

Graphs in the **gRbase** package

Søren Højsgaard

gRbase version 1.7-0 as of 2014-03-22

Contents

1	Introduction	1
2	Graphs	2
2.1	Undirected graphs	2
2.2	Directed acyclic graphs (DAGs)	3
2.3	Other types of graphs	4
2.4	Graph properties	5
2.5	Adjacency matrices	5
2.6	Graph coercion	6
2.7	Plotting graphs	7
3	Advanced graph operations	8
3.1	Moralization	8
3.2	Topological sort	9
3.3	Getting the cliques of an undirected graph	9
3.4	Perfect ordering, maximum cardinality search and triangulation	10
3.5	RIP ordering / junction tree	11
3.6	Minimal triangulation and maximum prime subgraph decomposition	12
4	Time and space considerations	13
4.1	Time	13
4.2	Space	14
5	Graph queries	14

1 Introduction

For the R community, the packages **graph**, **RBGL**, **Rgraphviz** and **igraph** are extremely useful tools for graph operations, manipulation and layout. The **gRbase** package adds some additional tools to these fine packages. The most important tools are:

1. Undirected and directed acyclic graphs can be specified using formulae or an adjacency list using the functions `ug()` and `dag()`.

This gives graphs represented in one of the following forms:¹

- A `graphNEL` object (the default),
- A dense adjacency matrix (a “standard” matrix in R).

¹There is a fourth form: **igraph** objects. These, however, will probably not be supported in the future.

- A sparse adjacency matrix(a `dgCMatrix` from the `Matrix` package).
2. Some graph algorithms are implemented in `gRbase`. These can be applied to graphs represented as `graphNELs` and matrices. The most important algorithms are:
- `moralize()`, (moralize a directed acyclic graph)
 - `mcs()`,(maximum cardinality search for undirected graph)
 - `triangulate()`, (triangulate undirected graph)
 - `rip()`, (RIP ordering of cliques of triangulated undirected graph)
 - `jTree()` (Create junction tree from triangualted undirected graph).
 - `getCliques()`, (get the (maximal) cliques of an undirected graph)
 - `minimalTriang()` (minimal triangulation of undirected graph)
 - `mpd()` (maximal prime subgraph decomposition of undirected graph)

The general scheme is the following: For example, for maximum cardinality search there is a `mcs()` function and a `mcs.default()` method performs maximum cardinality search for graphs represented as `graphNELs` and as sparse and dense matrices. The workhorse is the function `mcsMAT()` which takes a sparse or a dense matrix as input.

2 Graphs

Undirected graphs can be created by the `ug()` function and directed acyclic graphs (DAGs) by the `dag()` function. The graphs can be specified either using formulae or a list of vectors; see examples below.

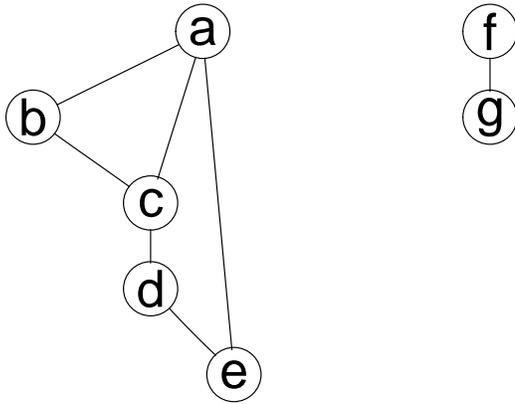
2.1 Undirected graphs

An undirected graph is created by the `ug()` function. The following specifications are equivalent (notice that “:” and “*” can be used interchangeably):

```
R> ug11 <- ug( ~a:b:c + c:d + d:e + a:e + f:g )
R> ug12 <- ug( c("a","b","c"),c("c","d"),c("d","e"),c("a","e"),c("f","g") )
R> ug13 <- ug(~a:b:c, ~c:d, ~d:e + a:e + f:g)
R> ug14 <- ugList( list(c("a","b","c"),c("c","d"),c("d","e"),c("a","e"),c("f","g")) )
```

