a particular area, to speed up the process you can use a more complex SQL query like

```
SELECT id, source, target from edge_table WHERE  
id < 17 and  
the_geom && (select st_buffer(the_geom,1) as myarea FROM edge_table where id=5)
```

Note that the same condition `id < 17` is used in all cases.

## 2.2 Sample Data

The documentation provides very simple example queries based on a small sample network. To be able to execute the sample queries, run the following SQL commands to create a table with a small network data set.

### Create table

```
CREATE TABLE edge_table (  
id serial,  
dir character varying,  
source integer,  
target integer,  
cost double precision,  
reverse_cost double precision,  
x1 double precision,  
y1 double precision,  
x2 double precision,  
y2 double precision,  
the_geom geometry  
);
```

### Insert network data

```
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1, 1, 2,0, 2,1);  
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES (-1, 1, 2,1, 3,1);  
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES (-1, 1, 3,1, 4,1);  
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1, 1, 2,1, 2,2);
```



```

INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1,-1, 3,1, 3,2);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1, 1, 0,2, 1,2);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1, 1, 1,2, 2,2);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1, 1, 2,2, 3,2);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1, 1, 3,2, 4,2);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1, 1, 2,2, 2,3);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1,-1, 3,2, 3,3);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1,-1, 2,3, 3,3);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1,-1, 3,3, 4,3);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1, 1, 2,3, 2,4);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1, 1, 4,2, 4,3);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1, 1, 4,1, 4,2);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1, 1, 0.5,3.5, 1.9999999999999999,3.5);
INSERT INTO edge_table (cost,reverse_cost,x1,y1,x2,y2) VALUES ( 1, 1, 3.5,2.3, 3.5,4);

UPDATE edge_table SET the_geom = st_makeline(st_point(x1,y1),st_point(x2,y2)),
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
WHEN (cost<0 and reverse_cost>0) THEN 'TF' -- reverse direction
ELSE '' END;

```

Before you test a routing function use this query to fill the source and target columns.

```
SELECT pgr_createTopology('edge_table',0.001);
```

This table is used in some of our examples

```

CREATE TABLE vertex_table (
  id serial,
  x double precision,
  y double precision
);

```

```

INSERT INTO vertex_table VALUES
(1,2,0), (2,2,1), (3,3,1), (4,4,1), (5,0,2), (6,1,2), (7,2,2),
(8,3,2), (9,4,2), (10,2,3), (11,3,3), (12,4,3), (13,2,4);

```

**Network for queries marked as directed and cost and reverse\_cost columns are used:**

Blue colored lines with line arrows, represent the cost of the edge table. Green colored lines with filled arrows, represent the reverse\_cost of the edge table.

**Network for queries marked as undirected and cost and reverse\_cost columns are used:**

**Network for queries marked as directed and only cost column is used:**

**Network for queries marked as undirected and only cost column is used:**

## 2.3 User's Recipes List

### 2.3.1 Comparing topology of a unnoded network with a noded network

**Author** pgRouting team.

**Licence** Open Source




Fig. 2.1: Graph 1: Directed, with cost and reverse cost

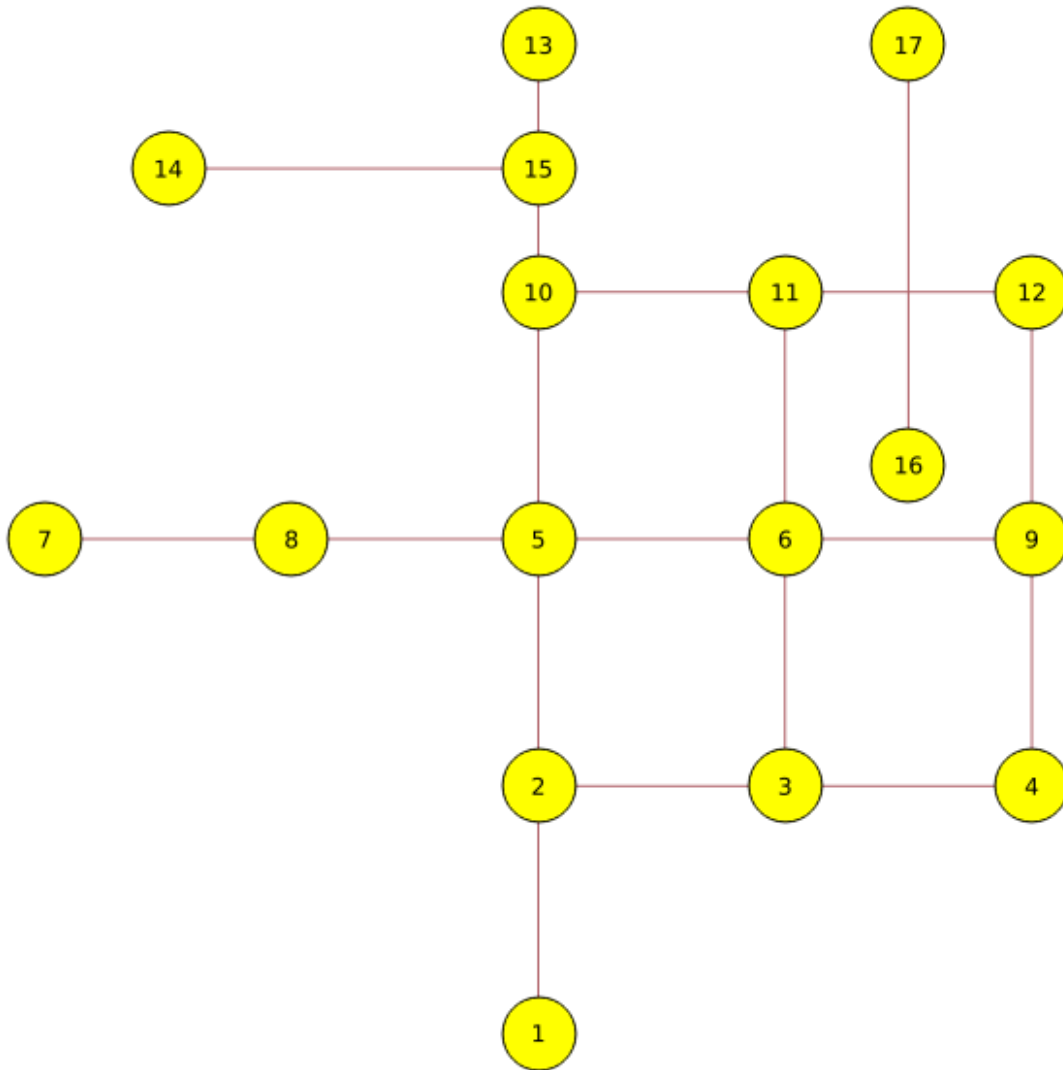


Fig. 2.2: Graph 2: Undirected, with cost and reverse cost




Fig. 2.3: Graph 3: Directed, with cost




Fig. 2.4: Graph 4: Undirected, with cost

This recipe uses the *Sample Data* network.

The purpose of this recipe is to compare a not noded network with a noded network.

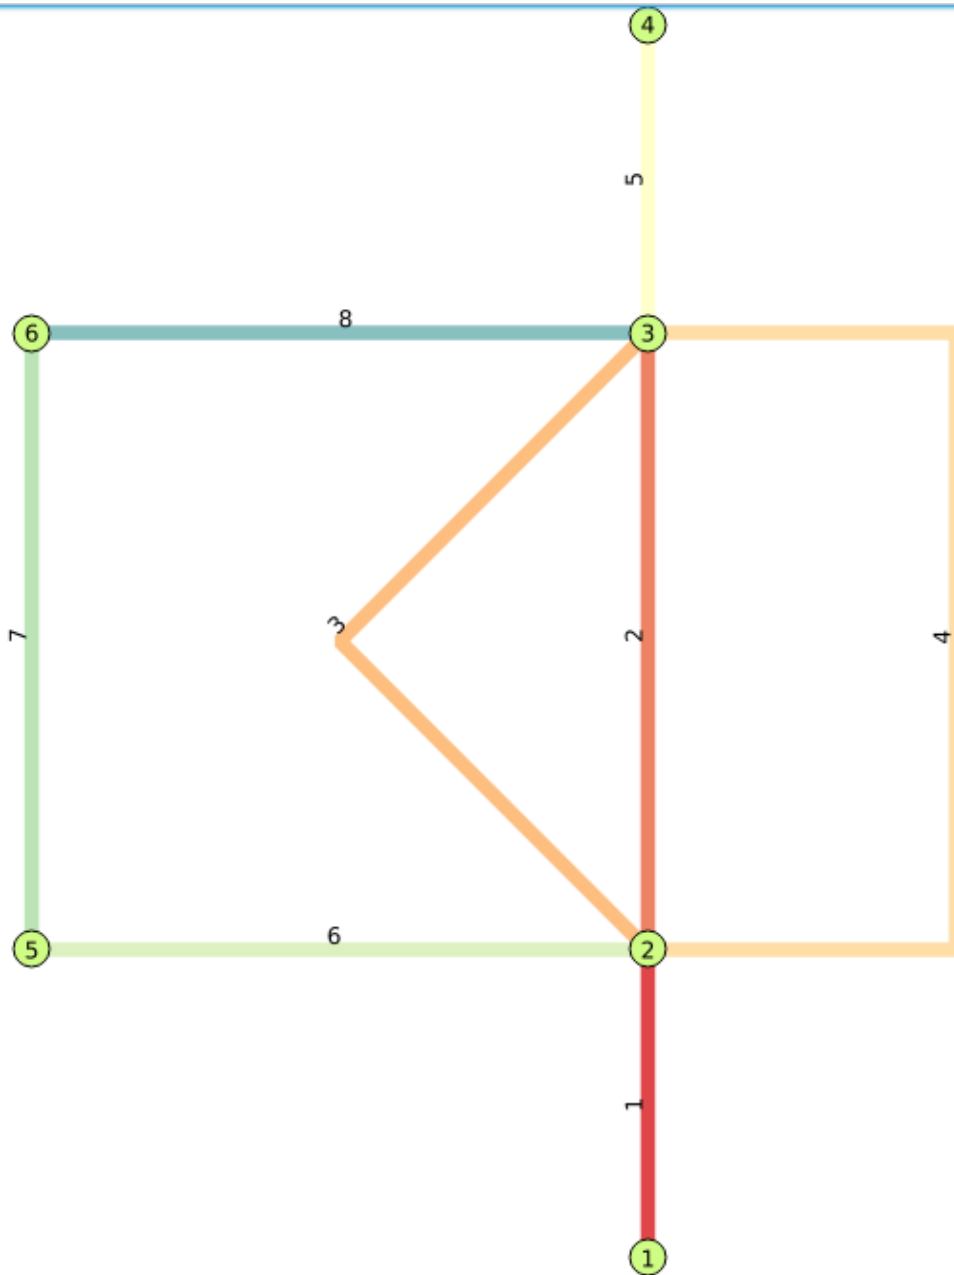
```
SELECT pgr_createTopology('edge_table', 0.001);
SELECT pgr_analyzegraph('edge_table', 0.001);
SELECT pgr_nodeNetwork('edge_table', 0.001);
SELECT pgr_createTopology('edge_table_noded', 0.001);
SELECT pgr_analyzegraph('edge_table_noded', 0.001);
```

### 2.3.2 Handling parallels after getting a path (pgr\_ksp focus)

**Author** pgRouting team.

**Licence** Open Source

## The graph



## Data

```

drop table if exists parallel;
CREATE TABLE parallel (
  id serial,
  source integer,
  target integer,
  cost double precision,
  reverse_cost double precision,
  x1 double precision,
  y1 double precision,
  x2 double precision,
  y2 double precision,
  the_geom geometry
);

```

```

INSERT INTO parallel (x1,y1,x2,y2)
  VALUES (1,0,1,1), (1,1,1,3), (1,1,1,3), (1,1,1,3), (1,3,1,4), (1,1,-1,1), (-1,1,-1,3), (-1,3,1,3);
UPDATE parallel SET the_geom = ST_makeline(ST_point(x1,y1),ST_point(x2,y2));
UPDATE parallel SET the_geom = ST_makeline(ARRAY[ST_point(1,1),ST_point(0,2),ST_point(1,3)]) WHERE id = 4;
UPDATE parallel SET the_geom = ST_makeline(ARRAY[ST_point(1,1),ST_point(2,1),ST_point(2,3),ST_point(1,3)])
  WHERE id = 4;
UPDATE parallel SET cost = ST_length(the_geom), reverse_cost = ST_length(the_geom);
SELECT pgr_createTopology('parallel',0.001);

```

### pg\_r\_ksp results

We ignore the costs because we want all the parallels

```

SELECT seq, path_id AS route, node, edge INTO routes
  FROM pgr_ksp('select id, source, target, cost, reverse_cost from parallel',
1, 4, 3);

select route, node, edge from routes;

```

route	node	edge
1	1	1
1	2	2
1	3	5
1	4	-1
2	1	1
2	2	6
2	5	7
2	6	8
2	3	5
2	4	-1

(10 rows)

We need an aggregate function:

```

CREATE AGGREGATE array_accum (anyelement)
(
  sfunc = array_append,
  stype = anyarray,
  initcond = '{}'
);

```

Now lets generate a table with the parallel edges.

```

select distinct seq,route,source,target, array_accum(id) as edges into paths
  from (select seq, route, source, target
        from parallel, routes where id = edge) as r
  join parallel using (source, target)
  group by seq,route,source,target order by seq;

select route, source, targets, edges from paths;

```

route	source	target	edges
1	1	2	{1}
2	1	2	{1}
2	2	5	{6}
1	2	3	{2,3,4}
2	5	6	{7}
1	3	4	{5}



```

2 |      6 |      3 | {8}
2 |      3 |      4 | {5}
(8 rows)

```

### Some more aggregate functions

To generate a table with all the combinations for parallel routes, we need some more aggregates

```

create or replace function multiply( integer, integer )
returns integer as
$$
  select $1 * $2;
$$
language sql stable;

create aggregate prod(integer)
(
  sfunc = multiply,
  stype = integer,
  initcond = 1
);

```

### And a function that “Expands” the table

```

CREATE OR REPLACE function    expand_parallel_edge_paths(tab text)
  returns TABLE (
    seq    INTEGER,
    route  INTEGER,
    source INTEGER, target INTEGER, -- this ones are not really needed
    edge   INTEGER ) AS

$body$
DECLARE
nroutes  INTEGER;
newroutes INTEGER;
rec      record;
seq2     INTEGER := 1;
rnum     INTEGER := 0;

BEGIN    -- get the number of distinct routes
  execute 'select count(DISTINCT route) from ' || tab INTO nroutes;
  FOR i IN 0..nroutes-1
  LOOP
    -- compute the number of new routes this route will expand into
    -- this is the product of the lengths of the edges array for each route
    execute 'select prod(array_length(edges, 1))-1 from '
    ||      quote_ident(tab) || ' where route=' || i INTO newroutes;
    -- now we generate the number of new routes for this route
    -- by repeatedly listing the route and swapping out the parallel edges
    FOR j IN 0..newroutes
    LOOP
      -- query the specific route
      FOR rec IN execute 'select * from ' || quote_ident(tab) || ' where route=' || i
      || ' order by seq'
      LOOP
        seq := seq2;
        route := rnum;
        source := rec.source;
        target := rec.target;
        -- using module arithmetic iterate through the various edge choices
        edge := rec.edges[(j % (array_length(rec.edges, 1)))+1];

```

```

        -- return a new record
        RETURN next;
        seq2 := seq2 + 1;    -- increment the record count
    END LOOP;
    seq := seq2;
    route := rnum;
    source := rec.target;
    target := -1;
    edge := -1;
    RETURN next; -- Insert the ending record of the route
    seq2 := seq2 + 1;

    rnum := rnum + 1; -- increment the route count
END LOOP;
END LOOP;
END;
$body$
language plpgsql volatile strict    cost 100 rows 100;

```

### Test it

```

select * from expand_parallel_edge_paths( 'paths' );
seq | route | source | target | edge
-----+-----+-----+-----+-----
  1 |     0 |     1 |     2 |     1
  2 |     0 |     2 |     3 |     2
  3 |     0 |     3 |     4 |     5
  4 |     0 |     4 |    -1 |    -1
  5 |     1 |     1 |     2 |     1
  6 |     1 |     2 |     3 |     3
  7 |     1 |     3 |     4 |     5
  8 |     1 |     4 |    -1 |    -1
  9 |     2 |     1 |     2 |     1
 10 |     2 |     2 |     3 |     4
 11 |     2 |     3 |     4 |     5
 12 |     2 |     4 |    -1 |    -1
 13 |     3 |     1 |     2 |     1
 14 |     3 |     2 |     5 |     6
 15 |     3 |     5 |     6 |     7
 16 |     3 |     6 |     3 |     8
 17 |     3 |     3 |     4 |     5
 18 |     3 |     4 |    -1 |    -1
(18 rows)

```

*No more contributions*

## 2.4 How to contribute.

### To add a recipe or a wrapper

The first steps are:

- Fork the repository
- Create a branch for your recipe or wrapper
- Create your recipe file

```
cd doc/src/recipes/  
vi myrecipe.rst  
git add myrecipe.rst  
# include the recipe in this file  
vi index.rst
```

### To create the test file of your recipe

```
cd test  
cp myrecipe.rst myrecipe.sql.test  
  
# make your test based on your recipe  
vi myrecipe.sql.test  
git add myrecipe.sql.test  
  
# create your test results file  
touch myrecipe.result  
git add myrecipe.result  
  
# add your test to the file  
vi test.conf
```

Leave only the SQL code on `myrecipe.sql.test` by deleting lines or by commenting lines.

### To get the results of your recipe

From the root directory execute:

```
tools/test-runner.pl -alg recipes -ignorenotice
```

Copy the results of your recipe and paste them in the file `myrecipe.result`, and remove the “>” from the file.

### make a pull request.

```
git commit -a -m 'myrecipe is for this and that'  
git push
```

From your fork in github make a pull request over develop

For a more complete introduction how to build a routing application read the [pgRouting Workshop](http://workshop.pgrouting.org)<sup>3</sup>.

---

<sup>3</sup><http://workshop.pgrouting.org>



*pgr\_version* - to get pgRouting's version information.

## 3.1 pgr\_version

### 3.1.1 Name

*pgr\_version* — Query for pgRouting version information.

### 3.1.2 Synopsis

Returns a table with pgRouting version information.

```
table() pgr_version();
```

### 3.1.3 Description

Returns a table with:

**version** varchar pgRouting version  
**tag** varchar Git tag of pgRouting build  
**hash** varchar Git hash of pgRouting build  
**branch** varchar Git branch of pgRouting build  
**boost** varchar Boost version

#### History

- New in version 2.0.0

### 3.1.4 Examples

- Query for full version string

```
SELECT pgr_version();  
  
          pgr_version  
-----
```

```
(2.0.0-dev,v2.0dev,290,c64bcb9,sew-devel-2_0,1.49.0)
(1 row)
```

- Query for version and boost attribute

```
SELECT version, boost FROM pgr_version();

 version | boost
-----+-----
 2.0.0-dev | 1.49.0
(1 row)
```

### 3.1.5 See Also

- *pgr\_versionless* to compare two version numbers

---

## Data Types

---

### *pgRouting Data Types*

- *pgr\_costResult[]* - A set of records to describe a path result with cost attribute.
- *pgr\_costResult3[]* - A set of records to describe a path result with cost attribute.
- *pgr\_geomResult* - A set of records to describe a path result with geometry attribute.

## 4.1 pgRouting Data Types

The following are commonly used data types for some of the pgRouting functions.

- *pgr\_costResult[]* - A set of records to describe a path result with cost attribute.
- *pgr\_costResult3[]* - A set of records to describe a path result with cost attribute.
- *pgr\_geomResult* - A set of records to describe a path result with geometry attribute.

### 4.1.1 pgr\_costResult[]

#### Name

`pgr_costResult[]` — A set of records to describe a path result with cost attribute.

#### Description

```
CREATE TYPE pgr_costResult AS
(
  seq integer,
  id1 integer,
  id2 integer,
  cost float8
);
```

**seq** sequential ID indicating the path order

**id1** generic name, to be specified by the function, typically the node id

**id2** generic name, to be specified by the function, typically the edge id

**cost** cost attribute

## 4.1.2 pgr\_costResult3[] - Multiple Path Results with Cost

### Name

pgr\_costResult3[] — A set of records to describe a path result with cost attribute.

### Description

```
CREATE TYPE pgr_costResult3 AS
(
  seq integer,
  id1 integer,
  id2 integer,
  id3 integer,
  cost float8
);
```

**seq** sequential ID indicating the path order

**id1** generic name, to be specified by the function, typically the path id

**id2** generic name, to be specified by the function, typically the node id

**id3** generic name, to be specified by the function, typically the edge id

**cost** cost attribute

### History

- New in version 2.0.0
- Replaces `path_result`

### See Also

- [Introduction](#)

## 4.1.3 pgr\_geomResult[]

### Name

pgr\_geomResult[] — A set of records to describe a path result with geometry attribute.

### Description

```
CREATE TYPE pgr_geomResult AS
(
  seq integer,
  id1 integer,
  id2 integer,
  geom geometry
);
```

**seq** sequential ID indicating the path order

**id1** generic name, to be specified by the function

**id2** generic name, to be specified by the function



**geom** geometry attribute

#### History

- New in version 2.0.0
- Replaces geoms

#### See Also

- *Introduction*



---

## Topology functions

---

### Topology Functions

- *pgr\_createTopology* - to create a topology based on the geometry.
- *pgr\_createVerticesTable* - to reconstruct the vertices table based on the source and target information.
- *pgr\_analyzeGraph* - to analyze the edges and vertices of the edge table.
- *pgr\_analyzeOneway* - to analyze directionality of the edges.
- *pgr\_nodeNetwork* -to create nodes to a not noded edge table.

## 5.1 Topology Functions

The pgRouting's topology of a network, represented with an edge table with source and target attributes and a vertices table associated with it. Depending on the algorithm, you can create a topology or just reconstruct the vertices table, You can analyze the topology, We also provide a function to node an unoded network.

- *pgr\_createTopology* - to create a topology based on the geometry.
- *pgr\_createVerticesTable* - to reconstruct the vertices table based on the source and target information.
- *pgr\_analyzeGraph* - to analyze the edges and vertices of the edge table.
- *pgr\_analyzeOneway* - to analyze directionality of the edges.
- *pgr\_nodeNetwork* -to create nodes to a not noded edge table.

### 5.1.1 pgr\_createTopology

#### Name

`pgr_createTopology` — Builds a network topology based on the geometry information.

#### Synopsis

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.

```
varchar pgr_createTopology(text edge_table, double precision tolerance,
                          text the_geom='the_geom', text id='id',
                          text source='source',text target='target',
                          text rows_where='true', boolean clean:=false)
```

## Description

### 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 vertices table is a requirement of the `pgr_analyzeGraph` and the `pgr_analyzeOneway` functions.

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. See `pgr_analyzeGraph`.
- chk** integer Indicator that the vertex might have a problem. See `pgr_analyzeGraph`.
- ein** integer Number of vertices in the `edge_table` that reference this vertex AS incoming. See `pgr_analyzeOneway`.

**out** integer Number of vertices in the edge\_table that reference this vertex AS outgoing. See *pgr\_analyzeOneway*.

**the\_geom** geometry Point geometry of the vertex.

## History

- Renamed in version 2.0.0

## Usage when the edge table's columns MATCH the default values:

The simplest way to use `pgr_createtopology` is:

```
SELECT pgr_createTopology('edge_table',0.001);
```

When the arguments are given in the order described in the parameters:

```
SELECT pgr_createTopology('edge_table',0.001,'the_geom','id','source','target');
```

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 `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('edge_table',0.001,'id','the_geom','source','target');
ERROR: Can not determine the srid of the geometry "id" in table public.edge_table
```

## When using the named notation

The order of the parameters do not matter:

```
SELECT pgr_createTopology('edge_table',0.001,the_geom:='the_geom',id:='id',source:='source',target:='target');
```

```
SELECT pgr_createTopology('edge_table',0.001,source:='source',id:='id',target:='target',the_geom:='the_geom');
```

Parameters defined with a default value can be omitted, AS long AS the value matches the default:

```
SELECT pgr_createTopology('edge_table',0.001,source:='source');
```

## Selecting rows using rows\_where parameter

Selecting rows based on the `id`.

