

- 1 -

A METHOD OF IMPLEMENTING AN ACYCLIC DIRECTED GRAPH STRUCTURE USING A RELATIONAL DATABASE

5 Field of the Invention

This invention relates to a method for implementing an acyclic directed graph structure, or acyclic diagraph, using a relational database. Acyclic diagraph structures are used frequently to model real world hierarchy systems. Examples of such hierarchy systems where an acyclic
10 diagraph structure may be used are family trees, Object-Oriented relationship models, network routing structure, etc.

Background of the Invention

15 Despite the prevalence of acyclic diagraph structures in the real world, few computer applications actually use the model. Application developers prefer instead to either use another similar model, the directed tree structure, to emulate acyclic diagraph systems, or alternatively generate "views" of an acyclic diagraph system and represent the views with a directed-tree structure.

20

The main reason why application programmers tend to avoid modelling the acyclic diagraph structure is due to the complexity of the structure, and the complexity of implementing and maintaining this structure.

25 A graph is a mathematical system consisting of nodes (or points or vertices) and edges (or links or arcs). The following is a mathematical definition of a graph taken from a book by Jean-Paul Tremblay and Paul G Sorenson (ISBN0-07-065157-4):-

30 "A *graph* G consist of a nonempty set V called the set of *nodes (points, vertices)* of the graph, a set E which is the set of edges of the *graph*, and a mapping from the set

of edges E to a set of pairs of elements of V .

We can assume throughout that both sets V and E of a graph are finite. It is also convenient to write a *graph* as $G = (V, E)$. Notice that the definition of a *graph* implies that to every
5 edge of the *graph* G , we can associate a pair of nodes of the *graph*. If an edge, $x \in E$ is thus associated with a pair of nodes (u, v) where $u, v \in V$ then we say that the edge x connects or joins the nodes u and v . Any two nodes which are connected by an edge in a *graph* are called adjacent nodes.

- 10 In a graph $G = (V, E)$ an *edge* which is directed from one node to another is called a *directed edge*, while an *edge* which has no specific direction is called an *undirected edge*. A *graph* in which every edge is directed is called a *directed graph*, or a *diagraph*".

In addition to the above definition, the same book gives the following additional definitions:-

- 15 "Any sequence of edges of a diagraph such that the terminal node of any edge in the sequence is the initial node of the edge, if any, appearing next in the sequence defines a *path* of the graph. A path is said to traverse through the nodes appearing in the sequence originating in the initial node of the first edge and ending in the terminal node of the last edge in the sequence."
- 20 "A path which originates and ends in the same node is called a *cycle (circuit)*."
- "A simple diagraph which does not have any cycles is called *acyclic*."

From the above definition, the properties of an acyclic diagraph can be summarised as:-

1. The structure consists only of nodes and edges.
- 25 2. Edges are uni-directional.
3. The direct relationship between any two nodes is represented by a maximum of one edge.
4. Indirect relationships, or paths, between any two nodes may exist through edges with other nodes.
- 30 5. Each node may have one or more child nodes.

6. Each node may have one or more parent nodes.
7. There may exist ancestor nodes for any given node, which consist of all the successive parents of a particular node.
8. There may exist descendant nodes for any given node, which consist of all the successive children of a particular node.
9. An ancestor of a node cannot be a descendant of the same node.
10. There exists an inverse structure to any given acyclic diagram structure. The inverse structure to any given acyclic diagram structure is the same structure but with all the edge directions reversed.

10

The traditional method to store and manipulate an acyclic directed graph structure in a computer is done using third generation programming languages. Basically this involves using defined memory storage areas as nodes, and pointers as edges. An example using C code is thus:-

```
15      struct Node {  
          char          nodename [10]  
          struct Node    *nextnode[4] }
```

Where nodename would represent the actual node, and nextnode would represent an edge to another node. Operations on the structure could be programmed using functions or procedures. Examples of such functions may be "Insert Node", "Delete Node", "Attach Node", etc.