Default is to return a `graphNEL` object (for which there is a plot method):

```
R> ug11
A graphNEL graph with undirected edges
Number of Nodes = 7
Number of Edges = 7
R> plot( ug11 )
```



A representation as a dense adjacency matrix can be obtained with:

```

R> ug11m <- ug(~a*b*c + c*d + d*e + a*e + f*g, result="matrix"); ug11m
  a b c d e f g
a 0 1 1 0 1 0 0
b 1 0 1 0 0 0 0
c 1 1 0 1 0 0 0
d 0 0 1 0 1 0 0
e 1 0 0 1 0 0 0
f 0 0 0 0 0 0 1
g 0 0 0 0 0 1 0

```

A representation as a sparse adjacency matrix is obtained with

```

R> ug11M <- ug(~a*b*c + c*d + d*e + a*e + f*g, result="dgCMatrix"); ug11M
7 x 7 sparse Matrix of class "dgCMatrix"
  a b c d e f g
a . 1 1 . 1 . .
b 1 . 1 . . . .
c 1 1 . 1 . . .
d . . 1 . 1 . .
e 1 . . 1 . . .
f . . . . . 1 .
g . . . . . 1 .

```

2.2 Directed acyclic graphs (DAGs)

A directed acyclic graph is created by the `dag()` function. The following specifications are equivalent (notice that “.” and “*” can be used interchangeably):

```

R> dag11 <- dag(~a + b:a + c:a:b + d:c:e + e:a + g:f)
R> dag12 <- dag("a", c("b","a"), c("c","a","b"), c("d","c","e"),
               c("e","a"),c("g","f"))
R> dag13 <- dag(~a, ~b:a, ~c:a:b, ~d:c:e, ~e:a, ~g:f)
R> dag14 <- dagList(list( "a", c("b","a"), c("c","a","b"), c("d","c","e"),
                          c("e","a"),c("g","f") ))

```

Notice: Above, `~a` (and `"a"`) means that “a” has no parents while `~c:a:b` (and `c("c","a","b")`) means that “c” has parents “a” and “b”.

```

R> dag11
A graphNEL graph with directed edges
Number of Nodes = 7
Number of Edges = 7

```

A representation as a dense adjacency matrix can be obtained with

```
R> dag11m <- dag(~a + b:a + c:a:b + d:c:e + e:a + g:f, result="matrix")
  a b c d e g f
a 0 1 1 0 1 0 0
b 0 0 1 0 0 0 0
c 0 0 0 1 0 0 0
d 0 0 0 0 0 0 0
e 0 0 0 1 0 0 0
g 0 0 0 0 0 0 0
f 0 0 0 0 0 1 0
```

A sparse matrix representation is obtained with:

```
R> dag11M <- dag(~a + b:a + c:a:b + d:c:e + e:a + g:f, result="dgCMatrix")
7 x 7 sparse Matrix of class "dgCMatrix"
  a b c d e g f
a . 1 1 . 1 . .
b . . 1 . . . .
c . . . 1 . . .
d . . . . . . .
e . . . 1 . . .
g . . . . . . .
f . . . . . 1 .
```

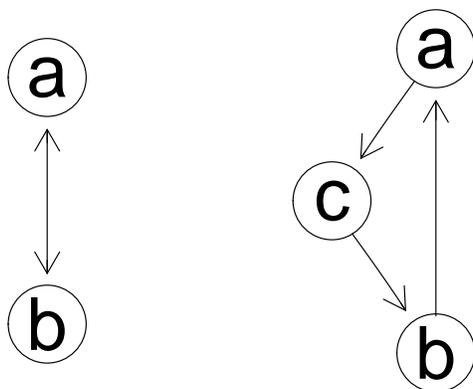
2.3 Other types of graphs

The `dag()` function allows for creation of directed graphs which are not DAGs. Consider

```
R> d1.bi <- dag(~a:b + b:a)
R> edgemode( d1.bi )
[1] "directed"
R> str( nodes(d1.bi) )
chr [1:2] "a" "b"
R> str( edges(d1.bi) )
List of 2
 $ a: chr "b"
 $ b: chr "a"
```