```
SELECT pgr_createTopology('edge_table',0.001,rows_where:='id < 10');
```

Selecting the rows where the geometry is near the geometry of row with `id=5`.

```
SELECT pgr_createTopology('edge_table',0.001,rows_where:='the_geom && (SELECT st_buffer(the_geom, 100) FROM other_table WHERE gid=100)');
```

Selecting the rows where the geometry is near the geometry of the row with `gid=100` of the table `othertable`.

```
DROP TABLE IF EXISTS otherTable;
CREATE TABLE otherTable AS (SELECT 100 AS gid, st_point(2.5,2.5) AS other_geom);
SELECT pgr_createTopology('edge_table',0.001,rows_where:='the_geom && (SELECT st_buffer(other_geom,1) FROM otherTable WHERE gid=100)');
```

### Usage when the edge table's columns DO NOT MATCH the default values:

For the following table

```
DROP TABLE IF EXISTS mytable;
CREATE TABLE mytable AS (SELECT id AS gid, the_geom AS mygeom, source AS src ,target AS tgt FROM edge_table);
```

#### Using positional notation:

The arguments need to be given in the order described in the parameters:

```
SELECT pgr_createTopology('mytable',0.001,'mygeom','gid','src','tgt');
```

#### Warning:

An error would occur when the arguments are not given in the appropriate order: In this example, the column `gid` 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_createTopology('mytable',0.001,'gid','mygeom','src','tgt');
ERROR: Can not determine the srid of the geometry "gid" in table public.mytable
```

#### When using the named notation

The order of the parameters do not matter:

```
SELECT pgr_createTopology('mytable',0.001,the_geom:='mygeom',id:='gid',source:='src',target:='tgt');
```

```
SELECT pgr_createTopology('mytable',0.001,source:='src',id:='gid',target:='tgt',the_geom:='mygeom');
```

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_createTopology('mytable',0.001,'mygeom','gid','src','tgt',rows_where:='gid < 10');
```

```
SELECT pgr_createTopology('mytable',0.001,source:='src',id:='gid',target:='tgt',the_geom:='mygeom',rows_where:='gid < 10');
```

Selecting the rows where the geometry is near the geometry of row with `id=5`.

```
SELECT pgr_createTopology('mytable',0.001,'mygeom','gid','src','tgt',
    rows_where:='mygeom && (SELECT st_buffer(mygeom,1) FROM mytable WHERE gid=5)');
```

```
SELECT pgr_createTopology('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)');
```

Selecting the rows where the geometry is near the geometry of the row with `gid=100` of the table `othertable`.

```
DROP TABLE IF EXISTS otherTable;
CREATE TABLE otherTable AS (SELECT 100 AS gid, st_point(2.5,2.5) AS other_geom) ;
SELECT pgr_createTopology('mytable',0.001,'mygeom','gid','src','tgt',
    rows_where:='mygeom && (SELECT st_buffer(other_geom,1) FROM otherTable WHERE gid=100)');
```

```
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 gi
```

## Examples

```
SELECT pgr_createTopology('edge_table', 0.001,rows_where:='id<10');
NOTICE:  PROCESSING:
NOTICE:  pgr_createTopology('edge_table',0.0001,'the_geom','id','source','target','id<10')
NOTICE:  Performing checks, pelase wait .....
NOTICE:  Creating Topology, Please wait...
NOTICE:  -----> TOPOLOGY CREATED FOR 9 edges
NOTICE:  Rows with NULL geometry or NULL id: 0
NOTICE:  Vertices table for table public.edge_table is: public.edge_table_vertices_pgr

pgr_createtopology
-----
OK
(1 row)

SELECT pgr_createTopology('edge_table', 0.001);
NOTICE:  PROCESSING:
NOTICE:  pgr_createTopology('edge_table',0.0001,'the_geom','id','source','target','true')
NOTICE:  Performing checks, pelase 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.edge_table is: public.edge_table_vertices_pgr

pgr_createtopology
-----
OK
(1 row)
```

The example uses the *Sample Data* network.

## See Also

- *Routing Topology* for an overview of a topology for routing algorithms.
- *pgr\_createVerticesTable* to reconstruct the vertices table based on the source and target information.
- *pgr\_analyzeGraph* to analyze the edges and vertices of the edge table.

## 5.1.2 pgr\_createVerticesTable

### Name

`pgr_createVerticesTable` — Reconstructs the vertices table based on the source and target information.

### Synopsis

The function returns:

- OK after the vertices table has been reconstructed.
- FAIL when the vertices table was not reconstructed due to an error.

```
varchar pgr_createVerticesTable(text edge_table, text the_geom:='the_geom'
text source:='source',text target:='target',text rows_where:='true')
```

## Description

### 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 exist, 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 vertices table is a requirement of the `pgr_analyzeGraph` and the `pgr_analyzeOneway` functions.

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. See `pgr_analyzeGraph`.
- chk** integer Indicator that the vertex might have a problem. See `pgr_analyzeGraph`.
- ein** integer Number of vertices in the `edge_table` that reference this vertex as incoming. See `pgr_analyzeOneway`.
- eout** integer Number of vertices in the `edge_table` that reference this vertex as outgoing. See `pgr_analyzeOneway`.
- the\_geom** geometry Point geometry of the vertex.

### History

- Renamed in version 2.0.0



**Usage when the edge table's columns MATCH the default values:**

The simplest way to use `pgr_createVerticesTable` is:

```
SELECT pgr_createVerticesTable('edge_table');
```

When the arguments are given in the order described in the parameters:

```
SELECT pgr_createVerticesTable('edge_table', 'the_geom', 'source', 'target');
```

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 column `source` 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('edge_table', 'source', 'the_geom', 'target');
```

**When using the named notation**

The order of the parameters do not matter:

```
SELECT pgr_createVerticesTable('edge_table', the_geom:='the_geom', source:='source', target:='target');
```

```
SELECT pgr_createVerticesTable('edge_table', source:='source', target:='target', the_geom:='the_geom');
```

Parameters defined with a default value can be omitted, as long as the value matches the default:

```
SELECT pgr_createVerticesTable('edge_table', source:='source');
```

**Selecting rows using rows\_where parameter**

Selecting rows based on the id.

```
SELECT pgr_createVerticesTable('edge_table', rows_where:='id < 10');
```

Selecting the rows where the geometry is near the geometry of row with `id=5`.

```
SELECT pgr_createVerticesTable('edge_table', rows_where:='the_geom && (select st_buffer(the_geom,
```

Selecting the rows where the geometry is near the geometry of the row with `gid=100` of the table `othertable`.

```
DROP TABLE IF EXISTS otherTable;
CREATE TABLE otherTable AS (SELECT 100 AS gid, st_point(2.5,2.5) AS other_geom) ;
SELECT pgr_createVerticesTable('edge_table', rows_where:='the_geom && (select st_buffer(othergeom,
```

**Usage when the edge table's columns DO NOT MATCH the default values:**

For the following table

```
DROP TABLE IF EXISTS mytable;
CREATE TABLE mytable AS (SELECT id AS gid, the_geom AS mygeom, source AS src ,target AS tgt FROM e
```

### Using positional notation:

The arguments need to be given in the order described in the parameters:

```
SELECT pgr_createVerticesTable('mytable','mygeom','src','tgt');
```

#### 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');
```

### When using the named notation

The order of the parameters do not matter:

```
SELECT pgr_createVerticesTable('mytable',the_geom:='mygeom',source:='src',target:='tgt');
```

```
SELECT pgr_createVerticesTable('mytable',source:='src',target:='tgt',the_geom:='mygeom');
```

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 gid.

```
SELECT pgr_createVerticesTable('mytable','mygeom','src','tgt',rows_where:='gid < 10');
```

```
SELECT pgr_createVerticesTable('mytable',source:='src',target:='tgt',the_geom:='mygeom',rows_where:='gid < 10');
```

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)');
```

```
SELECT pgr_createVerticesTable('mytable',source:='src',target:='tgt',the_geom:='mygeom',
    rows_where:='mygeom && (SELECT st_buffer(mygeom,0.5) FROM mytable WHERE gid=5)');
```

Selecting the rows where the geometry is near the geometry of the row with `gid=100` of the table `othertable`.

```
DROP TABLE IF EXISTS otherTable;
CREATE TABLE otherTable AS (SELECT 100 AS gid, st_point(2.5,2.5) AS other_geom);
SELECT pgr_createVerticesTable('mytable','mygeom','src','tgt',
    rows_where:='the_geom && (SELECT st_buffer(othergeom,0.5) FROM otherTable WHERE gid=100)');
```

```
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)');
```

### Examples

```
SELECT pgr_createVerticesTable('edge_table');
NOTICE: PROCESSING:
NOTICE: pgr_createVerticesTable('edge_table','the_geom','source','target','true')
NOTICE: Performing checks, please wait .....
NOTICE: Populating public.edge_table_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.edge_table is: public.edge_table_vertices_pgr
NOTICE:      -----

      pgr_createVerticesTable
      -----
      OK
      (1 row)

```

The example uses the *Sample Data* network.

## See Also

- *Routing Topology* for an overview of a topology for routing algorithms.
- *pgr\_createTopology* to create a topology based on the geometry.
- *pgr\_analyzeGraph* to analyze the edges and vertices of the edge table.
- *pgr\_analyzeOneway* to analyze directionality of the edges.

### 5.1.3 pgr\_analyzeGraph

#### Name

`pgr_analyzeGraph` — Analyzes the network topology.

#### Synopsis

The function returns:

- OK after the analysis has finished.
- FAIL when the analysis was not completed due to an error.

```

varchar pgr_analyzeGraph(text edge_table, double precision tolerance,
                        text the_geom='the_geom', text id='id',
                        text source='source', text target='target', text rows_where='true')

```

#### Description

#### 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.

- Use *pgr\_createVerticesTable* to create the vertices table.
- Use *pgr\_createTopology* to create the topology and the vertices table.

#### 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 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.

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 by `rows_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 vertices table can be created with *`pgr_createVerticesTable`* or *`pgr_createTopology`*

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. See *`pgr_analyzeOneway`*.
- out** integer Number of vertices in the `edge_table` that reference this vertex as outgoing. See *`pgr_analyzeOneway`*.
- the\_geom** geometry Point geometry of the vertex.

### History

- New in version 2.0.0

### Usage when the edge table's columns MATCH the default values:

The simplest way to use `pgr_analyzeGraph` is:

```
SELECT pgr_create_topology('edge_table',0.001);
SELECT pgr_analyzeGraph('edge_table',0.001);
```

When the arguments are given in the order described in the parameters:

```
SELECT pgr_analyzeGraph('edge_table',0.001,'the_geom','id','source','target');
```

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('edge_table',0.001,'id','the_geom','source','target');
ERROR: Can not determine the srid of the geometry "id" in table public.edge_table
```

### When using the named notation

The order of the parameters do not matter:

```
SELECT pgr_analyzeGraph('edge_table',0.001,the_geom:='the_geom',id:='id',source:='source',target:='target');
```

```
SELECT pgr_analyzeGraph('edge_table',0.001,source:='source',id:='id',target:='target',the_geom:='the_geom');
```

Parameters defined with a default value can be omitted, as long as the value matches the default:

```
SELECT pgr_analyzeGraph('edge_table',0.001,source:='source');
```

### Selecting rows using `rows_where` parameter

Selecting rows based on the `id`. Displays the analysis a the section of the network.

```
SELECT pgr_analyzeGraph('edge_table',0.001,rows_where:='id < 10');
```

Selecting the rows where the geometry is near the geometry of row with `id=5`.

```
SELECT pgr_analyzeGraph('edge_table',0.001,rows_where:='the_geom && (SELECT st_buffer(the_geom,0.001) && the_geom)');
```

Selecting the rows where the geometry is near the geometry of the row with `gid=100` of the table `othertable`.

```
DROP TABLE IF EXISTS otherTable;
CREATE TABLE otherTable AS (SELECT 100 AS gid, st_point(2.5,2.5) AS other_geom) ;
SELECT pgr_analyzeGraph('edge_table',0.001,rows_where:='the_geom && (SELECT st_buffer(other_geom,0.001) && the_geom)');
```

### Usage when the edge table's columns DO NOT MATCH the default values:

For the following table

```
DROP TABLE IF EXISTS mytable;
CREATE TABLE mytable AS (SELECT id AS gid, source AS src ,target AS tgt , the_geom AS mygeom FROM mygeom);
SELECT pgr_createTopology('mytable',0.001,'mygeom','gid','src','tgt');
```

### 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:**

An error would occur when the arguments are not given in the appropriate order: In this example, the column `gid` 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',0.001,'gid','mygeom','src','tgt');
ERROR: Can not determine the srid of the geometry "gid" in table public.mytable
```

**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');
```

```
SELECT pgr_analyzeGraph('mytable',0.001,source:='src',id:='gid',target:='tgt',the_geom:='mygeom');
```

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');
```

```
SELECT pgr_analyzeGraph('mytable',0.001,source:='src',id:='gid',target:='tgt',the_geom:='mygeom');
```

Selecting the rows WHERE the geometry is near the geometry of row with `id=5`.

```
SELECT pgr_analyzeGraph('mytable',0.001,'mygeom','gid','src','tgt',
    rows_where:='mygeom && (SELECT st_buffer(mygeom,1) FROM mytable WHERE id=5)');
```

```
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 id=5)');
```

Selecting the rows WHERE the geometry is near the `place='myhouse'` of the table `othertable`. (note the use of `quote_literal`)

```
DROP TABLE IF EXISTS otherTable;
CREATE TABLE otherTable AS (SELECT 'myhouse'::text AS place, st_point(2.5,2.5) AS other_geom);
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'))');
```

```
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'))');
```

**Examples**

```
SELECT pgr_create_topology('edge_table',0.001);
SELECT pgr_analyzeGraph('edge_table', 0.001);
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edge_table',0.001,'the_geom','id','source','target','true')
NOTICE: Performing checks, pelase 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('edge_table',0.001,rows_where:='id < 10');
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edge_table',0.001,'the_geom','id','source','target','id < 10')
NOTICE: Performing checks, pelase 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('edge_table',0.001,rows_where:='id >= 10');
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edge_table',0.001,'the_geom','id','source','target','id >= 10')
NOTICE: Performing checks, pelase 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)

-- Simulate removal of edges
SELECT pgr_createTopology('edge_table', 0.001,rows_where:='id <17');
SELECT pgr_analyzeGraph('edge_table', 0.001);
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edge_table',0.001,'the_geom','id','source','target','true')
NOTICE: Performing checks, pelase 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)
SELECT pgr_createTopology('edge_table', 0.001,rows_where='id <17');
NOTICE: PROCESSING:
NOTICE: pgr_createTopology('edge_table',0.001,'the_geom','id','source','target','id <17')
NOTICE: Performing checks, pelase wait .....
NOTICE: Creating Topology, Please wait...
NOTICE: -----> TOPOLOGY CREATED FOR 16 edges
NOTICE: Rows with NULL geometry or NULL id: 0
NOTICE: Vertices table for table public.edge_table is: public.edge_table_vertices_pgr
NOTICE: -----

pgr_analyzeGraph
-----

OK
(1 row)

SELECT pgr_analyzeGraph('edge_table', 0.001);
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edge_table',0.001,'the_geom','id','source','target','true')
NOTICE: Performing checks, pelase 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)

```

The examples use the *Sample Data* network.

### See Also

- *Routing Topology* for an overview of a topology for routing algorithms.
- *pgr\_analyzeOneway* to analyze directionality of the edges.
- *pgr\_createVerticesTable* to reconstruct the vertices table based on the source and target information.
- *pgr\_nodeNetwork* to create nodes to a not noded edge table.



## 5.1.4 pgr\_analyzeOneway

### Name

pgr\_analyzeOneway — Analyzes oneway Sstreets and identifies flipped segments.

### Synopsis

This function analyzes oneway streets in a graph and identifies any flipped segments.

```
text pgr_analyzeOneway(geom_table text,
                       text[] s_in_rules, text[] s_out_rules,
                       text[] t_in_rules, text[] t_out_rules,
                       text oneway='oneway', text source='source', text target='target',
                       boolean two_way_if_null=true);
```

### 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 a *sink* if all edges enter the node but none exit that node. For a *source* 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.

- Use [pgr\\_createVerticesTable](#) to create the vertices table.
- Use [pgr\\_createTopology](#) to create the topology and the vertices table.

### 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.

---

**Note:** It is strongly recommended to use the named notation. See *pgr\_createVerticesTable* or *pgr\_createTopology* for examples.

---

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 vertices table can be created with *pgr\_createVerticesTable* or *pgr\_createTopology*

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. See *pgr\_analyzeGraph*.
- chk** integer Indicator that the vertex might have a problem. See *pgr\_analyzeGraph*.
- 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.

### History

- New in version 2.0.0

### Examples

```
SELECT pgr_analyzeOneway('edge_table',
ARRAY['', 'B', 'TF'],
ARRAY['', 'B', 'FT'],
ARRAY['', 'B', 'FT'],
ARRAY['', 'B', 'TF'],
oneway:='dir');
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edge_table','{"",B,TF}','{"",B,FT}','{"",B,FT}','{"",B,TF}','dir','sou
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)
```

The queries use the *Sample Data* network.

## See Also

- *Routing Topology* for an overview of a topology for routing algorithms.
- *Graph Analytics* for an overview of the analysis of a graph.
- *pgr\_analyzeGraph* to analyze the edges and vertices of the edge table.
- *pgr\_createVerticesTable* to reconstruct the vertices table based on the source and target information.

## 5.1.5 pgr\_nodeNetwork

### Name

pgr\_nodeNetwork - Nodes an network edge table.

**Author** Nicolas Ribot

**Copyright** Nicolas Ribot, The source code is released under the MIT-X license.

### Synopsis

The function reads edges from a not “noded” network table and writes the “noded” edges into a new table.

```
text pgr_nodenetwork(text edge_table, float8, tolerance,
                    text id='id', text the_geom='the_geom', text table_ending='noded')
```

### Description

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** integer Empty source column to be used with *pgr\_createTopology* function
- target** integer Empty target column to be used with *pgr\_createTopology* function
- the\_geom** geometry Geometry column of the noded network

## History

- New in version 2.0.0

## Example

Let's create the topology for the data in *Sample Data*

```

SELECT pgr_createTopology('edge_table', 0.001);
NOTICE: PROCESSING:
NOTICE: pgr_createTopology('edge_table',0.001,'the_geom','id','source','target','true')
NOTICE: Performing checks, pelase 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.edge_table is: public.edge_table_vertices_pgr
NOTICE: -----
pgr_createtopology
-----
OK
(1 row)

```

Now we can analyze the network.

```

SELECT pgr_analyzegraph('edge_table', 0.001);
NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edge_table',0.001,'the_geom','id','source','target','true')
NOTICE: Performing checks, pelase 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)

```

The analysis tell us that the network has a gap and and an intersection. We try to fix the problem using:

```

SELECT pgr_nodeNetwork('edge_table', 0.001);
NOTICE: PROCESSING:
NOTICE: pgr_nodeNetwork('edge_table',0.001,'the_geom','id','noded')
NOTICE: Performing checks, pelase wait .....
NOTICE: Processing, pelase wait .....

```

```

NOTICE:      Splitted Edges: 3
NOTICE:      Untouched Edges: 15
NOTICE:      Total original Edges: 18
NOTICE:      Edges generated: 6
NOTICE:      Untouched Edges: 15
NOTICE:      Total New segments: 21
NOTICE:      New Table: public.edge_table_noded
NOTICE:      -----
pgr_nodenetwork
-----
OK
(1 row)

```

Inspecting the generated table, we can see that edges 13,14 and 18 has been segmented

```

SELECT old_id,sub_id FROM edge_table_noded ORDER BY old_id,sub_id;
old_id | sub_id
-----+-----
1      |      1
2      |      1
3      |      1
4      |      1
5      |      1
6      |      1
7      |      1
8      |      1
9      |      1
10     |      1
11     |      1
12     |      1
13     |      1
13     |      2
14     |      1
14     |      2
15     |      1
16     |      1
17     |      1
18     |      1
18     |      2
(21 rows)

```

We can create the topology of the new network

```

SELECT pgr_createTopology('edge_table_noded', 0.001);
NOTICE:  PROCESSING:
NOTICE:  pgr_createTopology('edge_table_noded',0.001,'the_geom','id','source','target','true')
NOTICE:  Performing checks, pelase wait .....
NOTICE:  Creating Topology, Please wait...
NOTICE:  -----> TOPOLOGY CREATED FOR 21 edges
NOTICE:  Rows with NULL geometry or NULL id: 0
NOTICE:  Vertices table for table public.edge_table_noded is: public.edge_table_noded_vertices_pgr
NOTICE:  -----
pgr_createtopology
-----
OK
(1 row)

```

Now let's analyze the new topology

```

SELECT pgr_analyzegraph('edge_table_noded', 0.001);
NOTICE:  PROCESSING:
NOTICE:  pgr_analyzeGraph('edge_table_noded',0.001,'the_geom','id','source','target','true')
NOTICE:  Performing checks, pelase wait...
NOTICE:  Analyzing for dead ends. Please wait...

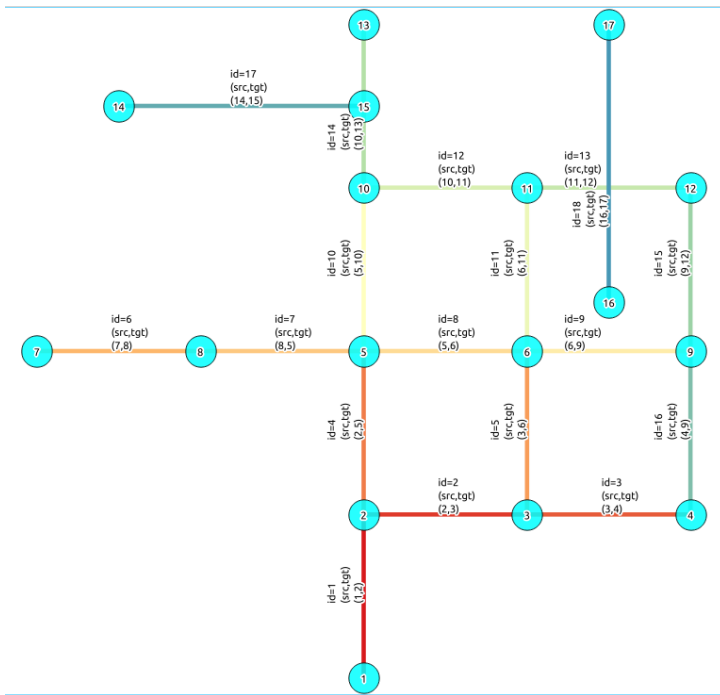
```

```

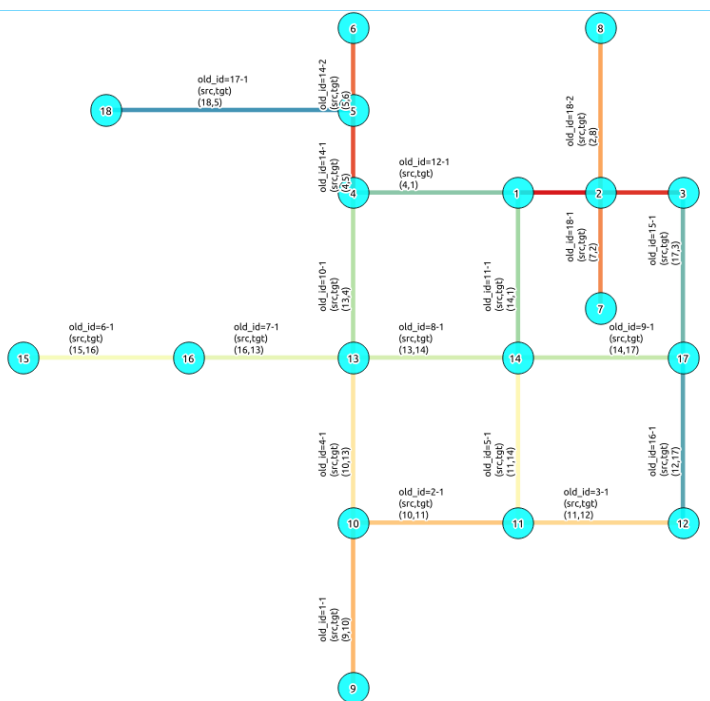
NOTICE: Analyzing for gaps. Please wait...
NOTICE: Analyzing for isolated edges. Please wait...
NOTICE: Analyzing for ring geometries. Please wait...
NOTICE: Analyzing for intersections. Please wait...
NOTICE: ANALYSIS RESULTS FOR SELECTED EDGES:
NOTICE: Isolated segments: 0
NOTICE: Dead ends: 6
NOTICE: Potential gaps found near dead ends: 0
NOTICE: Intersections detected: 0
NOTICE: Ring geometries: 0
pgr_createtopology
-----
OK
(1 row)
    
```

## Images

### Before Image



After Image



Comparing the results

Comparing with the Analysis in the original edge\_table, we see that.