Although this method can efficiently implement an acyclic diagram, there are several considerations to be made with regard to the complexity and reliability of such an implementation.

1. The maximum number of edges associated with a node has to be known in advance.
2. Difficult to find existence of paths between nodes, as all possible paths have to be traversed before a conclusion can be made.
3. No consideration is made on to the reliability of the structure with regards to

database concurrency and integrity.

4. No node locking mechanism is available, thus limiting the use of the structure to one user or operation at a time.

5. May encounter problems when using third generation languages with regard to memory allocation or lost pointers which would cause the whole program to crash.

Although a skilled programmer may be able to get around all these limitations, addressing each and every one of them only adds further complexity to the programming to be done. In addition, since everything is done programmatically, there is a greater possibility of hidden bugs in the systems, and the system maintenance may thus prove very difficult and time consuming.

Summary of the Invention

15 This invention offers an alternative method of storing and manipulating acyclic directed graph structures by using a relation database.

In accordance with the present invention, there is provided a method for generating an acyclic directed graph structure, or acyclic diagram, utilising a computer operated relational database, the acyclic diagram comprising a plurality of nodes directionally coupled to one another so that each node has at least one child node in a forward coupling direction and/or at least one parent node in a reverse coupling direction, comprising:

generating a node table structure indicating a relationship between each node in the acyclic diagram and at least one node property;

25 generating an edge table structure indicating a relationship between each directly coupled pair of nodes in the acyclic diagram; and

generating a path table structure indicating existence of a path between any two nodes in the acyclic diagram.

30 The present invention also provides a computer stored data structure in a relational database

system for representing an acyclic diagram comprising a plurality of nodes directionally coupled to one another so that each node has at least one child node in a forward coupling direction and/or at least one parent node in a reverse coupling direction, the data structure comprising:

5 a node table structure indicating a relationship between each node in the acyclic diagram and at least one node property, the node table structure including a record for each node in the acyclic diagram, each node table structure record including a node identification field and at least one node property field;

 an edge table structure indicating a relationship between each pair of coupled parent
10 and child nodes in the acyclic diagram, the edge table structure including a record for each pair of parent and child coupled nodes, each edge table structure record including respective fields to identify the parent node and the child node; and

 a path table structure indicating existence of a path between any two nodes in the acyclic diagram, the path table structure including a record for each pair of parent and child
15 coupled nodes and pair of ancestor and descendant coupled nodes wherein two nodes are coupled through an intermediate, each path table structure record including respective fields to identify each of the pair of nodes.

A method of maintaining a data structure as defined above is also provided by the present
20 invention, including:

 a node creation procedure for creating a new node in the acyclic diagram by inserting a new record to the node table structure with a unique value in said node identification field;

 a node linking procedure for linking a child node to a parent node in the acyclic diagram, comprising the steps of:

25 inserting a new record in the edge table structure identifying the child node and the parent node and a coupling direction from the parent node to the child node;

 inserting a new record in the path table structure identifying the child node and the parent node and a coupling direction from the parent node to the child
30 node;

- selecting from the path table structure each ancestor node coupled to the parent node and inserting a new record in the path table structure identifying the child node and each ancestor node and a coupling direction from the ancestor node to the child node; and
- 5 selecting from the path table structure each descendant node to which the child node is coupled and inserting a new record in the path table structure identifying the child node and each descendant node and a coupling direction from the child node to the descendant node;
- a node unlinking procedure for unlinking a child node from a parent node in the
- 10 acyclic diagraph, comprising the steps of:
- removing from the edge table structure the record identifying the child node and the parent node;
- removing from the path table structure each record identifying the child node;
- selecting and removing from the path table structure each record identifying
- 15 an ancestor node of the parent node and a descendant node of the child node; and
- systematically reviewing the edge table structure beginning with records identifying an ancestor node of the parent node to determine any path from an ancestor node of the parent node to a descendant node of the child node,
- 20 reviewing the path table structure to if any such path is identified by a path table structure record and inserting a new record into the path table structure identifying that path if none exists in the path table structure; and
- a remove node procedure for removing a node which has been unlinked from the acyclic diagraph, comprising removing from the node table structure the record identifying
- 25 the unlinked node.