This graph is not DAG because there is an edge from `a` to `b` and from `b` to `a`; i.e., the edge is bidirected. Likewise we may create:

```
R> d2.cyc <- dag(~a:b+b:c+c:a)
R> par(mfrow=c(1,2)); plot(d1.bi); plot(d2.cyc)
```



Notice: Supplying `dag()` with `forceCheck=TRUE` forces `dag()` to check if the graph is acyclic:

```
R> print( try( dag(~a:b+b:c+c:a, forceCheck=TRUE) ) )
[1] "Error : In dag/dagList: Graph is not a DAG\n"
attr("class")
[1] "try-error"
attr("condition")
<simpleError: In dag/dagList: Graph is not a DAG>
```

2.4 Graph properties

The `graphNEL` objects are very convenient to work with because there are many different methods for such objects, for example:

```
R> str( graph::nodes( ug11 ) )
chr [1:7] "a" "b" "c" "d" "e" "f" "g"
R> str( graph::edges( ug11 ) )
List of 7
 $ a: chr [1:3] "b" "c" "e"
 $ b: chr [1:2] "c" "a"
 $ c: chr [1:3] "d" "a" "b"
 $ d: chr [1:2] "e" "c"
 $ e: chr [1:2] "d" "a"
 $ f: chr "g"
 $ g: chr "f"
R> graph::edgeMatrix( ug11 )
      [,1] [,2] [,3] [,4] [,5] [,6] [,7]
from    1    1    1    2    3    4    6
to      2    3    5    3    4    5    7
```

and

```
R> str( RBGL::maxClique( ug11 ) )
List of 1
 $ maxCliques:List of 5
  ..$ : chr [1:3] "a" "b" "c"
  ..$ : chr [1:2] "a" "e"
  ..$ : chr [1:2] "d" "c"
  ..$ : chr [1:2] "d" "e"
  ..$ : chr [1:2] "f" "g"
```

2.5 Adjacency matrices

Graphs in the `graph` package (i.e. `graphNEL` objects) are represented as adjacency lists. However, there is a substantial overhead (in terms of computing time) for such objects. The graph algorithms in `gRbase` are mostly based on a representation as sparse adjacency matrices (which leads to faster code).

A non-zero value at entry (i, j) in an adjacency matrix A for a graph means that there is an edge from i to j . If also (j, i) is non-zero then there is also an edge from j to i . In this case we may think of a bidirected edge between i and j or we may think of the edge as being undirected. We do not distinguish between undirected and bidirected edges in the `gRbase` package. Put differently, in `gRbase`, edges are either directed or undirected/bidirected. In contrast, with `graphNEL` objects one can work with three types of edges: undirected, directed and bidirected edges.

- The function `is.UG()` checks if the adjacency matrix is symmetric (If applied to a `graphNEL`, the adjacency matrix is created and checked for symmetry.)

- The function `is.TUG()` checks if the graph is undirected and triangulated (also called chordal) by checking if the adjacency matrix is symmetric and the vertices can be given a perfect ordering using maximum cardinality search.
- The function `is.DG()` checks if a graph is directed, i.e., that there are no undirected edges. This is done by computing the elementwise product of `A` and the transpose of `A`; if there are no non-zero entries in this product then the graph is directed.
- The function `is.DAG()` will return `TRUE` if all edges are directed and if there are no cycles in the graph. (This is checked by checking if the vertices in the graph can be given a topological ordering which is based on identifying an undirected edge with a bidirected edge).

Notice a special case, namely if the graph has no edges at all (such that the adjacency matrix consists only of zeros). Such a graph is both undirected, triangulated, directed and directed acyclic.