	Before	After
Table name	edge_table	edge_table_noded
Fields	All original fields	Has only basic fields to do a topology analysis
Dead ends	<ul style="list-style-type: none"> <li>Edges with 1 dead end: 1,6,24</li> <li>Edges with 2 dead ends 17,18</li> </ul> Edge 17's right node is a dead end because there is no other edge sharing that same node. (cnt=1)	Edges with 1 dead end: 1-1 ,6-1,14-2, 18-1 17-1 18-2
Isolated segments	two isolated segments: 17 and 18 both they have 2 dead ends	<b>No Isolated segments</b> <ul style="list-style-type: none"> <li>Edge 17 now shares a node with edges 14-1 and 14-2</li> <li>Edges 18-1 and 18-2 share a node with edges 13-1 and 13-2</li> </ul>
Gaps	There is a gap between edge 17 and 14 because edge 14 is near to the right node of edge 17	Edge 14 was segmented Now edges: 14-1 14-2 17 share the same node The tolerance value was taken in account
Intersections	Edges 13 and 18 were intersecting	Edges were segmented, So, now in the interection's point there is a node and the following edges share it: 13-1 13-2 18-1 18-2

Now, we are going to include the segments 13-1, 13-2 14-1, 14-2 ,18-1 and 18-2 into our edge-table, copying the data for dir,cost,and reverse cost with the following steps:

- Add a column old\_id into edge\_table, this column is going to keep track the id of the original edge
- Insert only the segmented edges, that is, the ones whose max(sub\_id) >1

```
alter table edge_table drop column if exists old_id;
alter table edge_table add column old_id integer;
insert into edge_table (old_id,dir,cost,reverse_cost,the_geom)
  (with
   segmented as (select old_id,count(*) as i from edge_table_noded group by old_id)
  select  segments.old_id,dir,cost,reverse_cost,segments.the_geom
         from edge_table as edges join edge_table_noded as segments on (edges.id = segment
         where edges.id in (select old_id from segmented where i>1) );
```

We recreate the topology:

```
SELECT pgr_createTopology('edge_table', 0.001);

NOTICE:  PROCESSING:
NOTICE:  pgr_createTopology('edge_table',0.001,'the_geom','id','source','target','true')
NOTICE:  Performing checks, pelase wait .....
NOTICE:  Creating Topology, Please wait...
NOTICE:  -----> TOPOLOGY CREATED FOR 24 edges
NOTICE:  Rows with NULL geometry or NULL id: 0
NOTICE:  Vertices table for table public.edge_table is: public.edge_table_vertices_pgr
NOTICE:  -----
pgr_createtopology
-----
OK
(1 row)
```

To get the same analysis results as the topology of edge\_table\_noded, we do the following query:

```
SELECT pgr_analyzegraph('edge_table', 0.001,rows_where:='id not in (select old_id from edge_table
NOTICE:  PROCESSING:
NOTICE:  pgr_analyzeGraph('edge_table',0.001,'the_geom','id','source','target',
                        'id not in (select old_id from edge_table where old_id is not null)')
NOTICE:  Performing checks, pelase 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: 6
NOTICE:  Potential gaps found near dead ends: 0
NOTICE:  Intersections detected: 0
NOTICE:  Ring geometries: 0
pgr_createtopology
-----
OK
(1 row)
```

To get the same analysis results as the original edge\_table, we do the following query:

```
SELECT pgr_analyzegraph('edge_table', 0.001,rows_where:='old_id is null')
NOTICE:  PROCESSING:
NOTICE:  pgr_analyzeGraph('edge_table',0.001,'the_geom','id','source','target','old_id is null')
NOTICE:  Performing checks, pelase 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_createtopology
-----
OK
(1 row)

```

Or we can analyze everything because, maybe edge 18 is an overpass, edge 14 is an under pass and there is also a street level junction, and the same happens with edges 17 and 13.

```

SELECT pgr_analyzegraph('edge_table', 0.001);

NOTICE: PROCESSING:
NOTICE: pgr_analyzeGraph('edge_table',0.001,'the_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: 5
NOTICE: Ring geometries: 0
pgr_createtopology
-----
OK
(1 row)

```

## See Also

*Routing Topology* for an overview of a topology for routing algorithms. *pgr\_analyzeOneway* to analyze directionality of the edges. *pgr\_createTopology* to create a topology based on the geometry. *pgr\_analyzeGraph* to analyze the edges and vertices of the edge table.



---

## Routing functions

---

### *Routing Functions*

- *pgr\_apspJohnson*- All Pairs Shortest Path, Johnson's Algorithm
- *pgr\_apspWarshall* - All Pairs Shortest Path, Floyd-Warshall Algorithm
- *pgr\_astar* - Shortest Path A\*
- *pgr\_bdAstar* - Bi-directional A\* Shortest Path
- *pgr\_bdDijkstra* - Bi-directional Dijkstra Shortest Path
- *pgr\_dijkstra* - Shortest Path Dijkstra
- *pgr\_driving\_distance* - Driving Distance
- *pgr\_kDijkstra* - Multiple destination Shortest Path Dijkstra
- *pgr\_ksp* - K-Shortest Path
- *pgr\_trsp* - Turn Restriction Shortest Path (TRSP)
- *pgr\_tsp* - Traveling Sales Person

## 6.1 Routing Functions

- *pgr\_apspJohnson*- All Pairs Shortest Path, Johnson's Algorithm
- *pgr\_apspWarshall* - All Pairs Shortest Path, Floyd-Warshall Algorithm
- *pgr\_astar* - Shortest Path A\*
- *pgr\_bdAstar* - Bi-directional A\* Shortest Path
- *pgr\_bdDijkstra* - Bi-directional Dijkstra Shortest Path
- *pgr\_dijkstra* - Shortest Path Dijkstra
- *pgr\_drivingDistance* - Driving distance
- *pgr\_kDijkstra* - Multiple destination Shortest Path Dijkstra
- *pgr\_ksp* - K-Shortest Path
- *pgr\_tsp* - Traveling Sales Person
- *pgr\_trsp* - Turn Restriction Shortest Path (TRSP)

## 6.1.1 pgr\_apspJohnson - All Pairs Shortest Path, Johnson's Algorithm

### Name

pgr\_apspJohnson - Returns all costs for each pair of nodes in the graph.

### Synopsis

Johnson's algorithm is a way to find the shortest paths between all pairs of vertices in a sparse, edge weighted, directed graph. Returns a set of *pgr\_costResult* (seq, id1, id2, cost) rows for every pair of nodes in the graph.

```
pgr_costResult[] pgr_apspJohnson(sql text);
```

### Description

**sql** a SQL query that should return the edges for the graph that will be analyzed:

```
SELECT source, target, cost FROM edge_table;
```

**source** int4 identifier of the source vertex for this edge

**target** int4 identifier of the target vertex for this edge

**cost** float8 a positive value for the cost to traverse this edge

Returns set of *pgr\_costResult*[]):

**seq** row sequence

**id1** source node ID

**id2** target node ID

**cost** cost to traverse from id1 to id2

### History

- New in version 2.0.0

### Examples

```
SELECT seq, id1 AS from, id2 AS to, cost
  FROM pgr_apspJohnson(
    'SELECT source, target, cost FROM edge_table'
  );
```

seq	from	to	cost
0	1	1	0
1	1	2	1
2	1	5	2
3	1	6	3
[...]			

The query uses the *Sample Data* network.

## See Also

- *pgr\_costResult[]*
- *pgr\_apspWarshall* - All Pairs Shortest Path, Floyd-Warshall Algorithm
- [http://en.wikipedia.org/wiki/Johnson%27s\\_algorithm](http://en.wikipedia.org/wiki/Johnson%27s_algorithm)

## 6.1.2 pgr\_apspWarshall - All Pairs Shortest Path, Floyd-Warshall Algorithm

### Name

`pgr_apspWarshall` - Returns all costs for each pair of nodes in the graph.

### Synopsis

The Floyd-Warshall algorithm (also known as Floyd's algorithm and other names) is a graph analysis algorithm for finding the shortest paths between all pairs of nodes in a weighted graph. Returns a set of *pgr\_costResult* (`seq, id1, id2, cost`) rows for every pair of nodes in the graph.

```
pgr_costResult[] pgr_apspWarshall(sql text, directed boolean, reverse_cost boolean);
```

### Description

**sql** a SQL query that should return the edges for the graph that will be analyzed:

```
SELECT id, source, target, cost FROM edge_table;
```

**id** int4 identifier of the edge

**source** int4 identifier of the source vertex for this edge

**target** int4 identifier of the target vertex for this edge

**cost** float8 a positive value for the cost to traverse this edge

**directed** true if the graph is directed

**reverse\_cost** if true, the `reverse_cost` column of the SQL generated set of rows will be used for the cost of the traversal of the edge in the opposite direction.

Returns set of *pgr\_costResult[]*:

**seq** row sequence

**id1** source node ID

**id2** target node ID

**cost** cost to traverse from `id1` to `id2`

### History

- New in version 2.0.0

### Examples

```

SELECT seq, id1 AS from, id2 AS to, cost
  FROM pgr_apspWarshall(
    'SELECT id, source, target, cost FROM edge_table',
    false, false
  );

```

seq	from	to	cost
0	1	1	0
1	1	2	1
2	1	3	0
3	1	4	-1
[...]			

The query uses the *Sample Data* network.

## See Also

- *pgr\_costResult[]*
- *pgr\_apspJohnson* - All Pairs Shortest Path, Johnson's Algorithm
- [http://en.wikipedia.org/wiki/Floyd%E2%80%93Warshall\\_algorithm](http://en.wikipedia.org/wiki/Floyd%E2%80%93Warshall_algorithm)

## 6.1.3 pgr\_astar - Shortest Path A\*

### Name

`pgr_astar` — Returns the shortest path using A\* algorithm.

### Synopsis

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. Returns a set of *pgr\_costResult* (seq, id1, id2, cost) rows, that make up a path.

```

pgr_costResult[] pgr_astar(sql text, source integer, target integer,
                          directed boolean, has_rcost boolean);

```

### Description

**sql** a SQL query, which should return a set of rows with the following columns:

```

SELECT id, source, target, cost, x1, y1, x2, y2 [,reverse_cost] FROM edge_table

```

**id** int4 identifier of the edge

**source** int4 identifier of the source vertex

**target** int4 identifier of the target vertex

**cost** float8 value, of the edge traversal cost. A negative cost will prevent the edge from being inserted in the graph.

**x1** x coordinate of the start point of the edge

**y1** y coordinate of the start point of the edge

**x2** x coordinate of the end point of the edge

**y2** y coordinate of the end point of the edge

**reverse\_cost** (optional) the cost for the reverse traversal of the edge. This is only used when the `directed` and `has_rcost` parameters are `true` (see the above remark about negative costs).

**source** `int4` id of the start point

**target** `int4` id of the end point

**directed** `true` if the graph is directed

**has\_rcost** if `true`, the `reverse_cost` column of the SQL generated set of rows will be used for the cost of the traversal of the edge in the opposite direction.

Returns set of `pgr_costResult[]`:

**seq** row sequence

**id1** node ID

**id2** edge ID (-1 for the last row)

**cost** cost to traverse from `id1` using `id2`

## History

- Renamed in version 2.0.0

## Examples

- Without `reverse_cost`

```

SELECT seq, id1 AS node, id2 AS edge, cost
  FROM pgr_astar(
    'SELECT id, source, target, cost, x1, y1, x2, y2 FROM edge_table',
    4, 1, false, false
  );

```

seq	node	edge	cost
0	4	16	1
1	9	9	1
2	6	8	1
3	5	4	1
4	2	1	1
5	1	-1	0

(6 rows)

- With `reverse_cost`

```

SELECT seq, id1 AS node, id2 AS edge, cost
  FROM pgr_astar(
    'SELECT id, source, target, cost, x1, y1, x2, y2, reverse_cost FROM edge_table',
    4, 1, true, true
  );

```

seq	node	edge	cost
0	4	3	1
1	3	2	1
2	2	1	1
3	1	-1	0

(4 rows)

The queries use the *Sample Data* network.

## See Also

- *pgr\_costResult[]*
- [http://en.wikipedia.org/wiki/A\\*\\_search\\_algorithm](http://en.wikipedia.org/wiki/A*_search_algorithm)

## 6.1.4 pgr\_bdAstar - Bi-directional A\* Shortest Path

### Name

`pgr_bdAstar` - Returns the shortest path using Bidirectional A\* algorithm.

### Synopsis

This is a bi-directional A\* search algorithm. It searches from the source toward the destination and at the same time from the destination to the source and terminates when these two searches meet in the middle. Returns a set of *pgr\_costResult* (seq, id1, id2, cost) rows, that make up a path.

```
pgr_costResult[] pgr_bdAstar(sql text, source integer, target integer,
                             directed boolean, has_rcost boolean);
```

### Description

**sql** a SQL query, which should return a set of rows with the following columns:

```
SELECT id, source, target, cost, x1, y1, x2, y2 [,reverse_cost] FROM edge_table
```

**id** int4 identifier of the edge

**source** int4 identifier of the source vertex

**target** int4 identifier of the target vertex

**cost** float8 value, of the edge traversal cost. A negative cost will prevent the edge from being inserted in the graph.

**x1** x coordinate of the start point of the edge

**y1** y coordinate of the start point of the edge

**x2** x coordinate of the end point of the edge

**y2** y coordinate of the end point of the edge

**reverse\_cost** (optional) the cost for the reverse traversal of the edge. This is only used when the `directed` and `has_rcost` parameters are `true` (see the above remark about negative costs).

**source** int4 id of the start point

**target** int4 id of the end point

**directed** true if the graph is directed

**has\_rcost** if true, the `reverse_cost` column of the SQL generated set of rows will be used for the cost of the traversal of the edge in the opposite direction.

Returns set of *pgr\_costResult[]*:

**seq** row sequence

**id1** node ID

**id2** edge ID (-1 for the last row)



**cost** cost to traverse from id1 using id2

**Warning:** You must reconnect to the database after CREATE EXTENSION pgrouting. Otherwise the function will return Error computing path: std::bad\_alloc.

## History

- New in version 2.0.0

## Examples

- Without reverse\_cost

```
SELECT seq, id1 AS node, id2 AS edge, cost
  FROM pgr_bdAstar(
    'SELECT id, source, target, cost, x1, y1, x2, y2 FROM edge_table',
    4, 10, false, false
  );
```

seq	node	edge	cost
0	4	3	0
1	3	5	1
2	6	11	1
3	11	12	0
4	10	-1	0

(5 rows)

- With reverse\_cost

```
SELECT seq, id1 AS node, id2 AS edge, cost
  FROM pgr_bdAstar(
    'SELECT id, source, target, cost, x1, y1, x2, y2, reverse_cost FROM edge_table',
    4, 10, true, true
  );
```

seq	node	edge	cost
0	4	3	1
1	3	5	1
2	6	8	1
3	5	10	1
4	10	-1	0

(5 rows)

The queries use the *Sample Data* network.

## See Also

- [pgr\\_costResult\[\]](#)
- [pgr\\_bdDijkstra - Bi-directional Dijkstra Shortest Path](#)
- [http://en.wikipedia.org/wiki/Bidirectional\\_search](http://en.wikipedia.org/wiki/Bidirectional_search)
- [http://en.wikipedia.org/wiki/A\\*\\_search\\_algorithm](http://en.wikipedia.org/wiki/A*_search_algorithm)

## 6.1.5 pgr\_bdDijkstra - Bi-directional Dijkstra Shortest Path

### Name

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

### Synopsis

This is a bi-directional Dijkstra search algorithm. It searches from the source toward the destination and at the same time from the destination to the source and terminates when these two searches meet in the middle. Returns a set of *pgr\_costResult* (seq, id1, id2, cost) rows, that make up a path.

```
pgr_costResult[] pgr_bdDijkstra(sql text, source integer, target integer,
                                directed boolean, has_rcost boolean);
```

### Description

`sql` a SQL query, which should return a set of rows with the following columns:

```
SELECT id, source, target, cost [,reverse_cost] FROM edge_table
```

**id** int4 identifier of the edge

**source** int4 identifier of the source vertex

**target** int4 identifier of the target vertex

**cost** float8 value, of the edge traversal cost. A negative cost will prevent the edge from being inserted in the graph.

**reverse\_cost** (optional) the cost for the reverse traversal of the edge. This is only used when the `directed` and `has_rcost` parameters are `true` (see the above remark about negative costs).

**source** int4 id of the start point

**target** int4 id of the end point

**directed** true if the graph is directed

**has\_rcost** if true, the `reverse_cost` column of the SQL generated set of rows will be used for the cost of the traversal of the edge in the opposite direction.

Returns set of *pgr\_costResult*[:

**seq** row sequence

**id1** node ID

**id2** edge ID (-1 for the last row)

**cost** cost to traverse from `id1` using `id2`

### History

- New in version 2.0.0

### Examples

- Without `reverse_cost`

```

SELECT seq, id1 AS node, id2 AS edge, cost
FROM pgr_bdDijkstra(
  'SELECT id, source, target, cost FROM edge_table',
  4, 10, false, false
);

```

seq	node	edge	cost
0	4	3	0
1	3	5	1
2	6	11	1
3	11	12	0
4	10	-1	0

(5 rows)

- With `reverse_cost`

```

SELECT seq, id1 AS node, id2 AS edge, cost
FROM pgr_bdDijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  4, 10, true, true
);

```

seq	node	edge	cost
0	4	3	1
1	3	2	1
2	2	4	1
3	5	10	1
4	10	-1	0

(5 rows)

The queries use the *Sample Data* network.

## See Also

- `pgr_costResult[]`
- `pgr_bdAstar` - Bi-directional A\* Shortest Path
- [http://en.wikipedia.org/wiki/Bidirectional\\_search](http://en.wikipedia.org/wiki/Bidirectional_search)
- [http://en.wikipedia.org/wiki/Dijkstra%27s\\_algorithm](http://en.wikipedia.org/wiki/Dijkstra%27s_algorithm)

## 6.1.6 pgr\_dijkstra - Shortest Path Dijkstra

### Version 2.0 (deprecated)

- `pgr_dijkstra` - Shortest Path Dijkstra

### Version 2.1

- `pgr_dijkstra` - Shortest Path Dijkstra

## 6.1.7 The problem definition

Given the following query:

```
pgr_dijkstra(sql, start_vid, end_vid, directed)
```

where  $sql = \{(id_i, source_i, target_i, cost_i, reverse\_cost_i)\}$

and

- $source = \bigcup source_i$ ,
- $target = \bigcup target_i$ ,

The graphs are defined as follows:

### Directed graph

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

- the set of vertices  $V$ 
  - $V = source \cup target \cup start_{vid} \cup end_{vid}$

- the set of edges  $E$

$$- E = \begin{cases} \{(source_i, target_i, cost_i) \text{ when } cost \geq 0\} \\ \text{if } reverse\_cost = \\ \cup \\ \{(source_i, target_i, cost_i) \text{ when } cost \geq 0\} \\ \{(target_i, source_i, reverse\_cost_i) \text{ when } reverse\_cost_i \geq 0\} \\ \text{if } reverse\_cost \neq \end{cases}$$

### Undirected graph

The weighted undirected graph,  $G_u(V, E)$ , is defined by:

- the set of vertices  $V$ 
  - $V = source \cup target \cup start_{vid} \cup end_{vid}$

- the set of edges  $E$

$$- E = \begin{cases} \{(source_i, target_i, cost_i) \text{ when } cost \geq 0\} \\ \cup \\ \{(target_i, source_i, cost_i) \text{ when } cost \geq 0\} \\ \text{if } reverse\_cost = \\ \cup \\ \{(source_i, target_i, cost_i) \text{ when } cost \geq 0\} \\ \cup \\ \{(target_i, source_i, cost_i) \text{ when } cost \geq 0\} \\ \cup \\ \{(target_i, source_i, reverse\_cost_i) \text{ when } reverse\_cost_i \geq 0\} \\ \cup \\ \{(source_i, target_i, reverse\_cost_i) \text{ when } reverse\_cost_i \geq 0\} \\ \text{if } reverse\_cost \neq \end{cases}$$

### The problem

Given:

- $start_{vid} \in V$  a starting vertex
- $end_{vid} \in V$  an ending vertex
- $G(V, E) = \begin{cases} G_d(V, E) & \text{if } directed = true \\ G_u(V, E) & \text{if } directed = false \end{cases}$

Then:

$$\text{pgr\_dijkstra}(\text{sql}, \text{start}_{vid}, \text{end}_{vid}, \text{directed}) = \begin{cases} \text{shortest path } \pi \text{ between } \text{start}_{vid} \text{ and } \text{end}_{vid} & \text{if } \exists \pi \\ \text{otherwise} & \end{cases}$$

$$\pi = \{(path_{i,node_i,edge_i,cost_i,agg\_cost_i})\}$$

where:

- $path_{i=i}$
- $path_{|\pi|=|\pi|}$
- $node_i \in V$
- $node_1 = \text{start}_{vid}$
- $node_{|\pi|} = \text{end}_{vid}$
- $\forall i \neq |\pi|, (node_i, node_{i+1}, cost_i) \in E$
- $edge_i = \begin{cases} id_{(node_i, node_{i+1}, cost_i)} & \text{when } i \neq |\pi| \\ -1 & \text{when } i = |\pi| \end{cases}$
- $cost_i = cost_{(node_i, node_{i+1})}$
- $agg\_cost_i = \begin{cases} 0 & \text{when } i = 1 \\ \sum_{k=1}^i cost_{(node_{k-1}, node_k)} & \text{when } i \neq 1 \end{cases}$

**In other words: The algorithm returns a the shortest path between  $\text{start}_{vid}$  and  $\text{end}_{vid}$  , if it exists, in terms of a sequence of**

- $path$  indicates the relative position in the path of the  $node$  or  $edge$ .
- $cost$  is the cost of the edge to be used to go to the next node.
- $agg\_cost$  is the cost from the  $\text{start}_{vid}$  up to the node.

If there is no path, the resulting set is empty.

## pg\_routing (V 2.0)- Shortest Path Dijkstra

### Name

`pgr_dijkstra` — Returns the shortest path using Dijkstra algorithm.

### Synopsis

Dijkstra's algorithm, conceived by Dutch computer scientist Edsger Dijkstra in 1956. It is a graph search algorithm that solves the single-source shortest path problem for a graph with non-negative edge path costs, producing a shortest path tree. Returns a set of *pgr\_costResult* (seq, id1, id2, cost) rows, that make up a path.

```
pgr_costResult[] pgr_dijkstra(text sql, integer source, integer target,
                             boolean directed, boolean has_rcost);
```

**Warning:** This signature is being deprecated in version 2.1, Please use it without the `has_rcost` flag instead:

```
pgr_dijkstra(sql, source, target, directed)
```

See *pgr\_dijkstra - Shortest Path Dijkstra*

## Description

**sql** a SQL query, which should return a set of rows with the following columns:

```
SELECT id, source, target, cost [,reverse_cost] FROM edge_table
```

**id** int4 identifier of the edge

**source** int4 identifier of the source vertex

**target** int4 identifier of the target vertex

**cost** float8 value, of the edge traversal cost. A negative cost will prevent the edge from being inserted in the graph.