Brief Description of the Drawings

The invention is described in greater detail hereinafter, by way of example only, with

30 reference to the accompanying drawings, wherein:

Figure 1 is a representative diagram of an acyclic directed graph structure; and
Figure 2 is an example of another acyclic diagram structure.

Detailed Description of the Preferred Embodiments

5

Described herein is a method to implement acyclic diagram structures with the use of a commercially available relation database. Various operations are described on how to store, retrieve, update and maintain items in the structure.

10

The advantages of using commercially available databases are:-

1. Currently available in the market place at affordable prices.
2. Speeds up application development time.
3. Utilises the strengths of relational databases. These include the reliability, networking capability, indexing capabilities, etc.

15

4. Relational databases are based on the mathematical concepts of set theory and tuple calculus. The same concepts are used for operations on an acyclic diagram structure.

This preferred form of the invention employs three database tables. The tables are:-

1. Node Table

20

This table stores the properties of each node. A key field must be used to uniquely identify each node. Any number of additional fields can also be used to represent additional properties of the node. Thus each record in this table would represent a node in an acyclic directed tree structure.

25

2. Edge Table

This table stores the relationship between any two nodes. It must consist of at least two fields, a parent field (or initial node field) and a child field (or terminal node field). Thus each record within this table would thus represent an edge in an acyclic directed graph structure.

30

3. Path Table

This table list paths between any two nodes. This table must consist of at least two fields, an ancestor field, which is the initial node of the path, and a descendant field, which is the terminal node in the path. Thus each record within this table would represent a path in the acyclic directed graph structure. Note that this table does not list the routes between any two nodes, that has to be derived from a traverse of the edge table, but only the existence of a path between any two nodes. This is sufficient as it will show that there is a relationship between the two nodes.

10 It is important that the integrity between all three tables be kept. The following rules can be utilised in order to ensure integrity is maintained:-

1. Each value in the parent field, child field, ancestor field and descendant field must be able to be mapped to a key field record in the nodes table.
2. All paths as represented in the Path table must be able to be reproduced through traversing the nodes using the Edge table.
3. All possible paths in the structure which are traceable by traversing the nodes in the Edge table, are represented in the Path table.

If all these three rules are kept, then these three tables will thus represent an acyclic directed graph structure. Note that this method also allows more than one acyclic diagraph structure to exist in the database at any one time, and also allows the different structures to merge together, or split up into two different structures. Also note that this technique also allows us to not only represent an acyclic directed graph structure, but also its inverse structure. An example of such a structure in a relational database is represented by Tables 1, 2 and 3 below, and shown in diagrammatic form in Figure 1.

Node	Property (Colour)
1	Blue
2	Green
3	Blue
4	Red
5	Brown

5

10

6	Green
7	Red
8	Blue
9	Violet
10	Brown
11	Green
12	Blue
13	Violet
14	Blue
15	Red

Table 1: Node Table

15

20

25

30

Parent	Child
1	6
1	4
2	9
2	7
2	5
3	5
3	8
4	9
5	13
6	9
6	10
7	11
7	12
7	13
8	14
9	15
11	15

Table 2: Edge Table

35

40

Ancestor	Descendant
1	4
1	6
1	9
1	10
1	15
2	5
2	7

	2	9
	2	11
	2	12
5	2	13
	2	15
	3	5
	3	8
	3	13
10	3	14
	4	9
	4	15
	5	13
	6	9
15	6	10
	7	11
	7	12
	7	13
	8	14
	9	15
20	11	15

3

Table 3: Path Table

In order to efficiently manipulate the acyclic directed graph structure, several operations need
25 to be defined. The following is a description of the operations as well as sample Microsoft
SQL Server code on how to implement the operations.