Consider these examples (where `isDirected()` is a method from the **graph** package):

```
R> properties <- function(x){
  c(is.UG=is.UG(x), is.TUG=is.TUG(x), is.DG=is.DG(x), is.DAG=is.DAG(x),
    isD=graph::isDirected(x))
}
R> properties( ug11 )
  is.UG is.TUG  is.DG is.DAG  isD
  TRUE  FALSE  FALSE  FALSE  FALSE
R> properties( dag11 )
  is.UG is.TUG  is.DG is.DAG  isD
  FALSE  FALSE   TRUE   TRUE   TRUE
R> properties( d1.bi )
  is.UG is.TUG  is.DG is.DAG  isD
  TRUE  TRUE  FALSE  FALSE   TRUE
R> properties( d2.cyc )
  is.UG is.TUG  is.DG is.DAG  isD
  FALSE  FALSE   TRUE  FALSE   TRUE
```

2.6 Graph coercion

Graphs can be coerced between different representations using `as()`; for example:

```
R> (mat <- as(ug11, "matrix"))
  a b c d e f g
a 0 1 1 0 1 0 0
b 1 0 1 0 0 0 0
c 1 1 0 1 0 0 0
d 0 0 1 0 1 0 0
e 1 0 0 1 0 0 0
f 0 0 0 0 0 0 1
g 0 0 0 0 0 1 0
R> (Mat <- as(mat, "dgCMatrix"))
7 x 7 sparse Matrix of class "dgCMatrix"
  a b c d e f g
a . 1 1 . 1 . .
b 1 . 1 . . . .
c 1 1 . 1 . . .
d . . 1 . 1 . .
e 1 . . 1 . . .
f . . . . . 1
g . . . . . 1 .
```

```

R> (NEL <- as(Mat, "graphNEL"))
A graphNEL graph with undirected edges
Number of Nodes = 7
Number of Edges = 7

Coercion is generally faster to with coerceGraph():
R> (mat <- coerceGraph(ug11, "matrix"))
  a b c d e f g
a 0 1 1 0 1 0 0
b 1 0 1 0 0 0 0
c 1 1 0 1 0 0 0
d 0 0 1 0 1 0 0
e 1 0 0 1 0 0 0
f 0 0 0 0 0 0 1
g 0 0 0 0 0 1 0

R> (Mat <- coerceGraph(mat, "dgCMatrix"))
7 x 7 sparse Matrix of class "dgCMatrix"
  a b c d e f g
a . 1 1 . 1 . .
b 1 . 1 . . . .
c 1 1 . 1 . . .
d . . 1 . 1 . .
e 1 . . 1 . . .
f . . . . . 1
g . . . . . 1 .

R> (NEL <- as(Mat, "graphNEL"))
A graphNEL graph with undirected edges
Number of Nodes = 7
Number of Edges = 7

R> if( require(microbenchmark) ){
  microbenchmark( as(ug11, "matrix"), coerceGraph(ug11, "matrix") )
}

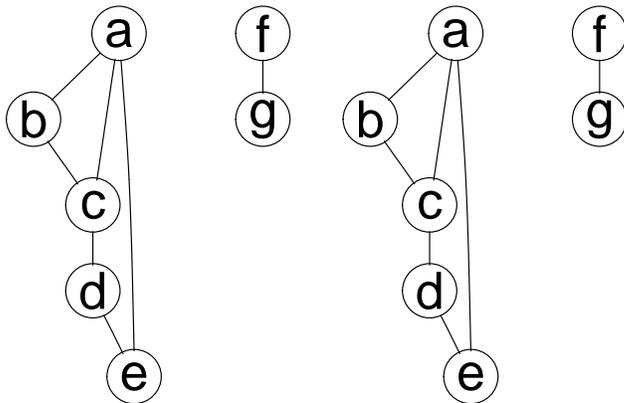
Unit: microseconds
      expr      min       lq   median       uq      max   neval
as(ug11, "matrix") 270.1 274.3 288.1 294.4 405.4    100
coerceGraph(ug11, "matrix") 137.6 140.0 141.8 156.3 353.6    100