**reverse\_cost** float8 (optional) the cost for the reverse traversal of the edge. This is only used when the `directed` and `has_rcost` parameters are `true` (see the above remark about negative costs).

**source** int4 id of the start point

**target** int4 id of the end point

**directed** true if the graph is directed

**has\_rcost** if `true`, the `reverse_cost` column of the SQL generated set of rows will be used for the cost of the traversal of the edge in the opposite direction.

Returns set of `pgr_costResult[]`:

**seq** row sequence

**id1** node ID

**id2** edge ID (-1 for the last row)

**cost** cost to traverse from `id1` using `id2`

## History

- Renamed in version 2.0.0

## Examples: Directed

- Without `reverse_cost`

```
SELECT seq, id1 AS node, id2 AS edge, cost
      FROM pgr_dijkstra(
          'SELECT id, source, target, cost, reverse_cost FROM edge_table',
          2,3, true, false
      );

seq | node | edge | cost
-----+-----+-----+-----
(0 rows)
```

- With `reverse_cost`

```
SELECT seq, id1 AS node, id2 AS edge, cost
      FROM pgr_dijkstra(
          'SELECT id, source, target, cost, reverse_cost FROM edge_table',
          2,3, true, true
      );

seq | node | edge | cost
-----+-----+-----+-----
```

```

0 | 2 | 4 | 1
1 | 5 | 8 | 1
2 | 6 | 9 | 1
3 | 9 | 16 | 1
4 | 4 | 3 | 1
5 | 3 | -1 | 0
(6 rows)

```

### Examples: Undirected

- Without `reverse_cost`

```

SELECT seq, id1 AS node, id2 AS edge, cost
  FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    2, 3, false, false
  );
seq | node | edge | cost
-----+-----+-----+-----
  0 | 2 | 4 | 1
  1 | 5 | 8 | 1
  2 | 6 | 5 | 1
  3 | 3 | -1 | 0
(4 rows)

```

- With `reverse_cost`

```

SELECT seq, id1 AS node, id2 AS edge, cost
  FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    2, 3, false, true
  );
seq | node | edge | cost
-----+-----+-----+-----
  0 | 2 | 2 | 1
  1 | 3 | -1 | 0
(2 rows)

```

The queries use the *Sample Data* network.

### See Also

- `pgr_costResult[]`
- [http://en.wikipedia.org/wiki/Dijkstra%27s\\_algorithm](http://en.wikipedia.org/wiki/Dijkstra%27s_algorithm)

### pgr\_dijkstra - Shortest Path Dijkstra

`pgr_dijkstra` — Returns the shortest path(s) using Dijkstra algorithm. In particular, the Dijkstra algorithm implemented by Boost.Graph.



Fig. 6.1: Boost Graph Inside

## Synopsis

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 (`start_vid`) to an ending vertex (`end_vid`). This implementation can be used with a directed graph and an undirected graph.

## Signatures

### Minimal signature

The minimal signature is for a **directed** graph from one `start_vid` to one `end_vid`:

```
pgr_dijkstra(text edges_sql, bigint start_vid, bigint end_vid)
  RETURNS SET OF (seq, path_seq, node, edge, cost, agg_cost) or EMPTY SET
```

### Dijkstra 1 to 1

This signature performs a Dijkstra from one `start_vid` to one `end_vid`:

- on a **directed** graph when `directed` flag is missing or is set to `true`.
- on an **undirected** graph when `directed` flag is set to `false`.

```
pgr_dijkstra(text edges_sql, bigint start_vid, bigint end_vid,
             boolean directed:=true);
  RETURNS SET OF (seq, path_seq, node, edge, cost, agg_cost) or EMPTY SET
```

### Dijkstra many to 1:

```
pgr_dijkstra(text edges_sql, array[ANY_INTEGER] start_vids, bigint end_vid,
             boolean directed:=true);
  RETURNS SET OF (seq, path_seq, start_vid, node, edge, cost, agg_cost) or EMPTY SET
```

This signature performs a Dijkstra from each `start_vid` in `start_vids` to one `end_vid`:

- on a **directed** graph when `directed` flag is missing or is set to `true`.
- on an **undirected** graph when `directed` flag is set to `false`.

Using this signature, will load once the graph and perform several one to one Dijkstra where the ending vertex is fixed. The result is the union of the results of the one to one dijkstra.

The extra `start_vid` in the result is used to distinguish to which path it belongs.

### Dijkstra 1 to many:

```
pgr_dijkstra(text edges_sql, bigint start_vid, array[ANY_INTEGER] end_vids,
             boolean directed:=true);
  RETURNS SET OF (seq, path_seq, end_vid, node, edge, cost, agg_cost) or EMPTY SET
```

This signature performs a Dijkstra from one `start_vid` to each `end_vid` in `end_vids`:

- on a **directed** graph when `directed` flag is missing or is set to `true`.
- on an **undirected** graph when `directed` flag is set to `false`.



Using this signature, will load once the graph and perform several 1 to 1 Dijkstra where the starting vertex is fixed. The result is the union of the results of the one to one dijkstra.

The extra `end_vid` in the result is used to distinguish to which path it belongs.

#### Dijkstra many to many:

```
pgr_dijkstra(text edges_sql, array[ANY_INTEGER] start_vids, array[ANY_INTEGER] end_vids,
             boolean directed:=true);
RETURNS SET OF (seq, path_seq, start_vid, end_vid, node, edge, cost, agg_cost) or EMPTY SET
```

This signature performs a Dijkstra from each `start_vid` in `start_vids` to each `end_vid` in `end_vids`:

- on a **directed** graph when `directed` flag is missing or is set to `true`.
- on an **undirected** graph when `directed` flag is set to `false`.

Using this signature, will load once the graph and perform all combinations for starting vertices and ending vertices.

The extra `start_vid` and `end_vid` in the result is used to distinguish to which path it belongs.

## Description of the Signatures

### Description of the SQL query

**edges\_sql** an SQL query, which should return a set of rows with the following columns:

**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 pont vertex of the edge.

**cost** ANY-NUMERICAL weight of the edge (*source*, *target*), if negative: edge (*source*, *target*) does not exist, therefore it's not part of the graph.

**reverse\_cost** ANY-NUMERICAL (optional) weight of the edge (*target*, *source*), if negative: edge (*target*, *source*) does not exist, therefore it's not part of the graph.

Where:

**ANY-INTEGER** smallint, int, bigint

**ANY-NUMERICAL** smallint, int, bigint, real, float

For example:

```
SELECT id, source, target, cost, reverse_cost from edge_table where geom && ST_Expand(ST_SetSRID(
```

### Description of the parameters of the signatures

**sql** SQL query as decribed above.

**start\_vid** BIGINT identifier of the starting vertex of the path.

**start\_vids** array[ANY-INTEGER] array of identifiers of starting vertices.

**end\_vid** BIGINT identifier of the ending vertex of the path.

**end\_vids** array[ANY-INTEGER] array of identifiers of ending vertices.

**directed** boolean (optional). When `false` the graph is considered as Undirected. Default is `true` which considers the graph as Directed.

### Description of the return values

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

- seq** INT isequential value starting from 1.
- path\_seq** INT relative position in the path. Has value 1 for the begining of a path.
- start\_vid** BIGINT id of the starting vertex. Used when multiple starting vetrices are in the query.
- end\_vid** BIGINT id of the ending vertex. Used when multiple ending vertices are in the query.
- node** BIGINT id of the node in the path from start\_vid to end\_v.
- edge** BIGINT id 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 total cost from start\_v to node.

### Examples

The examples of this section are based on the *Sample Data* network.

The examples include combinations from starting vertices 2 and 11 to ending vertices 3 and 5 in a directed and undirected graph with and with out reverse\_cost.

#### Examples for queries marked as directed with cost and reverse\_cost columns

The examples in this section use the following *Graph 1: Directed, with cost and reverse cost*

```

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    2, 3
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
  1 |         1 |    2 |    4 |    1 |         0
  2 |         2 |    5 |    8 |    1 |         1
  3 |         3 |    6 |    9 |    1 |         2
  4 |         4 |    9 |   16 |    1 |         3
  5 |         5 |    4 |    3 |    1 |         4
  6 |         6 |    3 |   -1 |    0 |         5
(6 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    2, 5
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
  1 |         1 |    2 |    4 |    1 |         0
  2 |         2 |    5 |   -1 |    0 |         1
(2 rows)

```

When you pass an array we get a combined result:

```

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    2, array[3,5]
);

```

```

seq | path_seq | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
 1 |          1 |        3 |    2 |    4 |    1 |          0
 2 |          2 |        3 |    5 |    8 |    1 |          1
 3 |          3 |        3 |    6 |    9 |    1 |          2
 4 |          4 |        3 |    9 |   16 |    1 |          3
 5 |          5 |        3 |    4 |    3 |    1 |          4
 6 |          6 |        3 |    3 |   -1 |    0 |          5
 7 |          1 |        5 |    2 |    4 |    1 |          0
 8 |          2 |        5 |    5 |   -1 |    0 |          1
(8 rows)

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  11, 3
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
 1 |          1 |   11 |   13 |    1 |          0
 2 |          2 |   12 |   15 |    1 |          1
 3 |          3 |    9 |   16 |    1 |          2
 4 |          4 |    4 |    3 |    1 |          3
 5 |          5 |    3 |   -1 |    0 |          4
(5 rows)

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  11, 5
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
 1 |          1 |   11 |   13 |    1 |          0
 2 |          2 |   12 |   15 |    1 |          1
 3 |          3 |    9 |    9 |    1 |          2
 4 |          4 |    6 |    8 |    1 |          3
 5 |          5 |    5 |   -1 |    0 |          4
(5 rows)

```

Some other combinations.

```

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  array[2,11], 5
);
seq | path_seq | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
 1 |          1 |          2 |    2 |    4 |    1 |          0
 2 |          2 |          2 |    5 |   -1 |    0 |          1
 3 |          1 |          11 |   11 |   13 |    1 |          0
 4 |          2 |          11 |   12 |   15 |    1 |          1
 5 |          3 |          11 |    9 |    9 |    1 |          2
 6 |          4 |          11 |    6 |    8 |    1 |          3
 7 |          5 |          11 |    5 |   -1 |    0 |          4
(7 rows)

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  array[2, 11], array[3,5]
);

```

```
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
 1 |          1 |          2 |          3 |      2 |      4 |      1 |          0
 2 |          2 |          2 |          3 |      5 |      8 |      1 |          1
 3 |          3 |          2 |          3 |      6 |      9 |      1 |          2
 4 |          4 |          2 |          3 |      9 |     16 |      1 |          3
 5 |          5 |          2 |          3 |      4 |      3 |      1 |          4
 6 |          6 |          2 |          3 |      3 |     -1 |      0 |          5
 7 |          1 |          2 |          5 |      2 |      4 |      1 |          0
 8 |          2 |          2 |          5 |      5 |     -1 |      0 |          1
 9 |          1 |         11 |          3 |     11 |     13 |      1 |          0
10 |          2 |         11 |          3 |     12 |     15 |      1 |          1
11 |          3 |         11 |          3 |      9 |     16 |      1 |          2
12 |          4 |         11 |          3 |      4 |      3 |      1 |          3
13 |          5 |         11 |          3 |      3 |     -1 |      0 |          4
14 |          1 |         11 |          5 |     11 |     13 |      1 |          0
15 |          2 |         11 |          5 |     12 |     15 |      1 |          1
16 |          3 |         11 |          5 |      9 |      9 |      1 |          2
17 |          4 |         11 |          5 |      6 |      8 |      1 |          3
18 |          5 |         11 |          5 |      5 |     -1 |      0 |          4
(18 rows)
```

Examples for queries marked as undirected with cost and reverse\_cost columns

The examples in this section use the following *Graph 2: Undirected, with cost and reverse cost*

```
SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    2, 3,
    false
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
 1 |          1 |      2 |      2 |      1 |          0
 2 |          2 |      3 |     -1 |      0 |          1
(2 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    2, 5,
    false
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
 1 |          1 |      2 |      4 |      1 |          0
 2 |          2 |      5 |     -1 |      0 |          1
(2 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    11, 3,
    false
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
 1 |          1 |     11 |     11 |      1 |          0
 2 |          2 |      6 |      5 |      1 |          1
 3 |          3 |      3 |     -1 |      0 |          2
(3 rows)
```

```

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  11, 5,
  false
);

```

seq	path_seq	node	edge	cost	agg_cost
1	1	11	11	1	0
2	2	6	8	1	1
3	3	5	-1	0	2

(3 rows)

```

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  array[2,11], 5,
  false
);

```

seq	path_seq	start_vid	node	edge	cost	agg_cost
1	1	2	2	4	1	0
2	2	2	5	-1	0	1
3	1	11	11	11	1	0
4	2	11	6	8	1	1
5	3	11	5	-1	0	2

(5 rows)

```

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, array[3,5],
  false
);

```

seq	path_seq	end_vid	node	edge	cost	agg_cost
1	1	3	2	2	1	0
2	2	3	3	-1	0	1
3	1	5	2	4	1	0
4	2	5	5	-1	0	1

(4 rows)

```

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  array[2, 11], array[3,5],
  false
);

```

seq	path_seq	start_vid	end_vid	node	edge	cost	agg_cost
1	1	2	3	2	2	1	0
2	2	2	3	3	-1	0	1
3	1	2	5	2	4	1	0
4	2	2	5	5	-1	0	1
5	1	11	3	11	11	1	0
6	2	11	3	6	5	1	1
7	3	11	3	3	-1	0	2
8	1	11	5	11	11	1	0
9	2	11	5	6	8	1	1
10	3	11	5	5	-1	0	2

(10 rows)

Examples for queries marked as directed with cost column

The examples in this section use the following *Graph 3: Directed, with cost*

```

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    2, 3
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
(0 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    2, 5
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
 1 |         1 |    2 |    4 |    1 |         0
 2 |         2 |    5 |   -1 |    0 |         1
(2 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    11, 3
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
(0 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    11, 5
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
(0 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    array[2,11], 5
);
seq | path_seq | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
 1 |         1 |          1 |    2 |    2 |    4 |         1 |         0
 2 |         2 |          2 |    2 |    5 |   -1 |         0 |         1
(2 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    2, array[3,5]
);
seq | path_seq | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
 1 |         1 |          1 |    5 |    2 |    4 |         1 |         0
 2 |         2 |          2 |    5 |    5 |   -1 |         0 |         1
(2 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    array[2, 11], array[3,5]

```

```
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
  1 |          1 |          2 |          5 |    2 |    4 |    1 |          0
  2 |          2 |          2 |          5 |    5 |   -1 |    0 |          1
(2 rows)
```

### Examples for queries marked as undirected with cost column

The examples in this section use the following *Graph 4: Undirected, with cost*

```
SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    2, 3,
    false
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
  1 |          1 |    2 |    4 |    1 |          0
  2 |          2 |    5 |    8 |    1 |          1
  3 |          3 |    6 |    5 |    1 |          2
  4 |          4 |    3 |   -1 |    0 |          3
(4 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    2, 5,
    false
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
  1 |          1 |    2 |    4 |    1 |          0
  2 |          2 |    5 |   -1 |    0 |          1
(2 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    11, 3,
    false
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
  1 |          1 |   11 |   11 |    1 |          0
  2 |          2 |    6 |    5 |    1 |          1
  3 |          3 |    3 |   -1 |    0 |          2
(3 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    11, 5,
    false
);
seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
  1 |          1 |   11 |   11 |    1 |          0
  2 |          2 |    6 |    8 |    1 |          1
  3 |          3 |    5 |   -1 |    0 |          2
(3 rows)
```

```

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    array[2,11], 5,
    false
);
seq | path_seq | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
 1 |         1 |         2 |    2 |    4 |    1 |         0
 2 |         2 |         2 |    5 |   -1 |    0 |         1
 3 |         1 |        11 |   11 |   11 |    1 |         0
 4 |         2 |        11 |    6 |    8 |    1 |         1
 5 |         3 |        11 |    5 |   -1 |    0 |         2
(5 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    2, array[3,5],
    false
);
seq | path_seq | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
 1 |         1 |         3 |    2 |    4 |    1 |         0
 2 |         2 |         3 |    5 |    8 |    1 |         1
 3 |         3 |         3 |    6 |    5 |    1 |         2
 4 |         4 |         3 |    3 |   -1 |    0 |         3
 5 |         1 |         5 |    2 |    4 |    1 |         0
 6 |         2 |         5 |    5 |   -1 |    0 |         1
(6 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost FROM edge_table',
    array[2, 11], array[3,5],
    false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
 1 |         1 |         2 |         3 |    2 |    4 |    1 |         0
 2 |         2 |         2 |         3 |    5 |    8 |    1 |         1
 3 |         3 |         2 |         3 |    6 |    5 |    1 |         2
 4 |         4 |         2 |         3 |    3 |   -1 |    0 |         3
 5 |         1 |         2 |         5 |    2 |    4 |    1 |         0
 6 |         2 |         2 |         5 |    5 |   -1 |    0 |         1
 7 |         1 |        11 |        3 |   11 |   11 |    1 |         0
 8 |         2 |        11 |         3 |    6 |    5 |    1 |         1
 9 |         3 |        11 |         3 |    3 |   -1 |    0 |         2
10 |         1 |        11 |         5 |   11 |   11 |    1 |         0
11 |         2 |        11 |         5 |    6 |    8 |    1 |         1
12 |         3 |        11 |         5 |    5 |   -1 |    0 |         2
(12 rows)

```

**Equivalences between signatures**

```

-- V2
SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    2, 3,
    true, -- directed flag
    true -- has_rcost
);

```



```

seq | id1 | id2 | cost
-----+-----+-----+-----
  0 |   2 |   4 |   1
  1 |   5 |   8 |   1
  2 |   6 |   9 |   1
  3 |   9 |  16 |   1
  4 |   4 |   3 |   1
  5 |   3 |  -1 |   0
(6 rows)

```

```
-- V3
```

```

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, 3,
  true    -- directed flag
);

```

```

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2,3
);

```

```

seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
  1 |         1 |   2 |   4 |   1 |         0
  2 |         2 |   5 |   8 |   1 |         1
  3 |         3 |   6 |   9 |   1 |         2
  4 |         4 |   9 |  16 |   1 |         3
  5 |         5 |   4 |   3 |   1 |         4
  6 |         6 |   3 |  -1 |   0 |         5
(6 rows)

```

```

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, array[3],
  true
);

```

```

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, array[3]
);

```

```

seq | path_seq | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
  1 |         1 |        3 |   2 |   4 |   1 |         0
  2 |         2 |        3 |   5 |   8 |   1 |         1
  3 |         3 |        3 |   6 |   9 |   1 |         2
  4 |         4 |        3 |   9 |  16 |   1 |         3
  5 |         5 |        3 |   4 |   3 |   1 |         4
  6 |         6 |        3 |   3 |  -1 |   0 |         5
(6 rows)

```

```

SELECT * FROM pgr_dijkstra(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  array[2], array[3],
  true
);

```

```

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    array[2], array[3]
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----+-----
  1 |         1 |         2 |       3 |    2 |    4 |    1 |         0
  2 |         2 |         2 |       3 |    5 |    8 |    1 |         1
  3 |         3 |         2 |       3 |    6 |    9 |    1 |         2
  4 |         4 |         2 |       3 |    9 |   16 |    1 |         3
  5 |         5 |         2 |       3 |    4 |    3 |    1 |         4
  6 |         6 |         2 |       3 |    3 |   -1 |    0 |         5
(6 rows)

```

**Equivalences between signatures**

```

-- V2
SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    2, 3,
    false, -- directed flag
    true   -- has_rcost
);

seq | id1 | id2 | cost
-----+-----+-----+-----
  0 |   2 |   2 |    1
  1 |   3 |  -1 |    0
(2 rows)

-- V3
SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    2, 3,
    false -- directed flag
);

seq | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
  1 |         1 |    2 |    2 |    1 |         0
  2 |         2 |    3 |   -1 |    0 |         1
(2 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    2, array[3],
    false
);

seq | path_seq | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
  1 |         1 |        3 |    2 |    2 |    1 |         0
  2 |         2 |        3 |    3 |   -1 |    0 |         1
(2 rows)

SELECT * FROM pgr_dijkstra(

```

```

        'SELECT id, source, target, cost, reverse_cost FROM edge_table',
        array[2], 3,
        false
    );
seq | path_seq | start_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
  1 |          1 |         2 |   2 |   2 |   1 |         0
  2 |          2 |         2 |   3 |  -1 |   0 |         1
(2 rows)

SELECT * FROM pgr_dijkstra(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    array[2], array[3],
    false
);
seq | path_seq | start_vid | end_vid | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
  1 |          1 |         2 |       3 |   2 |   2 |   1 |         0
  2 |          2 |         2 |       3 |   3 |  -1 |   0 |         1
(2 rows)

```

The queries use the *Sample Data* network.

## History

- Renamed in version 2.0.0
- Added functionality in version 2.1.0

## See Also

- [http://en.wikipedia.org/wiki/Dijkstra%27s\\_algorithm](http://en.wikipedia.org/wiki/Dijkstra%27s_algorithm)

## Indices and tables

- genindex
- search

## 6.1.8 Driving Distance

### Version 2.0 (deprecated)

- *pgr\_driving\_distance* - Driving Distance

### Version 2.1

- *pgr\_driving\_distance* - Driving Distance

## 6.1.9 Driving Distance post-processing

- *pgr\_alphaShape* - Alpha shape computation
- *pgr\_pointsAsPolygon* - Polygon around set of points

## pgr\_drivingDistance (V2.0)

### Name

pgr\_drivingDistance - Returns the driving distance from a start node.

### Synopsis

This function computes a Dijkstra shortest path solution then extracts the cost to get to each node in the network from the starting node. Using these nodes and costs it is possible to compute constant drive time polygons. Returns a set of *pgr\_costResult* (seq, id1, id2, cost) rows, that make up a list of accessible points.

```
pgr_costResult[] pgr_drivingDistance(text sql, integer source, double precision distance,
                                     boolean directed, boolean has_rcost);
```

**Warning:** This signature is being deprecated on version 2.1, Please use it without the `has_rcost` flag instead:

```
pgr_drivingDistance(sql, start_v, distance, directed)
```

See *pgr\_drivingDistance*

### Description

**sql** a SQL query, which should return a set of rows with the following columns:

```
SELECT id, source, target, cost [,reverse_cost] FROM edge_table
```

**id** int4 identifier of the edge

**source** int4 identifier of the source vertex

**target** int4 identifier of the target vertex

**cost** float8 value, of the edge traversal cost. A negative cost will prevent the edge from being inserted in the graph.

**reverse\_cost** (optional) the cost for the reverse traversal of the edge. This is only used when the `directed` and `has_rcost` parameters are `true` (see the above remark about negative costs).

**source** int4 id of the start point

**distance** float8 value in edge cost units (not in projection units - they might be different).

**directed** true if the graph is directed

**has\_rcost** if `true`, the `reverse_cost` column of the SQL generated set of rows will be used for the cost of the traversal of the edge in the opposite direction.

Returns set of *pgr\_costResult*[]):

**seq** row sequence

**id1** node ID

**id2** edge ID (this is probably not a useful item)

**cost** cost to get to this node ID

**Warning:** You must reconnect to the database after `CREATE EXTENSION pgrouting`. Otherwise the function will return `Error computing path: std::bad_alloc`.