Create Node

This operation creates a new node in the structure. A new node in the structure is represented
30 by a new record in the node table. A new key value must also be generated to uniquely
identify this new node. One or more additional properties may also be specified for the new
node. A new node once created, will exist independently from the rest of the structure. It
must be linked to a node already in the structure to become part of it.

Example:

```
INSERT INTO node_table (node, property)
5 VALUES (16, Blue);
```

Remove Node

This operation removes a node from the structure entirely. Note that the node must first be unlinked from all other nodes before it is removed.

10

Link Node

This operation links a node (Child node) to another node (Parent node) in the structure. By linking the two nodes, new paths may also be formed in the structure. Part of this algorithm is to ensure that there is integrity in the structure. The algorithm is:-

- 15 1. Determine whether a cyclic path would be formed by the Link operation.
(This can easily be done by ensuring that the Child node to be linked is not already an ancestor node of the Parent node to be linked to). If a cyclic path would be formed then the procedure is terminated since the resulting structure would not be acyclic.
2. Link the Child node to the Parent node.
- 20 3. Determine whether at least one of the new paths already exist in the structure. (If at least one new path already exists, then all possible new paths should already exist in structure).
4. End operation if paths already exist, else continue on to the next step.
5. Generate a list of all new paths and add corresponding records to the Path
- 25 table. This can be done by:
 - i) Adding all ancestors of the parent node as ancestors of the descendants of the child node into the Path table.
 - ii) Adding all the ancestors of the parent node as ancestors of the child node into the Path table.
 - 30 iii) Adding all the descendants of the child node as descendants of the parent node into the Path table.
 - iv) Adding the Parent node and Child node as ancestor and descendant

respectively into the Path table.

Example of Link Operation

```
5 CREATE PROCEDURE Link @parent_node, @child_node
  AS
  DECLARE    @ancestor int
  DECLARE    @testcount int
  DECLARE    nodeancestor    INSENSITIVE CURSOR
10      FOR  SELECT    DISTINCT ancestor
            FROM path_table
            WHERE descendant = @parent_node

  BEGIN
15      SELECT    @testcount = COUNT (ancestor)
            FROM    path_table
            WHERE descendant = @parent_node
            AND      ancestor = @child_node

20      IF @testcount > 0
          BEGIN
              RAISERROR ('Illegal Link, Cyclic path not allowed',16,-1)
              RETURN
          END

25      INSERT INTO    edge_table (parent, child)
            VALUES    @parent_node, @child_node

            SELECT    @testcount = COUNT (ancestor)
30      FROM    path_table
            WHERE    ancestor = @parent_node
            AND      descendant = @child_node

            IF @testcount > 0
35      BEGIN
          RETURN    /* Path already exist in structure */
        END
```

```
OPEN  nodeancestor
FETCH nodeancestor INTO @ancestor
WHILE @@fetch_status = 0
BEGIN
5      INSERT INTO      path_table (ancestor, descendant)
      SELECT            @ancestor, descendant
      FROM              path_table
      WHERE             ancestor = @child_node
      AND               descendant NOT IN
10      (               SELECT      descendant
                        FROM        path_table
                        WHERE ancestor = @ancestor)

      INSERT INTO      path_table (ancestor, descendant)
15      SELECT DISTINCT parent, @child_node
      FROM            edge_table
      WHERE           parent = @ancestor
      AND             child_node NOT IN
                    (      SELECT      descendant
20                        FROM        path_table
                        WHERE ancestor = @ancestor)

      FETCH nodeancestor INTO @ancestor
END

25      INSERT INTO      path_table (ancestor, descendant)
      SELECT            @parent_node, descendant
      FROM              path_table
      WHERE             ancestor = @child_node
      AND               descendant NOT IN
30      (               SELECT      descendant
                        FROM        path_table
                        WHERE ancestor = @parent_node)

      INSERT INTO      path_table (ancestor, descendant)
35      SELECT DISTINCT parent, @child_node
      FROM            edge_table
      WHERE           parent = @parent_node
      AND             @child_node NOT IN
                    (      SELECT      descendant
```

```
FROM      path_table
WHERE ancestor = @parent_node)
```