```

2.7 Plotting graphs

Graphs represented as `graphNEL` objects are displayed with `plot()`. There is no `plot()` method for graphs represented as adjacency matrices, so here coercion is one option:

```
R> par(mfrow=c(1,2)); plot(ug11); plot(as(ug11m,"graphNEL"))
```



An alternative for adjacency matrices is the `gplot()` function in the `sna` package:

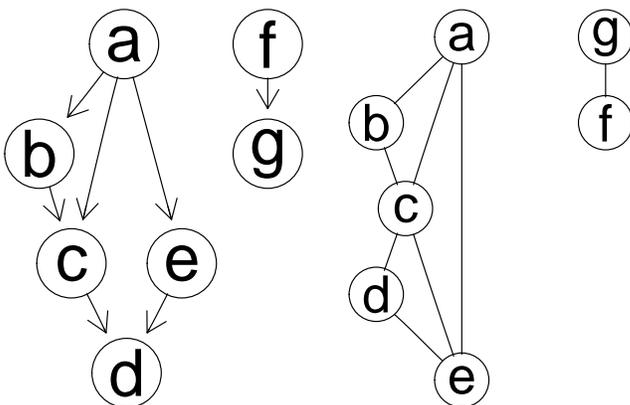
```
R> if (require(sna)){
  par(mfrow=c(1,2))
  gplot(ug11m, label=colnames(ug11m), gmode="graph")
  gplot(dag11m, label=colnames(dag11m))
}
```

3 Advanced graph operations

3.1 Moralization

A moralized directed acyclic graph is obtained with

```
R> dag11.mor <- moralize(dag11)
R> par(mfrow=c(1,2)); plot(dag11); plot(dag11.mor)
```



Default is that the output representation is the same as the input representation, but this can be changed:

```
R> class( moralize( dag11, result="matrix" ) )
[1] "matrix"
R> class( moralize( dag11, result="dgCMatrix" ) )
[1] "dgCMatrix"
attr(,"package")
[1] "Matrix"
```

3.2 Topological sort

A topological ordering of a directed graph is a linear ordering of its vertices such that, for every edge ($u \rightarrow v$), u comes before v in the ordering. A topological ordering is possible if and only if the graph has no directed cycles, that is, if it is a directed acyclic graph (DAG). Any DAG has at least one topological ordering.

```
R> topoSort(dag11)
[1] "a" "b" "c" "e" "f" "d" "g"
R> topoSort(dag11m)
[1] "a" "b" "c" "e" "f" "d" "g"
R> topoSort(dag11M)
[1] "a" "b" "c" "e" "f" "d" "g"
```

The following graph has a cycle and hence a topological ordering can not be made:

```
R> topoSort(dag(~a:b+b:c+c:a))
character(0)
```

Likewise for an undirected graph (recall that we regard an undirected edge as a bidirected edge):

```
R> topoSort( ug( ~a:b ) )
character(0)
```

3.3 Getting the cliques of an undirected graph

In graph theory, a clique is often a complete subset of a graph. A maximal clique is a clique which can not be enlarged. In statistics (and that is the convention we follow here) a clique is usually understood to be a maximal clique. Finding the cliques of a general graph is an NP-complete problem. Finding the cliques of triangulated graph is linear in the number of cliques.

```
R> str( getCliques(ug11) )
List of 5
 $ : chr [1:3] "a" "b" "c"
 $ : chr [1:2] "a" "e"
 $ : chr [1:2] "d" "c"
 $ : chr [1:2] "d" "e"
 $ : chr [1:2] "f" "g"
R> str( getCliques(ug11m) )
List of 5
 $ : chr [1:3] "a" "b" "c"
 $ : chr [1:2] "a" "e"
 $ : chr [1:2] "d" "c"
 $ : chr [1:2] "d" "e"
 $ : chr [1:2] "f" "g"
R> str( getCliques(ug11M) )
List of 5
 $ : chr [1:3] "a" "b" "c"
 $ : chr [1:2] "a" "e"
 $ : chr [1:2] "d" "c"
 $ : chr [1:2] "d" "e"
 $ : chr [1:2] "f" "g"
```