## History

- Renamed in version 2.0.0

## Examples

- Without `reverse_cost`

```
SELECT seq, id1 AS node, cost
  FROM pgr_drivingDistance(
    'SELECT id, source, target, cost FROM edge_table',
    7, 1.5, false, false
  );
```

seq	node	cost
0	7	0
1	8	1

(2 rows)

- With `reverse_cost`

```
SELECT seq, id1 AS node, cost
  FROM pgr_drivingDistance(
    'SELECT id, source, target, cost, reverse_cost FROM edge_table',
    7, 1.5, true, true
  );
```

seq	node	cost
0	7	0
1	8	1

(5 rows)

The queries use the *Sample Data* network.

## See Also

- `pgr_alphaShape` - Alpha shape computation
- `pgr_pointsAsPolygon` - Polygon around set of points

## pgr\_drivingDistance

### Name

`pgr_drivingDistance` - Returns the driving distance from a start node.



Fig. 6.2: Boost Graph Inside

## Synopsis

Using Dijkstra algorithm, extracts all the nodes that have costs less than or equal to the value `distance`. The edges extracted will conform the corresponding spanning tree.

### The minimal signature:

```
pgr_drivingDistance(sql text, start_v bigint, distance float8)
RETURNS SET OF (seq, node, edge, cost, agg_cost)
```

### Driving Distance from a single starting point:

```
pgr_drivingDistance(sql text, start_vid bigint, distance float8, directed boolean)
RETURNS SET OF (seq, node, edge, cost, agg_cost)
```

### Driving Distance from a multiple starting points:

```
pgr_drivingDistance(sql text, start_vids anyarray, distance float8,
    directed boolean default true,
    equicost boolean default false)
RETURNS SET OF (seq, start_vid, node, edge, cost, agg_cost)
```

## Description of the SQL query

`sql` a SQL query, which should return a set of rows with the following columns:

```
SELECT id, source, target, cost [,reverse_cost] FROM edge_table
```

**id** ANY-INTEGERS identifier of the edge.

**source** ANY-INTEGERS identifier of the source vertex of the edge.

**target** ANY-INTEGERS identifier of the target vertex of the edge.

**cost** ANY-NUMERICAL value of the edge traversal cost. A negative cost will prevent the edge (`source`, `target`) from being inserted in the graph.

**reverse\_cost** ANY-NUMERICAL (optional) the value for the reverse traversal of the edge. A negative cost will prevent the edge (`target`, `source`) from being inserted in the graph.

Where:

**ANY-INTEGERS** smallint, int, bigint

**ANY-NUMERICAL** smallint, int, bigint, real, float

## Description of the parameters of the signatures

**sql** SQL query as described above.

**start\_v** BIGINT id of the starting vertex.

**start\_v** array[ANY-INTEGERS] array of id of starting vertices.

**distance** FLOAT Upper limit for the inclusion of the node in the result.

**directed** boolean (optional). When `false` the graph is considered as Undirected. Default is `true` which considers the graph as Directed.

**equicost** boolean (optional). When `true` the node will only appear in the closest `start_v` list. Default is `false` which resembles several calls using the single starting point signatures. Tie brakes are arbitrarily.

### Description of the return values

Returns set of (`seq` [, `start_v`], `node`, `edge`, `cost`, `agg_cost`)

**seq** INT row sequence.

**start\_v** BIGINT id of the starting vertex. Used when multiple starting vertices are in the query.

**node** BIGINT id of the node within the limits from `start_v`.

**edge** BIGINT id of the edge used to arrive to node. 0 when the node is the `start_v`.

**cost** FLOAT cost to traverse edge.

**agg\_cost** FLOAT total cost from `start_v` to node.

### Examples for queries marked as `directed` with `cost` and `reverse_cost` columns

The examples in this section use the following *Graph 1: Directed, with cost and reverse cost*

```

SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, 3
);
seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----
  1 |    1 |    1 |    1 |         1
  2 |    2 |   -1 |    0 |         0
  3 |    5 |    4 |    1 |         1
  4 |    6 |    8 |    1 |         2
  5 |   11 |   12 |    1 |         3
  6 |   10 |   10 |    1 |         2
  7 |   13 |   14 |    1 |         3
  8 |    9 |    9 |    1 |         3
  9 |    7 |    6 |    1 |         3
 10 |    8 |    7 |    1 |         2
(10 rows)

SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  13, 3
);
seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----
  1 |    2 |    4 |    1 |         3
  2 |    5 |   10 |    1 |         2
  3 |    6 |    8 |    1 |         3
  4 |   11 |   12 |    1 |         2
  5 |   10 |   14 |    1 |         1
  6 |   12 |   13 |    1 |         3
  7 |   13 |   -1 |    0 |         0
  8 |    8 |    7 |    1 |         3
(8 rows)

SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  array[2,13], 3
);
seq | from_v | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----

```

```

-----+-----+-----+-----+-----+-----
 1 |      2 |      1 |      1 |      1 |      1
 2 |      2 |      2 |     -1 |      0 |      0
 3 |      2 |      5 |      4 |      1 |      1
 4 |      2 |      6 |      8 |      1 |      2
 5 |      2 |     11 |     12 |      1 |      3
 6 |      2 |     10 |     10 |      1 |      2
 7 |      2 |     13 |     14 |      1 |      3
 8 |      2 |      9 |      9 |      1 |      3
 9 |      2 |      7 |      6 |      1 |      3
10 |      2 |      8 |      7 |      1 |      2
11 |     13 |      2 |      4 |      1 |      3
12 |     13 |      5 |     10 |      1 |      2
13 |     13 |      6 |      8 |      1 |      3
14 |     13 |     11 |     12 |      1 |      2
15 |     13 |     10 |     14 |      1 |      1
16 |     13 |     12 |     13 |      1 |      3
17 |     13 |     13 |     -1 |      0 |      0
18 |     13 |      8 |      7 |      1 |      3
(18 rows)

SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  array[2,13], 3, equicost:=true
);
seq | from_v | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
 1 |      2 |      1 |      1 |      1 |      1
 2 |      2 |      2 |     -1 |      0 |      0
 3 |      2 |      5 |      4 |      1 |      1
 4 |      2 |      6 |      8 |      1 |      2
 5 |      2 |      7 |      6 |      1 |      3
 6 |      2 |      8 |      7 |      1 |      2
 7 |      2 |      9 |      9 |      1 |      3
 8 |      2 |     10 |     10 |      1 |      2
 9 |      2 |     11 |     12 |      1 |      3
10 |     13 |     13 |     -1 |      0 |      0
11 |     13 |     12 |     13 |      1 |      3
(11 rows)

```

**Examples for queries marked as undirected with cost and reverse\_cost columns**

The examples in this section use the following *Graph 2: Undirected, with cost and reverse cost*

```

SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, 3, false
);
seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----
 1 |      1 |      1 |      1 |      1
 2 |      2 |     -1 |      0 |      0
 3 |      3 |      2 |      1 |      1
 4 |      4 |      3 |      1 |      2
 5 |      5 |      4 |      1 |      1
 6 |      6 |      8 |      1 |      2
 7 |     11 |     12 |      1 |      3
 8 |     10 |     10 |      1 |      2
 9 |     13 |     14 |      1 |      3
10 |      9 |     16 |      1 |      3
11 |      7 |      6 |      1 |      3
12 |      8 |      7 |      1 |      2

```



```
(12 rows)
```

```
SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  13, 3, false
);
```

seq	node	edge	cost	agg_cost
1	2	4	1	3
2	5	10	1	2
3	6	11	1	3
4	11	12	1	2
5	10	14	1	1
6	12	13	1	3
7	13	-1	0	0
8	8	7	1	3

```
(8 rows)
```

```
SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  array[2,13], 3, false
);
```

seq	from_v	node	edge	cost	agg_cost
1	2	1	1	1	1
2	2	2	-1	0	0
3	2	3	2	1	1
4	2	4	3	1	2
5	2	5	4	1	1
6	2	6	8	1	2
7	2	11	12	1	3
8	2	10	10	1	2
9	2	13	14	1	3
10	2	9	16	1	3
11	2	7	6	1	3
12	2	8	7	1	2
13	13	2	4	1	3
14	13	5	10	1	2
15	13	6	11	1	3
16	13	11	12	1	2
17	13	10	14	1	1
18	13	12	13	1	3
19	13	13	-1	0	0
20	13	8	7	1	3

```
(20 rows)
```

```
SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  array[2,13], 3, false, equicost:=true
);
```

seq	from_v	node	edge	cost	agg_cost
1	2	1	1	1	1
2	2	2	-1	0	0
3	2	3	2	1	1
4	2	4	3	1	2
5	2	5	4	1	1
6	2	6	8	1	2
7	2	7	6	1	3
8	2	8	7	1	2
9	2	9	16	1	3
10	2	10	10	1	2
11	2	11	12	1	3

12		13		13		-1		0		0
13		13		12		13		1		3
(13 rows)										

**Examples for queries marked as directed with cost column**

The examples in this section use the following *Graph 3: Directed, with cost*

```

SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost FROM edge_table',
  2, 3
);
seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----
 1 |  2 |  -1 |  0 |         0
 2 |  5 |  4 |  1 |         1
 3 |  6 |  8 |  1 |         2
 4 | 11 | 11 |  1 |         3
 5 | 10 | 10 |  1 |         2
 6 | 13 | 14 |  1 |         3
 7 |  9 |  9 |  1 |         3
(7 rows)

SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost FROM edge_table',
  13, 3
);
seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----
 1 | 13 |  -1 |  0 |         0
(1 row)

SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost FROM edge_table',
  array[2,13], 3
);
seq | from_v | node | edge | cost | agg_cost
-----+-----+-----+-----+-----
 1 |  2 |  2 |  -1 |  0 |         0
 2 |  2 |  5 |  4 |  1 |         1
 3 |  2 |  6 |  8 |  1 |         2
 4 |  2 | 11 | 11 |  1 |         3
 5 |  2 | 10 | 10 |  1 |         2
 6 |  2 | 13 | 14 |  1 |         3
 7 |  2 |  9 |  9 |  1 |         3
 8 | 13 | 13 | -1 |  0 |         0
(8 rows)

SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost FROM edge_table',
  array[2,13], 3, equicost:=true
);
seq | from_v | node | edge | cost | agg_cost
-----+-----+-----+-----+-----
 1 |  2 |  2 |  -1 |  0 |         0
 2 |  2 |  5 |  4 |  1 |         1
 3 |  2 |  6 |  8 |  1 |         2
 4 |  2 |  9 |  9 |  1 |         3
 5 |  2 | 10 | 10 |  1 |         2
 6 |  2 | 11 | 11 |  1 |         3
 7 | 13 | 13 | -1 |  0 |         0

```

```
(7 rows)
```

### Examples for queries marked as undirected with cost column

The examples in this section use the following *Graph 4: Undirected, with cost*

```
SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost FROM edge_table',
  2, 3, false
);
```

seq	node	edge	cost	agg_cost
1	1	1	1	1
2	2	-1	0	0
3	3	5	1	3
4	5	4	1	1
5	6	8	1	2
6	11	12	1	3
7	10	10	1	2
8	13	14	1	3
9	9	9	1	3
10	7	6	1	3
11	8	7	1	2

(11 rows)

```
SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost FROM edge_table',
  13, 3, false
);
```

seq	node	edge	cost	agg_cost
1	2	4	1	3
2	5	10	1	2
3	6	11	1	3
4	11	12	1	2
5	10	14	1	1
6	12	13	1	3
7	13	-1	0	0
8	8	7	1	3

(8 rows)

```
SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost FROM edge_table',
  array[2,13], 3, false
);
```

seq	from_v	node	edge	cost	agg_cost
1	2	1	1	1	1
2	2	2	-1	0	0
3	2	3	5	1	3
4	2	5	4	1	1
5	2	6	8	1	2
6	2	11	12	1	3
7	2	10	10	1	2
8	2	13	14	1	3
9	2	9	9	1	3
10	2	7	6	1	3
11	2	8	7	1	2
12	13	2	4	1	3
13	13	5	10	1	2
14	13	6	11	1	3

```

15 |      13 |      11 |      12 |      1 |      2
16 |      13 |      10 |      14 |      1 |      1
17 |      13 |      12 |      13 |      1 |      3
18 |      13 |      13 |      -1 |      0 |      0
19 |      13 |       8 |       7 |      1 |      3
(19 rows)

SELECT * FROM pgr_drivingDistance(
  'SELECT id, source, target, cost FROM edge_table',
  array[2,13], 3, false, equicost:=true
);
 seq | from_v | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----
  1 |      2 |      1 |      1 |      1 |          1
  2 |      2 |      2 |     -1 |      0 |          0
  3 |      2 |      3 |      5 |      1 |          3
  4 |      2 |      5 |      4 |      1 |          1
  5 |      2 |      6 |      8 |      1 |          2
  6 |      2 |      7 |      6 |      1 |          3
  7 |      2 |      8 |      7 |      1 |          2
  8 |      2 |      9 |      9 |      1 |          3
  9 |      2 |     10 |     10 |      1 |          2
 10 |      2 |     11 |     12 |      1 |          3
 11 |     13 |     13 |     -1 |      0 |          0
 12 |     13 |     12 |     13 |      1 |          3
(12 rows)

```

The queries use the *Sample Data* network.

### History

- Renamed in version 2.0.0
- Added functionality in version 2.1

### See Also

- *pgr\_alphaShape* - Alpha shape computation
- *pgr\_pointsAsPolygon* - Polygon around set of points

### Indices and tables

- genindex
- search

## pgr\_alphaShape

### Name

pgr\_alphaShape — Core function for alpha shape computation.

### Synopsis

Returns a table with (x, y) rows that describe the vertices of an alpha shape.

```
table() pgr_alphaShape(text sql [, float8 alpha]);
```

### Description

**sql** text a SQL query, which should return a set of rows with the following columns:

```
SELECT id, x, y FROM vertex_table
```

**id** int4 identifier of the vertex

**x** float8 x-coordinate

**y** float8 y-coordinate

**alpha** (optional) float8 alpha value. If specified alpha value equals 0 (default), then optimal alpha value is used. For more information, see [CGAL - 2D Alpha Shapes<sup>3</sup>](#).

Returns a vertex record for each row:

**x** x-coordinate

**y** y-coordinate

If a result includes multiple outer/inner rings, return those with separator row (x=NULL and y=NULL).

### History

- Renamed in version 2.0.0
- Added alpha argument with default 0 (use optimal value) in version 2.1.0
- Supported to return multiple outer/inner ring coordinates with separator row (x=NULL and y=NULL) in version 2.1.0

### Examples

In the alpha shape code we have no way to control the order of the points so the actual output you might get could be similar but different. The simple query is followed by a more complex one that constructs a polygon and computes the areas of it. This should be the same as the result on your system. We leave the details of the complex query to the reader as an exercise if they wish to decompose it into understandable pieces or to just copy and paste it into a SQL window to run.

```
SELECT * FROM pgr_alphaShape('SELECT id, x, y FROM vertex_table');
```

```
x | y
---+---
2 | 4
0 | 2
2 | 0
4 | 1
4 | 2
4 | 3
(6 rows)
```

```
SELECT round(ST_Area(ST_MakePolygon(ST_AddPoint(foo.openline, ST_StartPoint(foo.openline))))::num
FROM (SELECT ST_MakeLine(points ORDER BY id) AS openline FROM
(SELECT ST_MakePoint(x, y) AS points, row_number() over() AS id
FROM pgr_alphaShape('SELECT id, x, y FROM vertex_table')
) AS a) AS foo;
```

<sup>3</sup>[http://doc.cgal.org/latest/Alpha\\_shapes\\_2/group\\_\\_PkgAlphaShape2.html](http://doc.cgal.org/latest/Alpha_shapes_2/group__PkgAlphaShape2.html)

```

st_area
-----
 10.00
(1 row)

SELECT * FROM pgr_alphaShape('SELECT id::integer, ST_X(the_geom)::float AS x, ST_Y(the_geom)::flo
  x | y
-----+-----
  2 |  4
0.5 | 3.5
  0 |  2
  2 |  0
  4 |  1
  4 |  2
  4 |  3
 3.5 |  4
(8 rows)

SELECT round(ST_Area(ST_MakePolygon(ST_AddPoint(foo.openline, ST_StartPoint(foo.openline))))::num
FROM (SELECT ST_MakeLine(points ORDER BY id) AS openline FROM
(SELECT ST_MakePoint(x, y) AS points, row_number() over() AS id
FROM pgr_alphaShape('SELECT id::integer, ST_X(the_geom)::float AS x, ST_Y(the_geom)::float AS y F
) AS a) AS foo;

st_area
-----
 11.75
(1 row)

```

The queries use the *Sample Data* network.

### See Also

- *Driving Distance* - Driving Distance
- *pgr\_pointsAsPolygon* - Polygon around set of points

## pgr\_pointsAsPolygon

### Name

`pgr_pointsAsPolygon` — Draws an alpha shape around given set of points.

### Synopsis

Returns the alpha shape as (multi)polygon geometry.

```
geometry pgr_pointsAsPolygon(text sql [, float8 alpha]);
```

### Description

`sql` text a SQL query, which should return a set of rows with the following columns:

```
SELECT id, x, y FROM vertex_result;
```

`id` int4 identifier of the vertex

**x** float8 x-coordinate

**y** float8 y-coordinate

**alpha** (optional) float8 alpha value. If specified alpha value equals 0 (default), then optimal alpha value is used. For more information, see [CGAL - 2D Alpha Shapes<sup>4</sup>](#).

Returns a (multi)polygon geometry (with holes).

### History

- Renamed in version 2.0.0
- Added alpha argument with default 0 (use optimal value) in version 2.1.0
- Supported to return a (multi)polygon geometry (with holes) in version 2.1.0

### Examples

In the following query there is no way to control which point in the polygon is the first in the list, so you may get similar but different results than the following which are also correct. Each of the `pgr_pointsAsPolygon` queries below is followed by one the compute the area of the polygon. This area should remain constant regardless of the order of the points making up the polygon.

```

SELECT ST_AsText (pgr_pointsAsPolygon('SELECT id, x, y FROM vertex_table'));

          st_astext
-----
POLYGON((2 4,4 3,4 2,4 1,2 0,0 2,2 4))
(1 row)

SELECT round(ST_Area(pgr_pointsAsPolygon('SELECT id, x, y FROM vertex_table'))::numeric, 2) AS st_area;

          st_area
-----
          10.00
(1 row)

SELECT ST_AsText (pgr_pointsAsPolygon('SELECT id::integer, ST_X(the_geom)::float AS x, ST_Y(the_geom)::float AS y
                                       FROM edge_table_vertices_pgr'));

          st_astext
-----
POLYGON((2 4,3.5 4,4 3,4 2,4 1,2 0,0 2,0.5 3.5,2 4))
(1 row)

SELECT round(ST_Area(pgr_pointsAsPolygon('SELECT id::integer, ST_X(the_geom)::float AS x, ST_Y(the_geom)::float AS y
                                       FROM edge_table_vertices_pgr'))::numeric, 2) AS st_area;

          st_area
-----
          11.75

```

The queries use the *Sample Data* network.

### See Also

- *Driving Distance* - Driving Distance
- *pgr\_alphaShape* - Alpha shape computation

<sup>4</sup>[http://doc.cgal.org/latest/Alpha\\_shapes\\_2/group\\_\\_PkgAlphaShape2.html](http://doc.cgal.org/latest/Alpha_shapes_2/group__PkgAlphaShape2.html)

## 6.1.10 pgr\_kDijkstra - Multiple destination Shortest Path Dijkstra

### Name

- `pgr_kdijkstraCost` - Returns the costs for K shortest paths using Dijkstra algorithm.
- `pgr_kdijkstraPath` - Returns the paths for K shortest paths using Dijkstra algorithm.

### Synopsis

These functions allow you to have a single start node and multiple destination nodes and will compute the routes to all the destinations from the source node. Returns a set of `pgr_costResult` or `pgr_costResult3`. `pgr_kdijkstraCost` returns one record for each destination node and the cost is the total cost of the route to that node. `pgr_kdijkstraPath` returns one record for every edge in that path from source to destination and the cost is to traverse that edge.

```
pgr_costResult[] pgr_kdijkstraCost(text sql, integer source,
                                integer[] targets, boolean directed, boolean has_rcost);

pgr_costResult3[] pgr_kdijkstraPath(text sql, integer source,
                                   integer[] targets, boolean directed, boolean has_rcost);
```

### Description

`sql` a SQL query, which should return a set of rows with the following columns:

```
SELECT id, source, target, cost [,reverse_cost] FROM edge_table
```

**id** int4 identifier of the edge

**source** int4 identifier of the source vertex

**target** int4 identifier of the target vertex

**cost** float8 value, of the edge traversal cost. A negative cost will prevent the edge from being inserted in the graph.

**reverse\_cost** (optional) the cost for the reverse traversal of the edge. This is only used when the `directed` and `has_rcost` parameters are `true` (see the above remark about negative costs).

**source** int4 id of the start point

**targets** int4[] an array of ids of the end points

**directed** true if the graph is directed

**has\_rcost** if `true`, the `reverse_cost` column of the SQL generated set of rows will be used for the cost of the traversal of the edge in the opposite direction.

`pgr_kdijkstraCost` returns set of `pgr_costResult[]`:

**seq** row sequence

**id1** path vertex source id (this will always be source start point in the query).

**id2** path vertex target id

**cost** cost to traverse the path from `id1` to `id2`. Cost will be -1.0 if there is no path to that target vertex id.

`pgr_kdijkstraPath` returns set of `pgr_costResult3[]` - *Multiple Path Results with Cost*:

**seq** row sequence

**id1** path target id (identifies the target path).



**id2** path edge source node id

**id3** path edge id (-1 for the last row)

**cost** cost to traverse this edge or -1.0 if there is no path to this target

## History

- New in version 2.0.0

## Examples

- Returning a cost result

```
SELECT seq, id1 AS source, id2 AS target, cost FROM pgr_kdijkstraCost(
  'SELECT id, source, target, cost FROM edge_table',
  10, array[4,12], false, false
);
```

seq	source	target	cost
0	10	4	4
1	10	12	2

```
SELECT seq, id1 AS path, id2 AS node, id3 AS edge, cost
FROM pgr_kdijkstraPath(
  'SELECT id, source, target, cost FROM edge_table',
  10, array[4,12], false, false
);
```

seq	path	node	edge	cost
0	4	10	12	1
1	4	11	13	1
2	4	12	15	1
3	4	9	16	1
4	4	4	-1	0
5	12	10	12	1
6	12	11	13	1
7	12	12	-1	0

(8 rows)

- Returning a path result

```
SELECT id1 as path, st_astext(st_linemerge(st_union(b.the_geom))) as the_geom
FROM pgr_kdijkstraPath(
  'SELECT id, source, target, cost FROM edge_table',
  10, array[4,12], false, false
) a,
edge_table b
WHERE a.id3=b.id
GROUP by id1
ORDER by id1;
```

path	the_geom
4	LINestring(2 3,3 3,4 3,4 2,4 1)
12	LINestring(2 3,3 3,4 3)

(2 rows)

There is no assurance that the result above will be ordered in the direction of flow of the route, ie: it might be reversed. You will need to check if `st_startPoint()` of the route is the same as the start node location and

if it is not then call `st_reverse()` to reverse the direction of the route. This behavior is a function of PostGIS functions `st_linemerge()` and `st_union()` and not pgRouting.