END

5

Unlink Node

This operation unlinks a node (Child node) from another node (Parent node) in the structure. By unlinking any two nodes, the paths table needs to be updated to reflect that some of the paths are no longer valid. The algorithm for the unlink operation can be expressed as:

- 10 1. Remove from the edge table the link between the Parent node and Child node.
Link the Child node to the Parent node.
2. Remove from the path table all links between ancestors of the parent node and descendants of the child node.
3. Remove from the path table the link between the parent node and child node.
- 15 4. Add back to the path table all links which may exist between the ancestors of the parent node and the descendants of the child node through alternative paths.
5. Add back to the path table the link between the parent node and child node if such a path does exist through an alternative route.

20 Example of Unlink Operation

```
CREATE PROCEDURE unlink @parent_node, @child_node
AS
DECLARE      @ancestor int
25 DECLARE   @testcount int
DECLARE      nodeancestor SCROLL CURSOR
FOR SELECT   DISTINCT path_table1.ancestor
FROM path_table as path_table1
INNER JOIN path_table as path_table2
30          ON path_table1.ancestor = path_table2.ancestor
WHERE       path_table1.descendant = @child_node
AND         (path_table2.descendant = @parent_node
OR path_table1.ancestor = @parent_node)

35 BEGIN
```

```
DELETE FROM      edge_table
WHERE            parent = @parent_node
AND             child = @child_node

5  OPEN          nodeancestor
   FETCH FIRST nodeancestor INTO @ancestor
   WHILE        @@fetch_status = 0
   BEGIN
       DELETE FROM      path_table
10      WHERE          ancestor = @ancestor
       AND             (descendant IN
                        (   SELECT      descendant
                          FROM path_table
                          WHERE        ancestor = @child_node)
                        OR descendant = @class_node)
       FETCH NEXT nodeancestor INTO @ancestor
   END

   FETCH FIRST nodeancestor INTO @ancestor
20  WHILE        @@fetch_status = 0
   BEGIN
       INSERT INTO      path_table (ancestor, descendant)
       SELECT           @ancestor, child
       FROM             edge_table
25      INNER JOIN      path_table ON parent = descendant
       WHERE            ancestor = @ancestor
       AND              child NOT IN (   SELECT descendant
                                         FROM path_table
                                         WHERE ancestor = @ancestor)

30      INSERT INTO      path_table (ancestor, descendant)
       SELECT           @ancestor, child
       FROM             edge_table
       WHERE            parent = @ancestor
       AND              child NOT IN (   SELECT descendant
                                         FROM path_table
35                                         WHERE ancestor = @ancestor)

       FETCH NEXT FROM nodeancestor INTO @ancestor
   END
END
```

An example of how this technique may be applied is be in a Classification System. A Business Directory may have several hundred products and services classifications. These classifications may be organised into a directed-graph structure. An illustration of a part of a sample classification structure is shown in Figure 2.

5

Using the described technique for representing this classification system, the following tables are developed.

10

Node	Classification
1	Flow Control Equipment
2	Valves
3	Pumps
4	Waste Treatment Equipment
5	Air Pumps
6	Acid Pumps
7	Sewerage Pumps
8	Sewerage Filters

15

Table 4: Node Table

20

Parent	Child
1	2
1	3
3	5
3	6
3	7
4	7
4	8

25

30

Table 5: Edge Table

5

10

Ancestor	Descendant
1	2
1	3
1	5
1	6
1	7
3	5
3	6
3	7
4	7
4	8

Table 6: Path Table

15

It is then possible to relate companies to any given classification, and in doing so, make sure that all ancestor classifications may also 'inherit' this relationship. For example, supposing it is desired to relate the company "Caterpillar Pumps Ltd" to "Air Pumps" by inserting a record into a table specially created to represent this relationship. The following table would
20 result:-