For `graphNEL` objects one may also use the `maxClique()` function in **RBGL**, but `getCliques()` applies also to matrices and it is substantially faster:

```
R> if (require(microbenchmark)){
  microbenchmark(
    RBGL::maxClique( ug11 ), getCliques( ug11 ), getCliques( ug11m ),
    getCliques( ug11M )
  )
}
Unit: microseconds
      expr      min       lq median      uq      max neval
RBGL::maxClique(ug11) 688.55 710.71 727.04 741.97 964.7   100
  getCliques(ug11) 243.05 253.54 260.77 272.90 2007.8  100
  getCliques(ug11m)  72.31  77.91  81.64  86.54  109.2   100
  getCliques(ug11M) 119.42 124.09 128.29 138.32  176.3   100
```

3.4 Perfect ordering, maximum cardinality search and triangulation

An undirected graph is triangulated (or chordal) if it has no cycles of length ≥ 4 without a chord. This is equivalent to that the vertices can be given a perfect ordering. A perfect ordering (if it exists) can be obtained with Maximum Cardinality Search. If `character(0)` is returned the graph is not triangulated. Otherwise a perfect ordering of the nodes is returned.

```
R> mcs(ug11)
character(0)
R> mcs(ug11m)
character(0)
R> mcs(ug11M)
character(0)
R> mcs( dag11.mor )
[1] "a" "b" "c" "e" "d" "g" "f"
R> mcs( as( dag11.mor, "matrix" ))
[1] "a" "b" "c" "e" "d" "g" "f"
R> mcs( as( dag11.mor, "dgCMatrix" ))
[1] "a" "b" "c" "e" "d" "g" "f"
R> mcs( dag11 )
character(0)
```

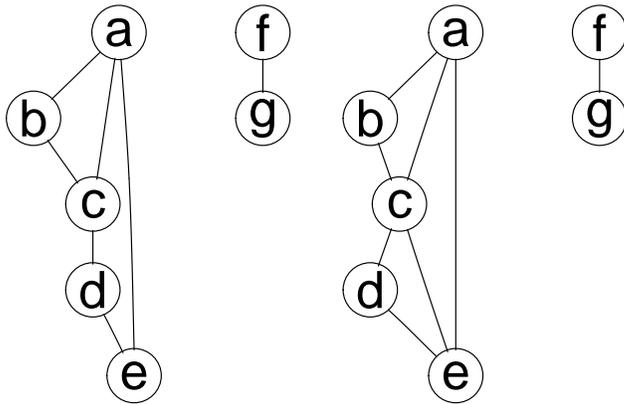
Any undirected graph can be triangulated by adding edges to the graph, so called fill-ins:

```
R> (tug11 <- triangulate(ug11))
A graphNEL graph with undirected edges
Number of Nodes = 7
Number of Edges = 8
R> (tug11m <- triangulate(ug11m))
  a b c d e f g
a 0 1 1 0 1 0 0
b 1 0 1 0 0 0 0
c 1 1 0 1 1 0 0
d 0 0 1 0 1 0 0
e 1 0 1 1 0 0 0
f 0 0 0 0 0 0 1
g 0 0 0 0 0 1 0
R> (tug11M <- triangulate(ug11M))
```

```

7 x 7 sparse Matrix of class "dgCMatrix"
  a b c d e f g
a . 1 1 . 1 . .
b 1 . 1 . . . .
c 1 1 . 1 1 . .
d . . 1 . 1 . .
e 1 . 1 1 . . .
f . . . . . 1
g . . . . . 1 .
R> par(mfrow=c(1,2)); plot(ug11); plot(tug11)