### See Also

- [pgr\\_costResult\[\]](#)
- [http://en.wikipedia.org/wiki/Dijkstra%27s\\_algorithm](http://en.wikipedia.org/wiki/Dijkstra%27s_algorithm)

## 6.1.11 pgr\_ksp - Multiple shortest Paths based on Dijkstra

### Version 2.0 (deprecated)

- [pgr\\_ksp](#) - Multiple shortest Paths based on Dijkstra

### Version 2.1

- [pgr\\_ksp](#) - Multiple shortest Paths based on Dijkstra

### pgr\_ksp (V 2.0) - K-Shortest Path

**Name** `pgr_ksp` — Returns the “K” shortest paths.

**Synopsis** The K shortest path routing algorithm based on Yen’s algorithm. “K” is the number of shortest paths desired. Returns a set of [pgr\\_costResult3](#) (seq, id1, id2, id3, cost) rows, that make up a path.

```
pgr_costResult3[] pgr_ksp(sql text, source integer, target integer,
                        paths integer, has_rcost boolean);
```

**Warning:** This signature is being deprecated in version 2.1, Please use it without the `has_rcost` flag instead.

- for undirected graph. `pgr_ksp(sql, source, target, distance, directed:=false)`
- for directed graph. `pgr_ksp(sql, source, target, distance, directed:=true)`

See [pgr\\_ksp - K-Shortest Path](#)

### Description

**sql** a SQL query, which should return a set of rows with the following columns:

```
SELECT id, source, target, cost, [,reverse_cost] FROM edge_table
```

- id** int4 identifier of the edge
- source** int4 identifier of the source vertex
- target** int4 identifier of the target vertex
- cost** float8 value, of the edge traversal cost. A negative cost will prevent the edge from being inserted in the graph.
- reverse\_cost** (optional) the cost for the reverse traversal of the edge. This is only used when `has_rcost` the parameter is `true` (see the above remark about negative costs).
- source** int4 id of the start point
- target** int4 id of the end point

**paths** int4 number of alternative routes

**has\_rcost** if true, the `reverse_cost` column of the SQL generated set of rows will be used for the cost of the traversal of the edge in the opposite direction.

Returns set of `pgr_costResult[]`:

**seq** sequence for ording the results

**id1** route ID

**id2** node ID

**id3** edge ID (0 for the last row)

**cost** cost to traverse from id2 using id3

KSP code base taken from <http://code.google.com/p/k-shortest-paths/source>.

## History

- New in version 2.0.0

## Examples

- Without `reverse_cost`

```
SELECT seq, id1 AS route, id2 AS node, id3 AS edge, cost
FROM pgr_ksp(
  'SELECT id, source, target, cost FROM edge_table',
  7, 12, 2, false
);
```

seq	route	node	edge	cost
0	0	7	6	1
1	0	8	7	1
2	0	5	8	1
3	0	6	9	1
4	0	9	15	1
5	0	12	-1	0
6	1	7	6	1
7	1	8	7	1
8	1	5	8	1
9	1	6	11	1
10	1	11	13	1
11	1	12	-1	0

(12 rows)

- With `reverse_cost`

```
SELECT seq, id1 AS route, id2 AS node, id3 AS edge, cost
FROM pgr_ksp(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  7, 12, 2, true
);
```

seq	route	node	edge	cost
0	0	7	6	1
1	0	8	7	1
2	0	5	8	1
3	0	6	9	1
4	0	9	15	1
5	0	12	-1	0
6	1	7	6	1
7	1	8	7	1

8		1		5		8		1
9		1		6		11		1
10		1		11		13		1
11		1		12		-1		0
(12 rows)								

The queries use the *Sample Data* network.

### See Also

- [pgr\\_costResult3\[\] - Multiple Path Results with Cost](#)
- [http://en.wikipedia.org/wiki/K\\_shortest\\_path\\_routing](http://en.wikipedia.org/wiki/K_shortest_path_routing)

### pgr\_ksp - K-Shortest Path

**Name** pgr\_ksp — Returns the “K” shortest paths.



Fig. 6.3: Boost Graph Inside

**Synopsis** The K shortest path routing algorithm based on Yen’s algorithm. “K” is the number of shortest paths desired.

#### The minimal signature:

```
pgr_ksp(TEXT sql_q, BIGINT start_vid, BIGINT end_vid, INTEGER k);
RETURNS SET OF (seq, path_id, path_seq, node, edge, cost, agg_cost) or EMPTY SET
```

#### The full signature:

```
pgr_ksp(TEXT edges_sql, BIGINT start_vid, BIGINT end_vid, INTEGER k,
        BOOLEAN directed:=true, BOOLEAN heap_paths:=false);
RETURNS SET OF (seq, path_id, path_seq, node, edge, cost, agg_cost) or EMPTY SET
```

**Description of the SQL query** General description of the edges\_sql

```
SELECT id, source, target, cost [,reverse_cost] FROM ...
```

**sql\_q** a SQL query, which returns a set of rows with the following columns:

- id** ANY-INTEGERS identifier of the edge.
- source** ANY-INTEGERS identifier of the source vertex of the edge.
- target** ANY-INTEGERS identifier of the target vertex of the edge.
- cost** ANY-NUMERICAL value of the edge traversal cost. A negative cost will prevent the edge (source, target) from being inserted in the graph.

**reverse\_cost** ANY-NUMERICAL (optional) the value for the reverse traversal of the edge. A negative cost will prevent the edge (target, source) from being inserted in the graph.

Where:

**ANY-INTEGER** smallint, int, bigint

**ANY-NUMERICAL** smallint, int, bigint, real, float

**Description of the parameters of the signatures**

**sql\_q** TEXT SQL query as described above.

**start\_vid** BIGINT id of the starting vertex.

**end\_vid** BIGINT id of the ending vertex.

**k** INTEGER The desired number of paths.

**directed** BOOLEAN (optional). When false the graph is considered as Undirected. Default is true which considers the graph as Directed.

**heap\_paths** BOOLEAN (optional). When true returns all the paths stored in the process heap. Default is false which only returns k paths.

Roughly, if the shortest path has N edges, the heap will contain about than  $N * k$  paths for small value of k and  $k > 1$ .

**Description of the return values** Returns set of (seq, path\_seq, path\_id, node, edge, cost, agg\_cost)

**seq** INT sequential number starting from 1.

**path\_seq** INT relative position in the path of node and edge. Has value 1 for the beginning of a path.

**path\_id** BIGINT path identifier. The ordering of the paths For two paths i, j if  $i < j$  then  $agg\_cost(i) \leq agg\_cost(j)$ .

**node** BIGINT id of the node in the path.

**edge** BIGINT id of the edge used to go from node to the next node in the path sequence. -1 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 total cost from start\_vid to node.

**Warning:** During the transition to 3.0, because pgr\_ksp version 2.0 doesn't have defined a directed flag nor a heap\_path flag, when pgr\_ksp is used with only one flag version 2.0 will be used.

**Examples to handle the one flag to choose signatures** The examples in this section use the following *Graph 1: Directed, with cost and reverse cost*

```
SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, 12, 2,
  true -- takes the (V2.0) signature (has_rcost = true and works on directed graph)
);
seq | id1 | id2 | id3 | cost
-----+-----+-----+-----+-----
  0 |   0 |   2 |   4 |   1
  1 |   0 |   5 |   8 |   1
  2 |   0 |   6 |   9 |   1
```

```

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

SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, 12, 2,
  directed:=true -- takes the new signature
);

SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, 12, 2
);
seq | path_id | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 0 | 1 | 2 | 4 | 1 | 0
2 | 0 | 2 | 5 | 8 | 1 | 1
3 | 0 | 3 | 6 | 9 | 1 | 2
4 | 0 | 4 | 9 | 15 | 1 | 3
5 | 0 | 5 | 12 | -1 | 0 | 4
6 | 1 | 1 | 2 | 4 | 1 | 0
7 | 1 | 2 | 5 | 8 | 1 | 1
8 | 1 | 3 | 6 | 11 | 1 | 2
9 | 1 | 4 | 11 | 13 | 1 | 3
10 | 1 | 5 | 12 | -1 | 0 | 4
(10 rows)

```

**Examples for queries marked as directed with cost and reverse\_cost columns** The examples in this section use the following *Graph 1: Directed, with cost and reverse cost*

```

SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, 12, 2
);
seq | path_id | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 2 | 4 | 1 | 0
2 | 1 | 2 | 5 | 8 | 1 | 1
3 | 1 | 3 | 6 | 9 | 1 | 2
4 | 1 | 4 | 9 | 15 | 1 | 3
5 | 1 | 5 | 12 | -1 | 0 | 4
6 | 2 | 1 | 2 | 4 | 1 | 0
7 | 2 | 2 | 5 | 8 | 1 | 1
8 | 2 | 3 | 6 | 11 | 1 | 2
9 | 2 | 4 | 11 | 13 | 1 | 3
10 | 2 | 5 | 12 | -1 | 0 | 4
(10 rows)

SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, 12, 2, heap_paths:=true
);

```

```

SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, 12, 2, true, true
);

```

seq	path_id	path_seq	node	edge	cost	agg_cost
1	1	1	2	4	1	0
2	1	2	5	8	1	1
3	1	3	6	9	1	2
4	1	4	9	15	1	3
5	1	5	12	-1	0	4
6	2	1	2	4	1	0
7	2	2	5	8	1	1
8	2	3	6	11	1	2
9	2	4	11	13	1	3
10	2	5	12	-1	0	4
11	3	1	2	4	1	0
12	3	2	5	10	1	1
13	3	3	10	12	1	2
14	3	4	11	13	1	3
15	3	5	12	-1	0	4

(15 rows)

**Examples for queries marked as undirected with cost and reverse\_cost columns** The examples in this section use the following *Graph 2: Undirected, with cost and reverse cost*

```

SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, 12, 2, directed:=false
);

```

seq	path_id	path_seq	node	edge	cost	agg_cost
1	1	1	2	2	1	0
2	1	2	3	3	1	1
3	1	3	4	16	1	2
4	1	4	9	15	1	3
5	1	5	12	-1	0	4
6	2	1	2	4	1	0
7	2	2	5	8	1	1
8	2	3	6	9	1	2
9	2	4	9	15	1	3
10	2	5	12	-1	0	4

(10 rows)

```

SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, 12, 2, directed:=false, heap_paths:=true
);

```

```

SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost, reverse_cost FROM edge_table',
  2, 12, 2, false, true
);

```

seq	path_id	path_seq	node	edge	cost	agg_cost
1	1	1	2	2	1	0
2	1	2	3	3	1	1
3	1	3	4	16	1	2
4	1	4	9	15	1	3
5	1	5	12	-1	0	4

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

(22 rows)

**Examples for queries marked as directed with cost column** The examples in this section use the following *Graph 3: Directed, with cost*

Empty path representation

```
SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost FROM edge_table',
  2, 3, 2
);
seq | path_id | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
(0 rows)
```

```
SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost FROM edge_table',
  2, 12, 2
);
seq | path_id | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
1 | 1 | 1 | 2 | 4 | 1 | 0
2 | 1 | 2 | 5 | 8 | 1 | 1
3 | 1 | 3 | 6 | 9 | 1 | 2
4 | 1 | 4 | 9 | 15 | 1 | 3
5 | 1 | 5 | 12 | -1 | 0 | 4
6 | 2 | 1 | 2 | 4 | 1 | 0
7 | 2 | 2 | 5 | 8 | 1 | 1
8 | 2 | 3 | 6 | 11 | 1 | 2
9 | 2 | 4 | 11 | 13 | 1 | 3
10 | 2 | 5 | 12 | -1 | 0 | 4
(10 rows)
```

```
SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost FROM edge_table',
  2, 12, 2, heap_paths:=true
);

SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost FROM edge_table',
  2, 12, 2, true, true
);
seq | path_id | path_seq | node | edge | cost | agg_cost
```



```

-----+-----+-----+-----+-----+-----+-----
  1 |         1 |         1 |         2 |         4 |         1 |         0
  2 |         1 |         2 |         5 |         8 |         1 |         1
  3 |         1 |         3 |         6 |         9 |         1 |         2
  4 |         1 |         4 |         9 |        15 |         1 |         3
  5 |         1 |         5 |        12 |        -1 |         0 |         4
  6 |         2 |         1 |         2 |         4 |         1 |         0
  7 |         2 |         2 |         5 |         8 |         1 |         1
  8 |         2 |         3 |         6 |        11 |         1 |         2
  9 |         2 |         4 |        11 |        13 |         1 |         3
 10 |         2 |         5 |        12 |        -1 |         0 |         4
 11 |         3 |         1 |         2 |         4 |         1 |         0
 12 |         3 |         2 |         5 |        10 |         1 |         1
 13 |         3 |         3 |        10 |        12 |         1 |         2
 14 |         3 |         4 |        11 |        13 |         1 |         3
 15 |         3 |         5 |        12 |        -1 |         0 |         4
(15 rows)

```

**Examples for queries marked as undirected with cost column** The examples in this section use the following *Graph 4: Undirected, with cost*

```

SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost FROM edge_table',
  2, 12, 2, directed:=false
);
 seq | path_id | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
  1 |         1 |         1 |         2 |         4 |         1 |         0
  2 |         1 |         2 |         5 |         8 |         1 |         1
  3 |         1 |         3 |         6 |         9 |         1 |         2
  4 |         1 |         4 |         9 |        15 |         1 |         3
  5 |         1 |         5 |        12 |        -1 |         0 |         4
  6 |         2 |         1 |         2 |         4 |         1 |         0
  7 |         2 |         2 |         5 |         8 |         1 |         1
  8 |         2 |         3 |         6 |        11 |         1 |         2
  9 |         2 |         4 |        11 |        13 |         1 |         3
 10 |         2 |         5 |        12 |        -1 |         0 |         4
(10 rows)

SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost FROM edge_table',
  2, 12, 2, directed:=false, heap_paths:=true
);

SELECT * FROM pgr_ksp(
  'SELECT id, source, target, cost FROM edge_table',
  2, 12, 2, false, true
);
 seq | path_id | path_seq | node | edge | cost | agg_cost
-----+-----+-----+-----+-----+-----+-----
  1 |         1 |         1 |         2 |         4 |         1 |         0
  2 |         1 |         2 |         5 |         8 |         1 |         1
  3 |         1 |         3 |         6 |         9 |         1 |         2
  4 |         1 |         4 |         9 |        15 |         1 |         3
  5 |         1 |         5 |        12 |        -1 |         0 |         4
  6 |         2 |         1 |         2 |         4 |         1 |         0
  7 |         2 |         2 |         5 |         8 |         1 |         1
  8 |         2 |         3 |         6 |        11 |         1 |         2
  9 |         2 |         4 |        11 |        13 |         1 |         3
 10 |         2 |         5 |        12 |        -1 |         0 |         4
 11 |         3 |         1 |         2 |         4 |         1 |         0

```

12		3		2		5		10		1		1
13		3		3		10		12		1		2
14		3		4		11		13		1		3
15		3		5		12		-1		0		4
(15 rows)												

The queries use the *Sample Data* network.

## History

- New in version 2.0.0
- Added functionality version 2.1

## See Also

- [http://en.wikipedia.org/wiki/K\\_shortest\\_path\\_routing](http://en.wikipedia.org/wiki/K_shortest_path_routing)

## Indices and tables

- [genindex](#)
- [search](#)

## 6.1.12 pgr\_tsp - Traveling Sales Person

### Name

- `pgr_tsp` - Returns the best route from a start node via a list of nodes.
- `pgr_tsp` - Returns the best route order when passed a distance matrix.
- `pgr_makeDistanceMatrix` - Returns a Euclidean distance Matrix from the points provided in the sql result.

### Synopsis

The travelling salesman problem (TSP) or travelling salesperson problem asks the following question: Given a list of cities and the distances between each pair of cities, what is the shortest possible route that visits each city exactly once and returns to the origin city? This algorithm uses simulated annealing to return a high quality approximate solution. Returns a set of *pgr\_costResult* (seq, id1, id2, cost) rows, that make up a path.

```
pgr_costResult[] pgr_tsp(sql text, start_id integer);
pgr_costResult[] pgr_tsp(sql text, start_id integer, end_id integer);
```

Returns a set of (seq integer, id1 integer, id2 integer, cost float8) that is the best order to visit the nodes in the matrix. `id1` is the index into the distance matrix. `id2` is the point id from the sql.

If no `end_id` is supplied or it is -1 or equal to the `start_id` then the TSP result is assumed to be a circular loop returning back to the start. If `end_id` is supplied then the route is assumed to start and end the the designated ids.

```
record[] pgr_tsp(matrix float[][], start integer)
record[] pgr_tsp(matrix float[][], start integer, end integer)
```

## Description

### With Euclidean distances

The TSP solver is based on ordering the points using straight line (euclidean) distance <sup>1</sup> between nodes. The implementation is using an approximation algorithm that is very fast. It is not an exact solution, but it is guaranteed that a solution is returned after certain number of iterations.

**sql** a SQL query, which should return a set of rows with the following columns:

```
SELECT id, x, y FROM vertex_table
```

**id** int4 identifier of the vertex

**x** float8 x-coordinate

**y** float8 y-coordinate

**start\_id** int4 id of the start point

**end\_id** int4 id of the end point, This is *OPTIONAL*, if include the route is optimized from start to end, otherwise it is assumed that the start and the end are the same point.

The function returns set of *pgr\_costResult[]*:

**seq** row sequence

**id1** internal index to the distance matrix

**id2** id of the node

**cost** cost to traverse from the current node to the next node.

### Create a distance matrix

For users that need a distance matrix we have a simple function that takes SQL in *sql* as described above and returns a record with *dmatrix* and *ids*.

```
SELECT dmatrix, ids from pgr_makeDistanceMatrix('SELECT id, x, y FROM vertex_table');
```

The function returns a record of *dmatrix*, *ids*:

**dmatrix** float8[][] a symetric Euclidean distance matrix based on *sql*.

**ids** integer[] an array of ids as they are ordered in the distance matrix.

### With distance matrix

For users, that do not want to use Euclidean distances, we also provide the ability to pass a distance matrix that we will solve and return an ordered list of nodes for the best order to visit each. It is up to the user to fully populate the distance matrix.

**matrix** float[][] distance matrix of points

**start** int4 index of the start point

**end** int4 (optional) index of the end node

<sup>1</sup> There was some thought given to pre-calculating the driving distances between the nodes using Dijkstra, but then I read a paper (unfortunately I don't remember who wrote it), where it was proved that the quality of TSP with euclidean distance is only slightly worse than one with real distance in case of normal city layout. In case of very sparse network or rivers and bridges it becomes more inaccurate, but still wholly satisfactory. Of course it is nice to have exact solution, but this is a compromise between quality and speed (and development time also). If you need a more accurate solution, you can generate a distance matrix and use that form of the function to get your results.

The end node is an optional parameter, you can just leave it out if you want a loop where the start is the depot and the route returns back to the depot. If you include the end parameter, we optimize the path from start to end and minimize the distance of the route while include the remaining points.

The distance matrix is a multidimensional PostgreSQL array type<sup>6</sup> that must be N × N in size.

The result will be N records of [ seq, id ]:

**seq** row sequence

**id** index into the matrix

### History

- Renamed in version 2.0.0
- GAUL dependency removed in version 2.0.0

### Examples

- Using SQL parameter (all points from the table, starting from 6 and ending at 5). We have listed two queries in this example, the first might vary from system to system because there are multiple equivalent answers. The second query should be stable in that the length optimal route should be the same regardless of order.

```

SELECT seq, id1, id2, round(cost::numeric, 2) AS cost
FROM pgr_tsp('SELECT id, x, y FROM vertex_table ORDER BY id', 6, 5);

seq | id1 | id2 | cost
-----+-----+-----+-----
  0 |   5 |   6 | 1.00
  1 |   6 |   7 | 1.00
  2 |   7 |   8 | 1.41
  3 |   1 |   2 | 1.00
  4 |   0 |   1 | 1.41
  5 |   2 |   3 | 1.00
  6 |   3 |   4 | 1.00
  7 |   8 |   9 | 1.00
  8 |  11 |  12 | 1.00
  9 |  10 |  11 | 1.41
 10 |  12 |  13 | 1.00
 11 |   9 |  10 | 2.24
 12 |   4 |   5 | 1.00
(13 rows)

SELECT round(sum(cost)::numeric, 4) as cost
FROM pgr_tsp('SELECT id, x, y FROM vertex_table ORDER BY id', 6, 5);

cost
-----
15.4787
(1 row)

```

- Using distance matrix (A loop starting from 1)

When using just the start node you are getting a loop that starts with 1, in this case, and travels through the other nodes and is implied to return to the start node from the last one in the list. Since this is a circle there are at least two possible paths, one clockwise and one counter-clockwise that will have the same length and be equally valid. So in the following example it is also possible to get back a sequence of ids = {1,0,3,2} instead of the {1,2,3,0} sequence listed below.

<sup>6</sup><http://www.postgresql.org/docs/9.1/static/arrays.html>

```
SELECT seq, id FROM pgr_tsp('{{0,1,2,3},{1,0,4,5},{2,4,0,6},{3,5,6,0}}'::float8[],1);
```

seq	id
0	1
1	2
2	3
3	0

(4 rows)

- Using distance matrix (Starting from 1, ending at 2)

```
SELECT seq, id FROM pgr_tsp('{{0,1,2,3},{1,0,4,5},{2,4,0,6},{3,5,6,0}}'::float8[],1,2);
```

seq	id
0	1
1	0
2	3
3	2

(4 rows)

- Using the vertices table `edge_table_vertices_pgr` generated by `pgr_createTopology`. Again we have two queries where the first might vary and the second is based on the overall path length.

```
SELECT seq, id1, id2, round(cost::numeric, 2) AS cost
FROM pgr_tsp('SELECT id::integer, st_x(the_geom) as x,st_x(the_geom) as y FROM edge_table_vertices_pgr');
```

seq	id1	id2	cost
0	5	6	0.00
1	10	11	0.00
2	2	3	1.41
3	3	4	0.00
4	11	12	0.00
5	8	9	0.71
6	15	16	0.00
7	16	17	2.12
8	1	2	0.00
9	14	15	1.41
10	7	8	1.41
11	6	7	0.71
12	13	14	2.12
13	0	1	0.00
14	9	10	0.00
15	12	13	0.00
16	4	5	1.41

(17 rows)

```
SELECT round(sum(cost)::numeric, 4) as cost
FROM pgr_tsp('SELECT id::integer, st_x(the_geom) as x,st_x(the_geom) as y FROM edge_table_vertices_pgr');
```

cost
11.3137

(1 row)

The queries use the *Sample Data* network.

### See Also

- `pgr_costResult[]`

- [http://en.wikipedia.org/wiki/Traveling\\_salesman\\_problem](http://en.wikipedia.org/wiki/Traveling_salesman_problem)
- [http://en.wikipedia.org/wiki/Simulated\\_annealing](http://en.wikipedia.org/wiki/Simulated_annealing)

### 6.1.13 pgr\_trsp - Turn Restriction Shortest Path (TRSP)

#### Name

pgr\_trsp — Returns the shortest path with support for turn restrictions.

#### Synopsis

The turn restricted shortest path (TRSP) is a shortest path algorithm that can optionally take into account complicated turn restrictions like those found in real work navigable road networks. Performance wise it is nearly as fast as the A\* search but has many additional features like it works with edges rather than the nodes of the network. Returns a set of *pgr\_costResult* (seq, id1, id2, cost) rows, that make up a path.

```
pgr_costResult[] pgr_trsp(sql text, source integer, target integer,
    directed boolean, has_rcost boolean [,restrict_sql text]);
```

```
pgr_costResult[] pgr_trsp(sql text, source_edge integer, source_pos float8,
    target_edge integer, target_pos float8,
    directed boolean, has_rcost boolean [,restrict_sql text]);
```

```
pgr_costResult3[] pgr_trspViaVertices(sql text, vids integer[],
    directed boolean, has_reverse_cost boolean
    [, turn_restrict_sql text]);
```

```
pgr_costResult3[] pgr_trspViaEdges(sql text, eids integer[], pcts float8[],
    directed boolean, has_reverse_cost boolean
    [, turn_restrict_sql text]);
```

#### Description

The Turn Restricted Shortest Path algorithm (TRSP) is similar to the *Shooting Star algorithm* in that you can specify turn restrictions.

The TRSP setup is mostly the same as *Dijkstra shortest path* with the addition of an optional turn restriction table. This provides an easy way of adding turn restrictions to a road network by placing them in a separate table.

**sql** a SQL query, which should return a set of rows with the following columns:

```
SELECT id, source, target, cost, [,reverse_cost] FROM edge_table
```

**id** int4 identifier of the edge

**source** int4 identifier of the source vertex

**target** int4 identifier of the target vertex

**cost** float8 value, of the edge traversal cost. A negative cost will prevent the edge from being inserted in the graph.

**reverse\_cost** (optional) the cost for the reverse traversal of the edge. This is only used when the **directed** and **has\_rcost** parameters are **true** (see the above remark about negative costs).