Company Classification Table

25

Classification Node	Company
5	Caterpillar Pumps Ltd

The Path Table can then be used to reflect the same relationship to the ancestor classifications. A sample SQL statement to do this would be:-

30

```
INSERT INTO company_classification_table
SELECT ancestor, "Caterpillar Pumps Ltd"
FROM path_table
WHERE descendant = 5
```

35 The resultant table would then be:-

Classification Node	Company
1	Caterpillar Pumps Ltd
3	Caterpillar Pumps Ltd
5	Caterpillar Pumps Ltd

If further companies were then added to the system, say "Calpeda Valves Ltd" to the classification "Valves" and "Green Treatment Ltd" to "Sewerage Pumps", then the following table would result:-

5

Classification Node	Company
1	Caterpillar Pumps Ltd
3	Caterpillar Pumps Ltd
5	Caterpillar Pumps Ltd
1	Calpeda Valves Ltd
2	Calpeda Valves Ltd
1	Green Treatment Ltd
3	Green Treatment Ltd
4	Green Treatment Ltd
7	Green Treatment Ltd

The same table, but now sorted by classifications would thus look like this:-

Classification Node	Company
1	Caterpillar Pumps Ltd
1	Calpeda Valves Ltd
1	Green Treatment Ltd
2	Calpeda Valves Ltd
3	Caterpillar Pumps Ltd
3	Green Treatment Ltd
4	Green Treatment Ltd
5	Caterpillar Pumps Ltd
7	Green Treatment Ltd

Thus it is possible to quickly maintain a proper "inheritance" structure without needing to find the root node, nor needing to traverse any paths.

The foregoing detailed description of the preferred implementation of the invention has
5 been presented by example only, and is not intended to be considered limiting to the invention as defined in the claims appended hereto.



Claims:

1. A method for generating an acyclic directed graph structure, or acyclic diagram, utilising a computer operated relational database, the acyclic diagram comprising a plurality
5 of nodes directionally coupled to one another so that each node has at least one child node in a forward coupling direction and/or at least one parent node in a reverse coupling direction, comprising:
 - generating a node table structure indicating a relationship between each node in the acyclic diagram and at least one node property;
 - 10 generating an edge table structure indicating a relationship between each directly coupled pair of nodes in the acyclic diagram; and
 - generating a path table structure indicating existence of a path between any two nodes in the acyclic diagram.
- 15 2. A method as claimed in claim 1, wherein the node table structure includes a record for each node in the acyclic diagram, each node table structure record including a node identification field and at least one node property field.
3. A method as claimed in claim 1, wherein the edge table structure includes a record for
20 each directly coupled pair of nodes, each edge table structure record including respective fields to identify each of the pair of nodes.
4. A method as claimed in claim 3, wherein the order of the fields in each edge table structure record indicates a direction of coupling of the nodes identified in the record.
25
5. A method as claimed in claim 1, wherein the path table structure includes a record for each directly coupled pair of nodes and each indirectly coupled pair of nodes, each path table structure record including respective fields to identify each of the pair of nodes.
- 30 6. A method as claimed in claim 5, wherein the data contained in the edge table structure



comprises a sub-set of the data contained in the path table structure.

7. A method as claimed in claim 5, wherein the order of the fields in each path table structure record indicates a direction of coupling of the nodes identified in the record.

5

8. A method as claimed in claim 1, wherein:

the node table structure includes a record for each node in the acyclic diagram, each node table structure record including a node identification field and at least one node property field;

10 the edge table structure includes a record for each directly coupled pair of nodes, each edge table structure record including respective fields to identify each of the pair of nodes, the order of the fields in each edge table structure record indicating a direction of coupling of the nodes identified in the record; and

the path table structure includes a record for each directly coupled pair of nodes and
15 each indirectly coupled pair of nodes, each path table structure record including respective fields to identify each of the pair of nodes, and the order of the fields in each path table structure record indicating a direction of coupling of the nodes identified in the record.