```



Default is that the output representation is the same as the input representation, but this can be changed:

```

R> class( triangulate( ug11, result = "matrix" ) )
[1] "matrix"
R> class( triangulate( ug11, result = "dgCMatrix" ) )
[1] "dgCMatrix"
attr(,"package")
[1] "Matrix"

```

3.5 RIP ordering / junction tree

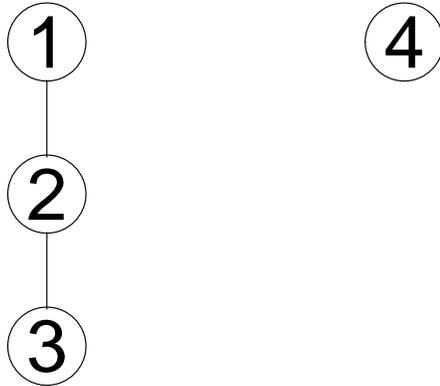
A RIP-ordering of the cliques of a triangulated graph can be obtained as:

```

R> rp <- rip(tug11); rp
cliques
 1 : c a b
 2 : c a e
 3 : c d e
 4 : f g
separators
 1 :
 2 : c a
 3 : c e
 4 :
parents
 1 : 0
 2 : 1
 3 : 2
 4 : 0

```

```
R> plot( rp )
```



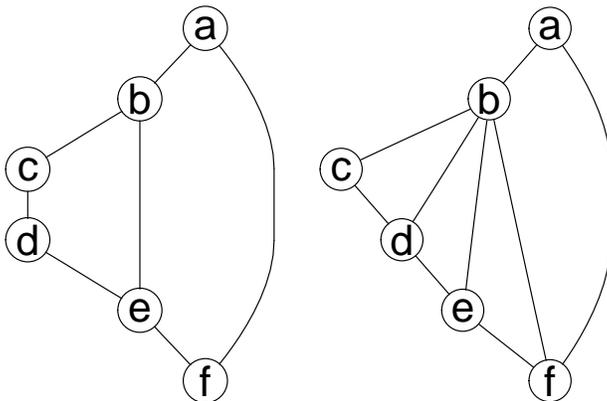
The function `jTree` takes an undirected graph as input; triangulates it if it is not already so and then finds a RIP-ordering.

3.6 Minimal triangulation and maximum prime subgraph decomposition

A triangulation is minimal if no fill-ins can be removed without breaking the property that the graph is triangulated. (A related concept is the minimum triangulation, which is the the graph with the smallest number of fill-ins. The minimum triangulation is unique, but finding the minimum triangulation is NP-hard.)

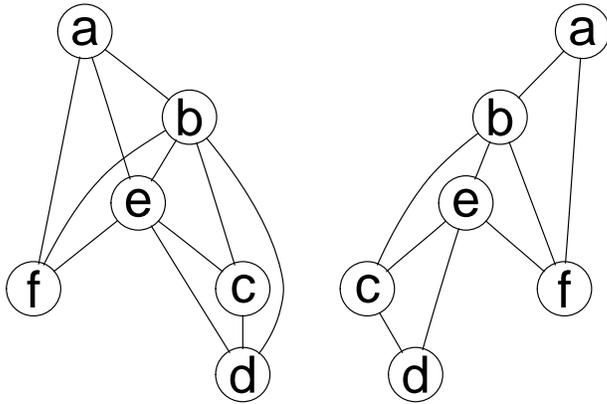
For example, this graph has two 4-cycles:

```
R> g1 <- ug(~a:b+b:c+c:d+d:e+e:f+a:f+b:e)
R> g1mt <- minimalTriang(g1) # A minimal triangulation
R> par(mfrow = c(1,2)); plot(g1); plot(g1mt)
```