**source** int4 **NODE id** of the start point

**target** int4 **NODE id** of the end point

**directed** **true** if the graph is directed

**has\_rcost** if `true`, the `reverse_cost` column of the SQL generated set of rows will be used for the cost of the traversal of the edge in the opposite direction.

**restrict\_sql** (optional) a SQL query, which should return a set of rows with the following columns:

```
SELECT to_cost, target_id, via_path FROM restrictions
```

**to\_cost** float8 turn restriction cost

**target\_id** int4 target id

**via\_path** text comma separated list of edges in the reverse order of rule

Another variant of TRSP allows to specify **EDGE id** of source and target together with a fraction to interpolate the position:

**source\_edge** int4 **EDGE id** of the start edge

**source\_pos** float8 fraction of 1 defines the position on the start edge

**target\_edge** int4 **EDGE id** of the end edge

**target\_pos** float8 fraction of 1 defines the position on the end edge

Returns set of *pgr\_costResult[]*:

**seq** row sequence

**id1** node ID

**id2** edge ID (-1 for the last row)

**cost** cost to traverse from `id1` using `id2`

## History

- New in version 2.0.0

## Support for Vias

**Warning:** The Support for Vias functions are prototypes. Not all corner cases are being considered.

We also have support for vias where you can say generate a from A to B to C, etc. We support both methods above only you pass an array of vertices or and array of edges and percentage position along the edge in two arrays.

**sql** a SQL query, which should return a set of rows with the following columns:

```
SELECT id, source, target, cost, [,reverse_cost] FROM edge_table
```

**id** int4 identifier of the edge

**source** int4 identifier of the source vertex

**target** int4 identifier of the target vertex

**cost** float8 value, of the edge traversal cost. A negative cost will prevent the edge from being inserted in the graph.

**reverse\_cost** (optional) the cost for the reverse traversal of the edge. This is only used when the `directed` and `has_rcost` parameters are `true` (see the above remark about negative costs).

**vids** int4[] An ordered array of **NODE id** the path will go through from start to end.

**directed** true if the graph is directed

**has\_rcost** if true, the `reverse_cost` column of the SQL generated set of rows will be used for the cost of the traversal of the edge in the opposite direction.

**restrict\_sql** (optional) a SQL query, which should return a set of rows with the following columns:

```
SELECT to_cost, target_id, via_path FROM restrictions
```

**to\_cost** float8 turn restriction cost

**target\_id** int4 target id

**via\_path** text comma separated list of edges in the reverse order of rule

Another variant of TRSP allows to specify **EDGE id** together with a fraction to interpolate the position:

**eids** int4 An ordered array of **EDGE id** that the path has to traverse

**pcts** float8 An array of fractional positions along the respective edges in `eids`, where 0.0 is the start of the edge and 1.0 is the end of the eadge.

Returns set of `pgr_costResult[]`:

**seq** row sequence

**id1** route ID

**id2** node ID

**id3** edge ID (-1 for the last row)

**cost** cost to traverse from `id2` using `id3`

## History

- Via Support prototypes new in version 2.1.0

## Examples

- Without turn restrictions

```
SELECT seq, id1 AS node, id2 AS edge, cost
FROM pgr_trsp(
    'SELECT id, source, target, cost FROM edge_table',
    7, 12, false, false
);
```

seq	node	edge	cost
0	7	6	1
1	8	7	1
2	5	8	1
3	6	11	1
4	11	13	1
5	12	-1	0

(6 rows)

- With turn restrictions

Turn restrictions require additional information, which can be stored in a separate table:

```
CREATE TABLE restrictions (
    rid serial,
    to_cost double precision,
    to_edge integer,
    from_edge integer,
    via text
```



```
);
INSERT INTO restrictions VALUES (1,100,7,4,null);
INSERT INTO restrictions VALUES (2,4,8,3,5);
INSERT INTO restrictions VALUES (3,100,9,16,null);
```

Then a query with turn restrictions is created as:

```
SELECT seq, id1 AS node, id2 AS edge, cost
FROM pgr_trsp(
  'SELECT id, source, target, cost FROM edge_table',
  7, 12, false, false,
  'SELECT to_cost, to_edge AS target_id,
  from_edge || coalesce('',' || via, '') AS via_path
FROM restrictions'
);
```

seq	node	edge	cost
0	7	6	1
1	8	7	1
2	5	8	1
3	6	11	1
4	11	13	1
5	12	-1	0

(6 rows)

An example query using vertex ids and via points:

```
SELECT * FROM pgr_trspViaVertices(
  'SELECT id, source::INTEGER, target::INTEGER, cost,
  reverse_cost FROM edge_table',
  ARRAY[1,8,13,5]::INTEGER[],
  true,
  true,
  'SELECT to_cost, to_edge AS target_id, FROM_edge ||
  coalesce('','||via, '') AS via_path FROM restrictions');
```

seq	id1	id2	id3	cost
1	1	1	1	1
2	1	2	4	1
3	1	5	8	1
4	1	6	9	1
5	1	9	16	1
6	1	4	3	1
7	1	3	5	1
8	1	6	8	1
9	1	5	7	1
10	2	8	7	1
11	2	5	10	1
12	2	10	14	1
13	3	13	14	1
14	3	10	10	1
15	3	5	-1	0

(15 rows)

An example query using edge ids and vias:

```
SELECT * FROM pgr_trspViaEdges(
  'SELECT id, source::INTEGER, target::INTEGER, cost,
  reverse_cost FROM edge_table',
  ARRAY[1,11,6]::INTEGER[],
```

```
ARRAY[0.5, 0.5, 0.5]::FLOAT8[],
true,
true,

'SELECT to_cost, to_edge AS target_id, FROM_edge ||
    coalesce('','||via,\'') AS via_path FROM restrictions');
```

seq	id1	id2	id3	cost
1	1	-1	1	0.5
2	1	2	4	1
3	1	5	8	1
4	1	6	11	1
5	2	11	13	1
6	2	12	15	1
7	2	9	9	1
8	2	6	8	1
9	2	5	7	1
10	2	8	6	0.5

(10 rows)

The queries use the *Sample Data* network.

### See Also

- `pgr_costResult[]`

---

## Pre processing or post processing helping functions

---

### *Driving Distance post-processing*

- *pgr\_alphaShape* - Alpha shape computation
- *pgr\_pointsAsPolygon* - Polygon around set of points



---

## Experimental and Proposed functions

---

**This section contains new experimental or proposed signatures for any of the following sections:**

- topology functions
- routing functions
- vehicle routing functions
- pre / post procesing helper functions

We are including them so that the pgRouting community can evaluate them before including them as an official function of pgRouting.

Some of them are unsupported like the GSoC functions.



---

## Experimental functions: Proposed by Steve Woodbridge

---

### Convenience Functions

- *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`.

## 9.1 Convenience Functions

The following functions are general purpose convenience functions that might be useful when building a larger application or handling input from say an Ajax handler.

- *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.

## 9.2 Distance Matrix Functions

These function may be helpful when you need to create or manipulate distance matrices, like for TSP or VRP related problems.

- *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`.

### 9.2.1 pgr\_pointToEdgeNode

#### Name

`pgr_pointToEdgeNode` - Converts a point to a `vertex_id` based on closest edge.

## Synopsis

The function returns:

- `integer` that is the vertex id of the closest edge in the `edges` table within the `tol` tolerance of `pnt`. The vertex is selected by projection the `pnt` onto the edge and selecting which vertex is closer along the edge.

```
integer pgr_pointToEdgeNode(edges text, pnt geometry, tol float8)
```

## Description

Given an table `edges` with a spatial index on `the_geom` and a point geometry search for the closest edge within `tol` distance to the edges then compute the projection of the point onto the line segment and select source or target based on whether the projected point is closer to the respective end and return the source or target value.

## Parameters

The function accepts the following parameters:

- `edges text` The name of the edge table or view. (may contain the schema name AS well).
- `pnt geometry` A point geometry object in the same SRID as `edges`.
- `tol float8` The maximum search distance for an edge.

**Warning:** If no edge is within `tol` distance then return -1

The `edges` table must have the following columns:

- `source`
- `target`
- `the_geom`

## History

- New in version 2.1.0

## Examples

```
select pgr_pointtoedgenode('edge_table', 'POINT(2 0)::geometry, 0.02);
pgr_pointtoedgenode
-----
                1
(1 row)
```

The example uses the *Sample Data* network.

## See Also

- *pgr\_pointsToVids* - convert an array of point geometries into vertex ids.



## 9.2.2 pgr\_flipEdges

### Name

pgr\_flipEdges -

### Synopsis

The function returns:

- `geometry[]` An array of the input geometries with the geometries flipped end to end such that the geometries are oriented as a path from start to end.

```
geometry[] pgr_flipEdges(ga geometry[])
```

### Description

Given an array of linestrings that are supposedly connected end to end like the results of a route, check the edges and flip any end for end if they do not connect with the previous segment and return the array with the segments flipped as appropriate.

### Parameters

`ga geometry[]` An array of geometries, like the results of a routing query.

#### Warning:

- No checking is done for edges that do not connect.
- Input geometries MUST be `LINestring` or `MULTILINESTRING`.
- Only the first `LINestring` of a `MULTILINESTRING` is considered.

### History

- New in version 2.1.0

### Examples

```
select st_astext(e) from (select unnest(pgr_flipedges(ARRAY[
'LINESTRING(2 1,2 2) '::geometry,
'LINESTRING(2 2,2 3) '::geometry,
'LINESTRING(2 2,2 3) '::geometry,
'LINESTRING(2 2,3 2) '::geometry,
'LINESTRING(3 2,4 2) '::geometry,
'LINESTRING(4 1,4 2) '::geometry,
'LINESTRING(3 1,4 1) '::geometry,
'LINESTRING(2 1,3 1) '::geometry,
'LINESTRING(2 0,2 1) '::geometry,
'LINESTRING(2 0,2 1) '::geometry]::geometry[])) as e) as foo;
 st_astext
-----
LINESTRING(2 1,2 2)
LINESTRING(2 2,2 3)
LINESTRING(2 3,2 2)
LINESTRING(2 2,3 2)
LINESTRING(3 2,4 2)
LINESTRING(4 2,4 1)
```

```
LINESTRING(4 1,3 1)
LINESTRING(3 1,2 1)
LINESTRING(2 1,2 0)
LINESTRING(2 0,2 1)
(10 rows)
```

## See Also

### 9.2.3 pgr\_textToPoints

#### Name

`pgr_textToPoints` - Converts a text string of the format “x,y;x,y;x,y;...” into an array of point geometries.

#### Synopsis

Given a text string of the format “x,y;x,y;x,y;...” and the srid to use, split the string and create an array of point geometries.

The function returns:

- 

```
integer pgr_textToPoints(pnts text, srid integer DEFAULT(4326))
```

#### Description

#### Parameters

**pnts** text A text string of the format “x,y;x,y;x,y;...” where x is longitude and y is latitude if use values in lat-lon.

**srid** integer The SRID to use when constructing the point geometry. If the parameter is absent it defaults to SRID:4326.

#### History

- New in version 2.1.0

#### Examples

```
select st_astext(g) from (
  select unnest(
    pgr_texttopoints('0,0;1,1;1,0;0,1;1,4;1,5;0,4;0,5', 0)
  ) as g
) as foo;
st_astext
-----
POINT(0 0)
POINT(1 1)
POINT(1 0)
POINT(0 1)
POINT(1 4)
POINT(1 5)
POINT(0 4)
```

```
POINT(0 5)
(8 rows)
```

## See Also

- `pgr_pointToEdgeNode` - convert a point geometry to a `node_id` based on closest edge.
- `pgr_pointsToVids` - convert an array of point geometries into vertex ids.

## 9.2.4 pgr\_pointsToVids

### Name

`pgr_pointsToVids` - Converts an array of point geometries into vertex ids.

### Synopsis

Given an array of point geometries and an edge table and a max search tol distance the function converts points into vertex ids using `pgr_pointtoedgenode()`.

The function returns:

- `integer[]` - An array of `vertex_id`.

```
integer[] pgr_pointsToVids(pnts geometry[], edges text, tol float8 DEFAULT(0.01))
```

### Description

#### Parameters

- pnts** `geometry[]` - An array of point geometries.
- edges** `text` - The edge table to be used for the conversion.
- tol** `float8` - The maximum search distance for locating the closest edge.

**Warning:** You need to check the results for any `vids=-1` which indicates if failed to locate an edge.

### History

- New in version 2.1.0

### Examples

```
select * from pgr_pointstovids(
  pgr_texttopoints('2,0;2,1;3,1;2,2;4,1;4,2;2,3;3,2', 0),
  'edge_table'
  -- tol is not specified, so we use the default of 0.01
);
 pgr_pointstovids
-----
 {1,2,3,5,4,9,10,6}
(1 row)
```

This example uses the *Sample Data* network.

## See Also

- *pgr\_pointToEdgeNode* - convert a point geometry to the closest vertex\_id of an edge..

## 9.2.5 pgr\_pointsToDMatrix

### Name

`pgr_pointsToDMatrix` - Creates a distance matrix from an array of points.

### Synopsis

Create a distance symmetric distance matrix suitable for TSP using Euclidean distances based on the `st_distance()`. You might want to create a variant of this the uses `st_distance_sphere()` or `st_distance_spheriod()` or some other function.

The function returns:

- **record** - with two fields as describe here
  - **dmatrix** float8[] - the distance matrix suitable to pass to `pgrTSP()` function.
  - **ids** integer[] - an array of ids for the distance matrix.

```
record pgr_pointsToDMatrix(pnts geometry[], OUT dmatrix double precision[], OUT ids integer[])
```

### Description

#### Parameters

**pnts** geometry[] - An array of point geometries.

**Warning:** The generated matrix will be symmetric as required for `pgr_TSP`.

### History

- New in version 2.1.0

### Examples

```
select unnest(dmatrix) from pgr_pointsToDMatrix(
  pgr_texttopoints('2,0;2,1;3,1;2,2;4,1;4,2;2,3;3,2', 0)
) limit 8;
  unnest
-----
         0
         1
  1.4142135623731
         2
  2.23606797749979
  2.82842712474619
         3
  2.23606797749979
(8 rows)

select ids from pgr_pointstodmatrix(
```

```

pgr_texttopoints('2,0;2,1;3,1;2,2;4,1;4,2;2,3;3,2', 0)
);
      ids
-----
{1,2,3,4,5,6,7,8}
(1 row)

```

## See Also

- *pgr\_vidsToDMatrix* - convert a point geometry to the closest vertex\_id of an edge..
- *pgr\_tsp* - Traveling Sales Person

## 9.2.6 pgr\_vidsToDMatrix

### Name

`pgr_vidsToDMatrix` - Creates a distances matrix from an array of `vertex_id`.

### Synopsis

This function takes an array of `vertex_id`, the original array of points used to generate the array of `vertex_id`, an edge table name and a `tol`. It then computes `kdijkstra()` distances for each vertex to all the other vertices and creates a symmetric distance matrix suitable for TSP. The `pnt` array and the `tol` are used to establish a BBOX for limiting selection of edges. The extents of the points is expanded by `tol`.

The function returns:

- **record** - with two fields as describe here
  - **dmatrix** `float8[]` - the distance matrix suitable to pass to `pgrTSP()` function.
  - **ids** `integer[]` - an array of ids for the distance matrix.

```
record pgr_vidsToDMatrix(IN vids integer[], IN pnts geometry[], IN edges text, tol float8 DEFAULT
```

### Description

#### Parameters

**vids** `integer[]` - An array of `vertex_id`.

**pnts** `geometry[]` - An array of point geometries that approximates the extents of the `vertex_id`.

**edges** `text` - The edge table to be used for the conversion.

**tol** `float8` - The amount to expand the BBOX extents of `pnts` when building the graph.

#### Warning:

- we compute a symmetric matrix because TSP requires that so the distances are better the Euclidean but but are not perfect
- `kdijkstra()` can fail to find a path between some of the vertex ids. We to not detect this other than the cost might get set to -1.0, so the `dmatrix` shoule be checked for this as it makes it invalid for TSP

## History

- New in version 2.1.0

## Examples

This example shows how this can be used in the context of feeding the results into `pgr_tsp()` function. We convert a text string of `x,y;x,y;...` into an array of points, then convert that into an array `vertex_id`, then create a distance matrix that gets feed into `pgr_tsp()` that returns the final result.

```
select * from pgr_tsp(
  (select dmatrix::float8[]
   from pgr_vidstodmatrix(
     pgr_pointstovids(
       pgr_texttopoints('2,0;2,1;3,1;2,2;4,1;4,2;2,3;3,2', 0),
       'edge_table'),
     pgr_texttopoints('2,0;2,1;3,1;2,2;4,1;4,2;2,3;3,2', 0),
     'edge_table')
   ),
  1
);
seq | id
-----+-----
 0 | 1
 1 | 3
 2 | 7
 3 | 5
 4 | 4
 5 | 2
 6 | 6
 7 | 0
(8 rows)
```

This example uses the *Sample Data* network.

## See Also

- [pgr\\_textToPoints](#) - Create an array of points from a text string.
- [pgr\\_tsp](#) - Traveling Sales Person

## 9.2.7 pgr\_vidsToDMatrix

### Name

`pgr_vidsToDMatrix` - Creates a distances matrix from an array of `vertex_id`.

### Synopsis

This function takes an array of `vertex_id`, a sql statement to select the edges, and some boolean arguments to control the behavior. It then computes `kdiijkstra()` distances for each vertex to all the other vertices and creates a distance matrix suitable for TSP.

The function returns:

- **dmatrix** float8[] - the distance matrix suitable to pass to `pgr_TSP()` function.

```
pgr_vidsToDMatrix(IN sql text, IN vids integer[], IN directed boolean, IN has_reverse_cost boolean)
```

## Description

### Parameters

**sql** text - A SQL statement to select the edges needed for the solution.

**vids** integer[] - An array of vertex\_id.

**directed** boolean - A flag to indicate if the graph is directed.

**has\_reverse\_cost** boolean - A flag to indicate if the SQL has a column reverse\_cost.

**want\_symmetric** boolean - A flag to indicate if you want a symmetric or asymmetric matrix. You will need a symmetric matrix for pgr\_TSP(). If the matrix is asymmetric, then the cell(i,j) and cell(j,i) will be set to the average of those two cells except if one or the other are -1.0 then it will take the value of the other cell. If both are negative they will be left alone.

#### Warning:

- kdijkstra() can fail to find a path between some of the vertex ids. We do not detect this other than the cost might get set to -1.0, so the dmatrix should be checked for this as it makes it invalid for TSP

### History

- New in version 2.1.0

### Examples

This example shows how this can be used in the context of feeding the results into pgr\_tsp() function. We convert a text string of x,y;x,y;... into an array of points, then convert that into an array vertex\_id, then create a distance matrix that gets feed into pgr\_tsp() that returns the final result.

```
select * from pgr_tsp(
  (select dmatrix::float8[]
   from pgr_vidstodmatrix(
     'select id, source, target, cost, reverse_cost from edge_table',
     pgr_pointstovids(
       pgr_texttopoints('2,0;2,1;3,1;2,2;4,1;4,2;2,3;3,2', 0),
       'edge_table'),
     true, true, true) as dmatrix
  ),
  1
);
seq | id
-----+-----
 0 | 1
 1 | 0
 2 | 6
 3 | 3
 4 | 7
 5 | 5
 6 | 4
 7 | 2
(8 rows)
```

This example uses the *Sample Data* network.

### See Also

- [pgr\\_vidsToDMatrix](#) - - Creates a distances matrix from an array of vertex\_id.

- *pgr\_textToPoints* - Create an array of points from a text string.
- *pgr\_tsp* - Traveling Sales Person



---

## Experimental functions: by GSoC

---

### The following functions are experimental

- They may lack documentation,
- Were created by GSoC students.
- they are unsupported.
- *pgr\_vrpOneDepot* - VRP One Depot
- *pgr\_vrppdtw* - *Pickup and Delivery problem* - Pickup and Delivery problem

## 10.1 pgr\_vrpOneDepot

No documentation available from the original developer

**Warning:** this function is experimental and there is no assurance that it will work

- *pgr\_costResult[]*
- [http://en.wikipedia.org/wiki/Vehicle\\_routing\\_problem](http://en.wikipedia.org/wiki/Vehicle_routing_problem)

## 10.2 pgr\_vrppdtw - Pickup and Delivery problem

### 10.2.1 Name

`pgr_gsoc_vrppdtw` — Returns optimized solution

### 10.2.2 Synopsis

Vehicle Routing Problem with Pickup and Delivery (VRPPD): A number of goods need to be moved from certain pickup locations to other delivery locations. The goal is to find optimal routes for a fleet of vehicles to visit the pickup and drop-off locations.

```
pgr_gsoc_vrppdtw(text sql, integer , integer;
```

### 10.2.3 Description

`sql` a SQL query, which should return a set of rows with the following columns:

```
select * from pgr_gsoc_vrppdtw(
    'select * from customer order by id'::text, 25,200
);
```

Returns set of *pgr\_costResult[]*:

- seq** row sequence
- rid** route ID
- nid** node ID (-1 for the last row)
- cost** cost to traverse to seq

### 10.2.4 Examples

```
SELECT * from pgr_gsoc_vrppdtw(
    'select * from customer order by id'::text, 25,200
);
```

seq	rid	nid	cost
0	7	8	1
1	8	9	1
2	9	15	1
3	12	-1	0
.	.	.	.
.	.	.	.

### 10.2.5 See Also

- *pgr\_costResult[]*
- [http://en.wikipedia.org/wiki/Vehicle\\_routing\\_problem](http://en.wikipedia.org/wiki/Vehicle_routing_problem)

---

## Proposed functions: Proposed by Zia Mohammed

---

### About this proposal:

- Author: Zia Mohammed.
- Status: Needs a lot of testing. I am working on that.
- I did not add automated test.
- Temporary name: `pgr_labelGraph`
- Need: I need feedback from the community.
- *pgr\_labelGraph* - Analyze / label subgraphs within a network

## 11.1 pgr\_labelGraph

### 11.1.1 Name

`pgr_labelGraph` — Locates and labels sub-networks within a network which are not topologically connected. Must be run after `pgr_createTopology()`. No use of geometry column. Only `id`, `source` and `target` columns are required.

### 11.1.2 Synopsis

The function returns:

- OK when a column with provided name has been generated and populated successfully. All connected edges will have unique similar integer values. In case of `rows_where` condition, non participating rows will have -1 integer values.
- FAIL when the processing cannot be finished due to some error. Notice will be thrown accordingly.
- `rows_where` condition generated 0 rows when passed SQL condition has not been fulfilled by any row.

```
varchar pgr_labelGraph(text, text, text, text, text, text)
```

### 11.1.3 Description

A network behind any routing query may consist of sub-networks completely isolated from each other. Possible reasons could be:

- An island with no bridge connecting to the mainland.

- An edge or mesh of edges failed to connect to other networks because of human negligence during data generation.
- The data is not properly noded.
- Topology creation failed to succeed.

`pgr_labelGraph()` will create an integer column (with the name provided by the user) and will assign same integer values to all those edges in the network which are connected topologically. Thus better analysis regarding network structure is possible. In case of `rows_where` condition, non participating rows will have -1 integer values.

Prerequisites: Must run `pgr_createTopology()` in order to generate source and target columns. Primary key column `id` should also be there in the network table.