9 A method as claimed in claim 8, including a node creation procedure for creating a
20 new node in the acyclic diagram by inserting a new record to the node table structure with a unique value in said node identification field.

10. A method as claimed in claim 9, including a node linking procedure for linking the new node as a child node to a parent node in the acyclic diagram, comprising the steps of:

25 inserting a new record in the edge table structure identifying the child node and the parent node and a coupling direction from the parent node to the child node;

inserting a new record in the path table structure identifying the child node and the parent node and a coupling direction from the parent node to the child node;

selecting from the path table structure each ancestor node coupled to the parent node
30 and inserting a new record in the path table structure identifying the child node and each



ancestor node and a coupling direction from the ancestor node to the child node.

11. A method as claimed in claim 8, including a node linking procedure for linking a
5 child node to a parent node in the acyclic diagram, comprising the steps of:

inserting a new record in the edge table structure identifying the child node and the
parent node and a coupling direction from the parent node to the child node;

inserting a new record in the path table structure identifying the child node and the
parent node and a coupling direction from the parent node to the child node;

10 selecting from the path table structure each ancestor node coupled to the parent
node and inserting a new record in the path table structure identifying the child node and
each ancestor node and a coupling direction from the ancestor node to the child node; and

selecting from the path table structure each descendant node to which the child
node is coupled and inserting a new record in the path table structure identifying the parent
15 node and each descendant node and a coupling direction from the parent node to the
descendant node.

12. A method as claimed in claim 8, including a node unlinking procedure for
unlinking a child node from a parent node in the acyclic diagram, comprising the steps of:

20 removing from the edge table structure the record identifying the child node and
the parent node;

removing from the path table structure each record identifying the child node;

selecting and removing from the path table structure each record identifying an
ancestor node of the parent node and a descendant node of the child node; and

25 systematically reviewing the edge table structure beginning with records identifying
an ancestor node of the parent node to determine the existence of any other path from an
ancestor node of the parent node to a descendant node of the child node, reviewing the
path table structure to if any such path is identified by a path table structure record and
inserting a new record into the path table structure identifying that path if none exists in
30 the path table structure.

13. A method as claimed in claim 12, including a remove node procedure for removing
a node which has been unlinked from the acyclic diagram, comprising removing from the



node table structure the record identifying the unlinked node.

14. A computer stored data structure in a relational database system for representing an acyclic diagram comprising a plurality of nodes directionally coupled to one another so that
5 each node has at least one child node in a forward coupling direction and/or at least one parent node in a reverse coupling direction, the data structure comprising:

a node table structure storing a relationship between each node in the acyclic diagram and at least one node property, the node table structure including a record for each node in the acyclic diagram, each node table structure record including a node identification field and
10 at least one node property field;

an edge table structure storing a relationship between each pair of coupled parent and child nodes in the acyclic diagram, the edge table structure including a record for each pair of parent and child coupled nodes, each edge table structure record including respective fields to identify the parent node and the child node; and

15 a path table structure storing representation of a path between any two nodes in the acyclic diagram, the path table structure including a record for each pair of parent and child coupled nodes and pair of ancestor and descendant coupled nodes wherein two nodes are coupled through an intermediate, each path table structure record including respective fields to identify each of the pair of nodes.

20

15. A product and/or service classification system comprising a computer stored data structure as defined in claim 14, wherein the at least one node property in said node table structure represents a particular product or service associated with the corresponding node of the acyclic diagram.

25

16. A product and/or service classification system as claimed in claim 15, including a classification table structure for storing a relationship between a supplier or customer and a node in the acyclic diagram, the classification table structure having at least one record including a node identification field for identifying a said node in the acyclic diagram and a
30 company identification field for identifying a corresponding supplier or customer.



17. A product and/or service classification system as claimed in claim 16, wherein the classification table structure includes a record having a corresponding company identification field for each ancestor node of any node identified in the classification table, according to the path table structure.