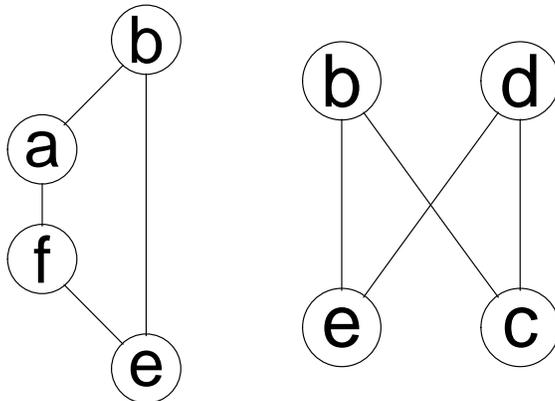
The following graph is also a triangulation of `g1`, and from this a minimal triangulation can be obtained:

```
R> g2 <- ug(~a:b:e:f+b:c:d:e)
R> g1mt2 <- minimalTriang(g1, tobject=g2)
R> par(mfrow = c(1,2)); plot(g2); plot(g1mt2)
```



The junction tree of a maximum prime subgraph decomposition is obtained with:

```
R> mm <- mpd( g1 ); mm
cliques
 1 : b a f e
 2 : b d e c
separators
 1 :
 2 : b e
parents
 1 : 0
 2 : 1
R> par(mfrow = c(1,2))
R> plot(subGraph(mm$cliques[[1]], g1))
R> plot(subGraph(mm$cliques[[2]], g1))
```



4 Time and space considerations

4.1 Time

It is worth noticing that working with graphs represented as `graphNEL` objects can be somewhat slower than working with graphs represented as adjacency matrices. Consider finding the cliques of an undirected graph represented as a `graphNEL` object or as a matrix:

```
R> if(require(microbenchmark)){
  microbenchmark(
    RBGL::maxClique(ug11),
    getCliques(ug11),
    getCliques(ug11m),
    getCliques(ug11M) ) }
Unit: microseconds
      expr      min       lq  median       uq      max  neval
RBGL::maxClique(ug11) 685.75 703.95 712.58 727.5 907.3   100
getCliques(ug11)     238.38 250.04 256.81 266.8 306.5   100
getCliques(ug11m)    73.24 77.91 80.94 83.5 133.4   100
getCliques(ug11M)   118.49 122.92 127.36 133.2 253.3   100
```

4.2 Space

The **graphNEL** representation is – at least – in principle more economic in terms of space requirements than the adjacency matrix representation (because the adjacency matrix representation uses a 0 to represent a “missing edge”). The sparse matrix representation is clearly only superior to the standard matrix representation if the graph is sparse:

```
R> V <- 1:300
R> M <- 1:10
R> ## Sparse graph
R> ##
R> g1 <- randomGraph(V, M, 0.05)
R> length(edgeList(g1))
[1] 896
R> s<-c(NEL=object.size(g1), dense=object.size(as(g1, "matrix")),
       sparse=object.size(as(g1, "dgCMatrix")))
R> s/max(s)
      NEL  dense  sparse
1.00000 0.93593 0.07175
R> ## More dense graph
R> ##
R> g1 <- randomGraph(V, M, 0.5)
R> length(edgeList(g1))
[1] 41872
R> s <- c(NEL=object.size(g1), dense=object.size(as(g1, "matrix")),
         sparse=object.size(as(g1, "dgCMatrix")))
R> s/max(s)
      NEL  dense  sparse
1.00000 0.02385 0.03293
```

5 Graph queries

The **graph** and **RBGL** packages implement various graph operations for **graphNEL** objects. See the documentation for these packages. The **gRbase** implements a few additional functions, see Section 1. An additional function in **gRbase** for graph operations is **querygraph()**. This function is intended as a wrapper for the various graph operations available in **gRbase**, **graph** and **RBGL**. There are two main virtues of **querygraph()**: 1) **querygraph()** operates on any of the three graph representations described above² and 2) **querygraph()** provides a unified interface to the graph operations. The general syntax is

²Actually not quite yet, but it will be so in the future.

```
R> args(querygraph)
function (object, op, set = NULL, set2 = NULL, set3 = NULL)
NULL
```