Function accepts the following parameters:

**edge\_table** text Network table name, with optional schema name.

**id** text Primary key column name of the network table. Default is `id`.

**source** text Source column name generated after `pgr_createTopology()`. Default is `source`.

**target** text Target column name generated after `pgr_createTopology()`. Default is `target`.

**subgraph** text Column name which will hold the integer labels for each sub-graph. Default is `subgraph`.

**rows\_where** text The SQL where condition. Default is `true`, means the processing will be done on the whole table.

### 11.1.4 Possible Usage

– The following should be OK

```
select pgr_labelGraph('ways');
select pgr_labelGraph('Ways');
select pgr_labelGraph('ways', 'id');
select pgr_labelGraph('ways', 'id', 'source');
select pgr_labelGraph('ways', 'id', 'source', 'target');
select pgr_labelGraph('ways', 'id', 'source', 'target', 'subgraph');
select pgr_labelGraph('ways', 'id', 'source', 'target', 'subgraph', 'id<100');
```

– When table located in another schema `e03`

```
select pgr_labelGraph('e03.ways');
select pgr_labelGraph('e03.Ways');
select pgr_labelGraph('e03.ways', 'id');
select pgr_labelGraph('e03.ways', 'id', 'source');
select pgr_labelGraph('e03.ways', 'id', 'source', 'target');
select pgr_labelGraph('e03.ways', 'id', 'source', 'target', 'subgraph');
select pgr_labelGraph('e03.ways', 'id', 'source', 'target', 'subgraph', 'id<100');
```

– When using the named notation

```
select pgr_labelGraph('e03.calles', target:='destino', subgraph:='subgraph', id:='gido', source:=
select pgr_labelGraph('e03.calles', rows_where:='gido<100', id:='gido', source:='salida', target:=
```

– The following should FAIL

```
select pgr_labelGraph('id', 'ways');
select pgr_labelGraph('ways', 'id', 'sourc', 'target');
select pgr_labelGraph('ways', 'id', 'source', 'Target');
select pgr_labelGraph('ways', 'id', 'source', 'target', 'subgraph', 'id<');
```

– When table located in another schema e03

```
select pgr_labelGraph('e03.calles');
select pgr_labelGraph('e03.Calles');
select pgr_labelGraph('id', 'e03.calles');
select pgr_labelGraph('e03.calles', 'id', 'sourc', 'target');
select pgr_labelGraph('e03.calles', 'gido', 'source', 'target', 'subgraph', 'id<');
select pgr_labelGraph('e03.calles', 'gid', 'salida', 'target', 'subgraph', 'id<10');
select pgr_labelGraph('e03.calles', 'gid', 'salida', 'destino', 'subgraph', 'id<10 AND id>100');
```

– When using the named notation

```
select pgr_labelGraph('e03.calles', target:='destino', subgraph:='subgraph', id:='gido');
select pgr_labelGraph('e03.calles', target:='destino', subgraph:='subgraph', id:='gido', source:=);
select pgr_labelGraph(rows_where:='gido<100', id:='gido', source:='salida', 'e03.calles', target:=);
```

– The following should return “rows\_where condition generated 0 rows”

```
select pgr_labelGraph('ways', 'id', 'source', 'target', 'subgraph', 'id<10 AND id>100');
select pgr_labelGraph('e03.calles', id:='gido', rows_where:='gido<100 AND gido>200', source:=);
```

## 11.1.5 Examples Output pane Messages

```
NOTICE: Processing:
NOTICE: pgr_labelGraph('ways', 'id', 'source', 'target', 'subgraph', 'true')
NOTICE: Performing initial checks, please hold on ...
NOTICE: Starting - Checking table ...
NOTICE: Ending - Checking table
NOTICE: Starting - Checking columns
NOTICE: Ending - Checking columns
NOTICE: Starting - Checking rows_where condition
NOTICE: Ending - Checking rows_where condition
NOTICE: Starting - Calculating subgraphs
NOTICE: Successfully complicated calculating subgraphs
NOTICE: Ending - Calculating subgraphs

Total query runtime: 5426 ms.
1 row retrieved.

pgr_labelgraph
character varying
-----
OK
(1 row)
```

## 11.1.6 See Also

- [pgr\\_createTopology](#)<sup>1</sup> to create the topology of a table based on its geometry and tolerance value.

## 11.2 Developers’s Functions

Functions to help you develop a wrapper or a recipe

- [pgr\\_getColumnName](#) - to get the name of the column as is stored in the postgres administration tables.
- [pgr\\_getTableName](#) - to retrieve the name of the table as is stored in the postgres administration tables.
- [pgr\\_isColumnIndexed](#) - to check if the column is indexed.

<sup>1</sup>[https://github.com/Zia-/pgrouting/blob/develop/src/common/sql/pgrouting\\_topology.sql](https://github.com/Zia-/pgrouting/blob/develop/src/common/sql/pgrouting_topology.sql)

- *pgr\_isColumnInTable* - to check only for the existence of the column.
- *pgr\_pointToId* -to insert/get the id of the inserted point in a vertices table.
- *pgr\_quote\_ident* - to quotes the input text to be used as an identifier in an SQL statement string.
- *pgr\_versionless* - to compare two version numbers.
- *pgr\_startPoint* - to get the start point of a (multi)linestring.
- *pgr\_endPoint* - to get the end point of a (multi)linestring.

## 11.2.1 pgr\_getColumnName

### Name

`pgr_getColumnName` — Retrieves the name of the column as is stored in the postgres administration tables.

---

**Note:** This function is intended for the developer’s aid.

---

**Warning:** This function is deprecated in 2.1. Use `_pgr_getColumnName` instead

### Synopsis

Returns a text contining the registered name of the column.

```
text pgr_getColumnName(tab text, col text);
```

### Description

#### Parameters

**tab** text table name with or without schema component.

**col** text column name to be retrived.

#### Returns

- text containing the registered name of the column.
- NULL when :
  - The table “tab” is not found or
  - Column “col” is not found in table “tab” in the postgres administration tables.

### History

- New in version 2.0.0

### Examples

```
SELECT pgr_getColumnName('edge_table','the_geom');

pgr_iscolumnintable
-----
the_geom
(1 row)
```

```

SELECT pgr_getColumnName('edge_table', 'The_Geom');

pgr_iscolumnintable
-----
the_geom
(1 row)

```

The queries use the *Sample Data* network.

### See Also

- *Developer's Guide* for the tree layout of the project.
- *pgr\_isColumnInTable* to check only for the existence of the column.
- *pgr\_getTableName* to retrieve the name of the table as is stored in the postgres administration tables.

## 11.2.2 pgr\_getTableName

### Name

`pgr_getTableName` — Retrieves the name of the column as is stored in the postgres administration tables.

**Note:** This function is intended for the developer's aid.

**Warning:** This function is deprecated in 2.1 Use `_pgr_getTableName` instead

### Synopsis

Returns a record containing the registered names of the table and of the schema it belongs to.

```
(text sname, text tname) pgr_getTableName(text tab)
```

### Description

#### Parameters

**tab** text table name with or without schema component.

#### Returns

##### sname

- text containing the registered name of the schema of table “tab”.
  - when the schema was not provided in “tab” the current schema is used.
- NULL when :
  - The schema is not found in the postgres administration tables.

##### tname

- text containing the registered name of the table “tab”.
- NULL when :
  - The schema is not found in the postgres administration tables.
  - The table “tab” is not registered under the schema `sname` in the postgres administration tables

## History

- New in version 2.0.0

## Examples

```
SELECT * from pgr_getTableName('edge_table');

sname | tname
-----+-----
public | edge_table
(1 row)

SELECT * from pgr_getTableName('EdgeTable');

sname | tname
-----+-----
public |
(1 row)

SELECT * from pgr_getTableName('data.Edge_Table');
sname | tname
-----+-----
      |
(1 row)
```

The examples use the *Sample Data* network.

## See Also

- *Developer's Guide* for the tree layout of the project.
- `pgr_isColumnInTable` to check only for the existence of the column.
- `pgr_getTableName` to retrieve the name of the table as is stored in the postgres administration tables.

## 11.2.3 pgr\_isColumnIndexed

### Name

`pgr_isColumnIndexed` — Check if a column in a table is indexed.

---

**Note:** This function is intended for the developer's aid.

---

**Warning:** This function is deprecated in 2.1 Use `_pgr_isColumnIndexed` instead

### Synopsis

Returns `true` when the column “col” in table “tab” is indexed.

```
boolean pgr_isColumnIndexed(text tab, text col);
```



## Description

**tab** text Table name with or without schema component.

**col** text Column name to be checked for.

Returns:

- true when the column “col” in table “tab” is indexed.
- false when:
  - The table “tab” is not found or
  - Column “col” is not found in table “tab” or
  - Column “col” in table “tab” is not indexed

## History

- New in version 2.0.0

## Examples

```
SELECT pgr_isColumnIndexed('edge_table', 'x1');

pgr_iscolumnindexed
-----
f
(1 row)

SELECT pgr_isColumnIndexed('public.edge_table', 'cost');

pgr_iscolumnindexed
-----
f
(1 row)
```

The example use the *Sample Data* network.

## See Also

- *Developer’s Guide* for the tree layout of the project.
- *pgr\_isColumnInTable* to check only for the existence of the column in the table.
- *pgr\_getColumnName* to get the name of the column as is stored in the postgres administration tables.
- *pgr\_getTableName* to get the name of the table as is stored in the postgres administration tables.

## 11.2.4 pgr\_isColumnInTable

### Name

pgr\_isColumnInTable — Check if a column is in the table.

**Note:** This function is intended for the developer’s aid.

**Warning:** This function is deprecated in 2.1 Use *\_pgr\_isColumnInTable* instead

## Synopsis

Returns true when the column “col” is in table “tab”.

```
boolean pgr_isColumnInTable(text tab, text col);
```

## Description

**tab** text Table name with or without schema component.

**col** text Column name to be checked for.

Returns:

- true when the column “col” is in table “tab”.
- false when:
  - The table “tab” is not found or
  - Column “col” is not found in table “tab”

## History

- New in version 2.0.0

## Examples

```
SELECT pgr_isColumnInTable('edge_table', 'x1');

 pgr_iscolumnintable
-----
t
(1 row)

SELECT pgr_isColumnInTable('public.edge_table', 'foo');

 pgr_iscolumnintable
-----
f
(1 row)
```

The example use the *Sample Data* network.

## See Also

- *Developer’s Guide* for the tree layout of the project.
- *pgr\_isColumnIndexed* to check if the column is indexed.
- *pgr\_getColumnName* to get the name of the column as is stored in the postgres administration tables.
- *pgr\_getTableName* to get the name of the table as is stored in the postgres administration tables.

## 11.2.5 pgr\_pointToId

### Name

pgr\_pointToId — Inserts a point into a vertices table and returns the corresponig id.

---

**Note:** This function is intended for the developer's aid.  
Use `pgr_createTopology` or `pgr_createVerticesTable` instead.

**Warning:** This function is deprecated in 2.1

## Synopsis

This function returns the `id` of the row in the vertices table that corresponds to the `point` geometry

```
bigint pgr_pointToId(geometry point, double precision tolerance, text vertname text, integer srid)
```

## Description

**point** geometry “POINT” geometry to be inserted.

**tolerance** float8 Snapping tolerance of disconnected edges. (in projection unit)

**vertname** text Vertices table name WITH schema included.

**srid** integer SRID of the geometry point.

This function returns the `id` of the row that corresponds to the `point` geometry

- When the `point` geometry already exists in the vertices table `vertname`, it returns the corresponding `id`.
- When the `point` geometry is not found in the vertices table `vertname`, the function inserts the `point` and returns the corresponding `id` of the newly created vertex.

**Warning:** The function do not perform any checking of the parameters. Any validation has to be done before calling this function.

## History

- Renamed in version 2.0.0

## See Also

- *Developer's Guide* for the tree layout of the project.
- `pgr_createVerticesTable` to create a topology based on the geometry.
- `pgr_createTopology` to create a topology based on the geometry.

## 11.2.6 pgr\_quote\_ident

### Name

`pgr_quote_ident` — Quotes the input text to be used as an identifier in an SQL statement string.

**Note:** This function is intended for the developer's aid.

**Warning:** This function is deprecated in 2.1 Use `_pgr_quote_ident` instead

## Synopsis

Returns the given identifier `idname` suitably quoted to be used as an identifier in an SQL statement string.

```
text pgr_quote_ident (text idname);
```

## Description

### Parameters

**idname** text Name of an SQL identifier. Can include . dot notation for schemas.table identifiers

Returns the given string suitably quoted to be used as an identifier in an SQL statement string.

- When the identifier `idname` contains on or more . separators, each component is suitably quoted to be used in an SQL string.

### History

- New in version 2.0.0

### Examples

Everything is lower case so nothing needs to be quoted.

```
SELECT pgr_quote_ident ('the_geom');

pgr_quote_ident
-----
 the_geom
(1 row)

SELECT pgr_quote_ident ('public.edge_table');

pgr_quote_ident
-----
 public.edge_table
(1 row)
```

The column is upper case so its double quoted.

```
SELECT pgr_quote_ident ('edge_table.MYGEOM');

pgr_quote_ident
-----
 edge_table."MYGEOM"
(1 row)

SELECT pgr_quote_ident ('public.edge_table.MYGEOM');

pgr_quote_ident
-----
 public.edge_table."MYGEOM"
(1 row)
```

The schema name has a capital letter so its double quoted.

```
SELECT pgr_quote_ident('Myschema.edge_table');

 pgr_quote_ident
-----
"Myschema".edge_table
(1 row)
```

Ignores extra . separators.

```
SELECT pgr_quote_ident('Myschema...edge_table');

 pgr_quote_ident
-----
"Myschema".edge_table
(1 row)
```

### See Also

- *Developer's Guide* for the tree layout of the project.
- `pgr_getTableName` to get the name of the table as is stored in the postgres administration tables.

## 11.2.7 pgr\_versionless

### Name

`pgr_versionless` — Compare two version numbers.

**Note:** This function is intended for the developer's aid.

**Warning:** This function is deprecated in 2.1. Use `_pgr_versionless` instead

### Synopsis

Returns `true` if the first version number is smaller than the second version number. Otherwise returns `false`.

```
boolean pgr_versionless(text v1, text v2);
```

### Description

- v1** text first version number
- v2** text second version number

### History

- New in version 2.0.0

### Examples

```
SELECT pgr_versionless('2.0.1', '2.1');

 pgr_versionless
-----
 t
(1 row)
```

### See Also

- *Developer's Guide* for the tree layout of the project.
- *pgr\_version* to get the current version of pgRouting.

## 11.2.8 pgr\_startPoint

### Name

`pgr_startPoint` — Returns a start point of a (multi)linestring geometry.

---

**Note:** This function is intended for the developer's aid.

---

**Warning:** This function is deprecated in 2.1 Use `_pgr_startPoint` instead

### Synopsis

Returns the geometry of the start point of the first LINESTRING of `geom`.

```
geometry pgr_startPoint(geometry geom);
```

### Description

#### Parameters

**geom** `geometry` Geometry of a MULTILINESTRING or LINESTRING.

Returns the geometry of the start point of the first LINESTRING of `geom`.

#### History

- New in version 2.0.0

### See Also

- *Developer's Guide* for the tree layout of the project.
- *pgr\_endPoint* to get the end point of a (multi)linestring.

## 11.2.9 pgr\_endPoint

### Name

`pgr_endPoint` — Returns an end point of a (multi)linestring geometry.

**Note:** This function is intended for the developer's aid.

**Warning:** This function is being deprecated on 2.1. Use `_pgr_endPoint` instead

### Synopsis

Returns the geometry of the end point of the first LINESTRING of `geom`.

```
text pgr_startPoint(geometry geom);
```

### Description

#### Parameters

**geom** geometry Geometry of a MULTILINESTRING or LINESTRING.

Returns the geometry of the end point of the first LINESTRING of `geom`.

#### History

- New in version 2.0.0

#### See Also

- *Developer's Guide* for the tree layout of the project.
- `pgr_startPoint` to get the start point of a (multi)linestring.





---

## Discontinued Functions

---

### 12.1 Discontinued Functions

Especially with new major releases functionality may change and functions may be discontinued for various reasons. Functionality that has been discontinued will be listed here.

#### 12.1.1 Shooting Star algorithm

**Version** Removed with 2.0.0

**Reasons** Unresolved bugs, no maintainer, replaced with *pgr\_trsp - Turn Restriction Shortest Path (TRSP)*

**Comment** Please *contact us* if you're interested to sponsor or maintain this algorithm. The function signature is still available in *common\_legacy* but it is just a wrapper that throws an error. We have not included any of the old code for this in this release.



*Developers's Functions*

**Warning:** In V3.0 This function are going to be discontinued. Use the already available underscored version instead.

**Warning:** *Developers's Functions* documentation is going to be deleted from the pgRouting documentation in V3.0

The following functions are used internally the topology functions.

- *pgr\_getColumnName* - to get the name of the column as is stored in the postgres administration tables.
- *pgr\_getTableName* - to retrieve the name of the table as is stored in the postgres administration tables.
- *pgr\_isColumnIndexed* - to check if the column is indexed.
- *pgr\_isColumnInTable* - to check only for the existance of the column.
- *pgr\_pointToId* -to insert/get the id of the inserted point in a vertices table.
- *pgr\_quote\_ident* - to quotes the input text to be used as an identifier in an SQL statement string.
- *pgr\_versionless* - to compare two version numbers.
- *pgr\_startPoint* - to get the start point of a (multi)linestring.
- *pgr\_endPoint* - to get the end point of a (multi)linestring.

## 13.1 Developer's Guide

### 13.1.1 Source Tree Layout

*cmake/* cmake scripts used as part of our build system.

*src/* This is the algorithm source tree. Each algorithm is to be contained in its on sub-tree with */doc*, */sql*, */src*, and */test* sub-directories.

For example:

- *src/dijkstra* Main direcotry for dijkstra algorithm.
- *src/dijkstra/doc* Dijkstra's documentation directory.
- *src/dijkstra/src* Dijkstra's C and/or C++ code.
- *src/dijkstra/sql* Dijkstra's sql code.
- *src/dijkstra/test* Dijkstra's tests.

## 13.1.2 Tools

*tools/* Miscellaneous scripts and tools.

### pre-commit

To keep version/branch/commit up to date install pelase do the following:

```
cp tools/pre-commit .git/hooks/pre-commit
```

After each commit a the file **VERSION** will remain. (The hash number will be one behind)

### tester

The tester is executed from the top level of the tree layout:

```
tools/test-runner.pl --help
```

### doxygen

**Warning:** *Developers's Functions* documentation is going to be deleted from the pgRouting documentation and included in the doxygen documentation.

To use doxygen:

```
cd tools/doxygen/  
make
```

The code's documentation can be found in:

```
build/doxy/html/index.html
```

### cpplint

We try to follow the following guidelines for C++ coding:

<https://google-styleguide.googlecode.com/svn/trunk/cppguide.html>

Sample use:

```
python cpplint.py ../src/dijkstra/src/dijkstra_driver.h  
../src/dijkstra/src/dijkstra_driver.h:34: Lines should be <= 80 characters long [whitespace/lin  
../src/dijkstra/src/dijkstra_driver.h:40: Line ends in whitespace. Consider deleting these ext  
Done processing ../src/dijkstra/src/dijkstra_driver.h  
Total errors found: 2
```

- Maybe line 34 is a very complicated calculation so you can just ignore the message
- Delete whitespace at end of line is easy fix.
- Use your judgement!!!

Some files like `postgres.h` are system dependant so don't include the directory.

## Other tools

### Tools like:

- `doit`
- `winnie`
- `publish_doc.sh`

are very specific for the deployment of new versions, so please ask first!

### 13.1.3 Documentation Layout

---

**Note:** All documentation should be in reStructuredText format. See: <http://docutils.sf.net/rst.html> for introductory docs.

---

Documentation is distributed into the source tree. This top level “doc” directory is intended for high level documentation cover subjects like:

- Compiling and testing
- Installation
- Tutorials
- Users’ Guide front materials
- Reference Manual front materials
- etc


Since the algorithm specific documentation is contained in the source tree with the algorithm specific files, the process of building the documentation and publishing it will need to assemble the details with the front material as needed.

Also, to keep the “doc” directory from getting cluttered, each major book like those listed above, should be contained in a separate directory under “doc”. Any images or other materials related to the book should also be kept in that directory.

### 13.1.4 Testing Infrastructure

There is a very basic testing infrastructure put in place. Here are the basics of how it works. We need more test cases. Longer term we should probably get someone to setup travis-ci or jenkins testing frameworks.

Here is the graph for the TRSP tests.



Tests are run via the script at the top level `tools/test-runner.pl` and it runs all the test configured tests and at the moment just dumps the results structure of the test. This can be prettied up later.

It also assumes that you have installed the libraries as it tests using the installed postgresql. This probably needs to be made smarter so we can test out of the build tree. I'll need to think about that.

Basically each `.../test/` directory should include one `test.conf` file that is a perl script fragment that defines what data files to load and what tests to run. I have built in some mechanisms to allow test and data to be pg version and postgis version specific, but I'm not using that yet. So for example, `core/trsp/test/test-any-00.data` is a sql plain text dump that will load and needed data for a set of tests. This is also the graph in the image above. You can specify multiple files to load, but as a group they need to have unique names.

`core/trsp/test/test-any-00.test` is a sql command to be run. It will get run as:

```
psql ... -A -t -q -f file.test dbname > tmpfile
diff -w file.rest tmpfile
```

Then if there is a difference then an test failure is reported.

#### Release Notes

- [pgRouting 2.1.0 Release Notes](#)

- [pgRouting 2.0 Release Notes](#)
- [pgRouting 1.x Release Notes](#)

## 13.2 Release Notes

- [pgRouting 2.1.0 Release Notes](#)
- [pgRouting 2.0 Release Notes](#)
- [pgRouting 1.x Release Notes](#)

### 13.2.1 pgRouting 2.1.0 Release Notes

With the release of pgRouting 2.1.0 fixes some bugs and issues.

- To see the full list of changes check the list of [Git commits](#)<sup>1</sup> on Github.
- To see the issues closed by this release see the [Git closed issues](#)<sup>2</sup> on Github.
- For important changes see the following release notes.

#### Release Notes

- 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
- Three Functions where completely re-factored using the developers library and functionality was added.
  - pgr\_dijkstra
  - pgr\_ksp
  - pgr\_drivingDistance
- Improved - Alphashape function now can generate better (multi)polygon with holes and alpha parameter.
- Instead of generating many libraries: - All functions are encapsulated in one library - The library has a the prefix 2-1-0
- Added proposed functions to be evaluated:
  - 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.

<sup>1</sup><https://github.com/pgRouting/pgrouting/commits>

<sup>2</sup><https://github.com/pgRouting/pgrouting/issues?q=is%3Aissue+milestone%3A%22Release+2.1.0%22+is%3Aclosed>

\* `pgr_vidsToDMatrix` - Create a distance matrix from an array of `vertex_id`.

- Added proposed functions from GSoc Projects:
  - `pgr_vrppdtw`
- Removed the 1.x legacy functions
- Some bug fixes in other functions

## 13.2.2 pgRouting 2.0 Release Notes

With the release of pgRouting 2.0 the library has abandoned backwards compatibility to *pgRouting 1.x* releases. We did this so we could restructure pgRouting, standardize the function naming, and prepare the project for future development. As a result of this effort, we have been able to simplify pgRouting, add significant new functionality, integrate documentation and testing into the source tree and make it easier for multiple developers to make contribution.

For important changes see the following release notes. To see the full list of changes check the list of [Git commits](#)<sup>3</sup> on Github.

### Changes for 2.0.0

- 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 refactored and unified
- Support for table SCHEMA in function parameters
- Support for `st_` PostGIS function prefix
- Added `pgr_` prefix to functions and types
- Better documentation: <http://docs.pgrouting.org>

---

<sup>3</sup><https://github.com/pgRouting/pgrouting/commits>



### 13.2.3 pgRouting 1.x Release Notes

The following release notes have been copied from the previous `RELEASE_NOTES` file and are kept as a reference. Release notes starting with *version 2.0.0* will follow a different schema.

#### Changes for release 1.05

- Bugfixes

#### Changes for release 1.03

- Much faster topology creation
- Bugfixes

#### Changes for release 1.02

- Shooting\* bugfixes
- Compilation problems solved

#### Changes for release 1.01

- Shooting\* bugfixes

#### Changes for release 1.0

- Core and extra functions are separated
- Cmake build process
- Bugfixes

#### Changes for release 1.0.0b

- Additional SQL file with more simple names for wrapper functions
- Bugfixes

#### 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 couldn't 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

### **Indices and tables**

- genindex
- search