5

18. A method of maintaining a data structure as defined in claim 14, including:

a node creation procedure for creating a new node in the acyclic diagram by inserting a new record to the node table structure with a unique value in said node identification field;

a node linking procedure for linking a child node to a parent node in the acyclic
10 diagram, comprising the steps of:

inserting a new record in the edge table structure identifying the child node and the parent node and a coupling direction from the parent node to the child node;

inserting a new record in the path table structure identifying the child node and the parent node and a coupling direction from the parent node to the child
15 node;

selecting from the path table structure each ancestor node coupled to the parent node and inserting a new record in the path table structure identifying the child node and each ancestor node and a coupling direction from the ancestor node to the child node; and
20

selecting from the path table structure each descendant node to which the child node is coupled and inserting a new record in the path table structure identifying the child node and each descendant node and a coupling direction from the child node to the descendant node;

25 a node unlinking procedure for unlinking a child node from a parent node in the acyclic diagram, comprising the steps of:

removing from the edge table structure the record identifying the child node and the parent node;

removing from the path table structure each record identifying the child node;

30 selecting and removing from the path table structure each record identifying



- 25 -

an ancestor node of the parent node and a descendant node of the child node;
and

5 systematically reviewing the edge table structure beginning with records
identifying an ancestor node of the parent node to determine any path from an
ancestor node of the parent node to a descendant node of the child node,
reviewing the path table structure to if any such path is identified by a path
table structure record and inserting a new record into the path table structure
identifying that path if none exists in the path table structure; and

a remove node procedure for removing a node which has been unlinked from the
10 acyclic diagraph, comprising removing from the node table structure the record identifying
the unlinked node.

19. A method of maintaining a product and/or service classification system as claimed in
any one of claims 15 to 17, including:

15 a node creation procedure for creating a new node in the acyclic diagraph by inserting
a new record to the node table structure with a unique value in said node identification field
and a product or service identification in said node property field;

a node linking procedure for linking a child node to a parent node in the acyclic
diagraph, comprising the steps of:

20 inserting a new record in the edge table structure identifying the child node
and the parent node and a coupling direction from the parent node to the child
node;

inserting a new record in the path table structure identifying the child node and
the parent node and a coupling direction from the parent node to the child
25 node;

selecting from the path table structure each ancestor node coupled to the parent
node and inserting a new record in the path table structure identifying the child
node and each ancestor node and a coupling direction from the ancestor node
to the child node; and

30 selecting from the path table structure each descendant node to which the child



node is coupled and inserting a new record in the path table structure identifying the child node and each descendant node and a coupling direction from the child node to the descendant node;

a node unlinking procedure for unlinking a child node from a parent node in the
5 acyclic diagram, comprising the steps of:

removing from the edge table structure the record identifying the child node and the parent node;

removing from the path table structure each record identifying the child node;

selecting and removing from the path table structure each record identifying
10 an ancestor node of the parent node and a descendant node of the child node;
and

systematically reviewing the edge table structure beginning with records identifying an ancestor node of the parent node to determine any path from an ancestor node of the parent node to a descendant node of the child node,

15 reviewing the path table structure to if any such path is identified by a path table structure record and inserting a new record into the path table structure identifying that path if none exists in the path table structure; and

a remove node procedure for removing a node which has been unlinked from the acyclic diagram, comprising removing from the node table structure the record identifying
20 the unlinked node.

20. A method of maintaining a product and/or service classification system as claimed in claim 19, including a company classification procedure for generating a company classification table structure by:

25 specifying at least one supplier or customer company associated with at least one node in the acyclic diagram;

systematically reviewing the path table structure to determine each ancestor node to said at least one node associated with the specified company; and

generating a company classification table structure having a record identifying each
30 determined ancestor node and the at least one specified company associated therewith.



- 27 -

21. A method of maintaining a product and/or service classification system as claimed in claim 20, wherein the company classification table includes a stored correspondence between the at least one specified company and the products and/or services corresponding to each determined ancestor node from node table structure.

5