N-Dimensional Reasoning Using a Tesseral Representation

Frans Coenen, Bernard (Diz) Diaz, Michael Shave,
Bridget Beattie and Trevor Bench-Capon
Department of Computer Science,
The University of Liverpool,
Chadwick Building, P.O. Box 147,
Liverpool L69 3BX, England.
Tel: 0151 794 3698
Fax: 0151 794 3715
email: frans@csc.liv.ac.uk

A multi-dimensional quantitative spatial reasoning system founded on a tesseral representation of space is described. This representation allows space to be linearised so that the same technique is applicable to problems in any number of dimensions. The principal advantages offered by the system are that of general applicability and computational effectiveness. With respect to Cartesian and other tesseral representations it provides for efficient translation and rotation of space using standard arithmetic operations, effective data storage using "block" encoding and fast comparison of groups of addresses. The system has been successfully targeted at a number of application areas, but particularly Environmental Impact Assessment (EIA) and Geographic Information Systems (GIS).

Keywords: N-dimensional Spatio-Temporal Reasoning, Tesseral Representations.

1. Introduction

A universally applicable spatio-temporal reasoning system, the SPARTA (SPAtial Reasoning using Tesseral Addressing) system, The system is founded on a dedicated tesseral representation of space and uses a heuristically guided constraint satisfaction mechanism to solve N-dimensional spatial problems. Using the system N-dimensional problems are defined in terms of an object and constraint scripting language founded on a PROLOG predicate style of syntax (Coenen et al. 1996a). The resulting script is then passed to the SPARTA system which processes the constraints using a heuristically guided depth first constraint satisfaction strategy. The result (assuming that a solution exists) is one or more configurations of all objects, defined in the script, that satisfy the given constraints. The resulting numerically referenced configuration(s) may then be output in some format compatible with the application. Examples include two- and three- dimensional images, application specific file formats such as PBM or DXF (Drawing eXchange Format), or encodings such as HTML.

The resulting system has been found to have a great many applications and has been successfully targeted at a number areas, particularly Geographic Information Systems (GIS) (see Coenen et al. 1994, 1996b and Beattie et al. 1995 for further detail) and Environmental Impact Assessment (EIA) (Beattie et al. 1996). In addition the system has been used to address classic AI problems, such as shape fitting problems (Coenen et al. 1997), and also a number of further applications which do not generally fall under the remit of spatial reasoning, for example timetabling and scheduling (Coenen et al. 1995). This paper is illustrated using the well known N-queens AI scenario, presented as a spatial problem.

2. Object space

Given a particular application, all spatial entities of interest are considered to exist in some N-dimensional orthogonal parallelepiped space referred to as the *object space*. This space is tesselated (sub-divided) into a set of isohedral (same shape) N-dimensional cubic subspaces down to some desired resolution. Each subspace is then allocated a unique integer address (or reference) that reflects its location within the space with respect to some arbitrarily selected corner origin. With respect to this paper we will assume that the maximum number of dimensions to be considered is four (e.g. three spatial and one temporal) and that addresses are represented as 32bit signed integers whose bit pattern is defined as follows:

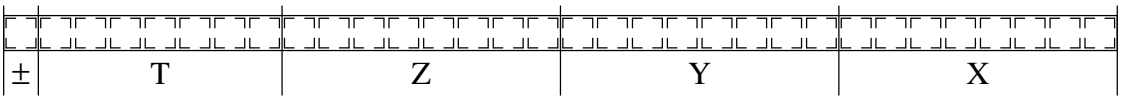

Figure 1: Address bit pattern

Of course any other bit pattern may be selected if it is more appropriate to a particular application. Similarly the number of dimensions to be considered can be increased or decreased, or longer or shorter integers may be used (for example 64 bit or 16 bit integers). Further, no particular relevance need be attached to the labelling of the bit groupings in Figure 1; the X, Y, Z, T labelling has simply been used here to facilitate understanding.

In Figure 2 an example object space is given defined by two dimensions, `X` and `Y`, such that `X = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9}` and `Y = {0, 1, 2, 3, 4}`. Note that each address is made up of two coordinates using the bit pattern presented in Figure 1. Note also that, as a consequence of the addressing mechanism, the space is linearised commencing at the 0 bottom-left corner address (the *origin*) to the top-right address.

The advantages offered by the tesseral addressing mechanism, particularly with respect to the Cartesian system, are as follows:

1) Any cell can be addressed using a single integer. Consequently the computer storage requirements using tesseral addressing are both low and constant.
2) Translation through the space can be achieved through simple integer addition and subtraction. For example, in Figure 2, to move one tile to the "north-east" we always add the address 257 regardless of where we are in the space (to move to the south-west we add the fictitious address -257, and so on). The approach thus avoids recourse to computationally expensive incrementation and decrementation of individual coordinates or cumbersome trigonometric techniques.
3) The N-dimensional space in question is linearised in a broad bottom-left to top-right manner. Knowledge of this linearisation can be utilised as follows:
    - Block encoding of addresses: Groups of addresses contained in a single orthogonal parallelogram can be stored in terms of two addresses representing the corners closest to and furthest away from the origin. Consequently further savings in computer storage requirements are obtained. (We refer to this as *block encoding* so as to relate it to the more usual *run-line encoding*.)

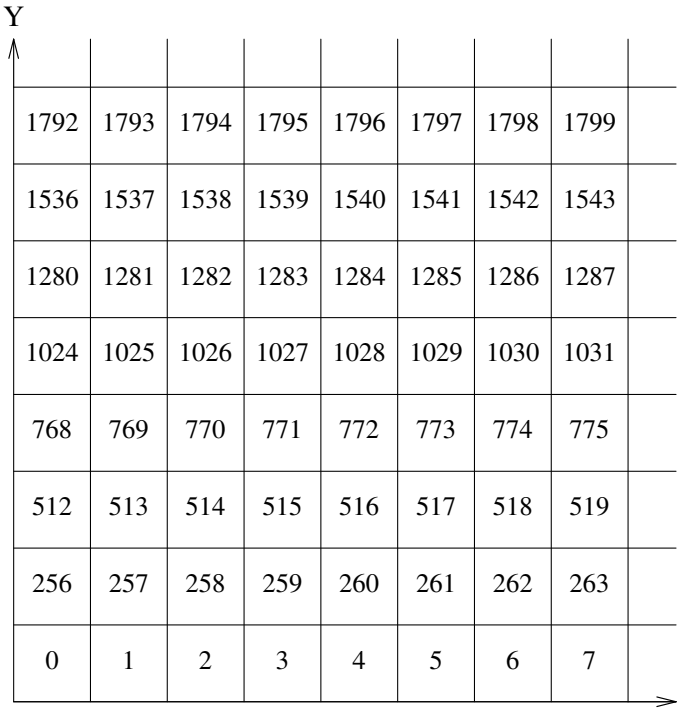

Figure 2: Tesselated 2-D space

- Simple comparison of sets of addresses: To determine whether one set of addresses stands in some relation to another set of addresses it is not necessary to compare each address in one with each address in the other. Instead, knowledge of the relative location within the linearisation can be used to achieve efficient comparison of sets of addresses.
- Simple addition (and subtraction) of blocks of addresses: When adding (or subtracting) one block of addresses to another it is only necessary to add (or subtract) the two pairs of corner addresses.

An additional advantage of the approach is that it is fully compatible with established Raster representations found in GIS (Geographic Information Systems), image encoding and other computer map representations. For further discussion relating to tesseral addressing and its merits interested readers are referred to Diaz and Bell (1986) or Samet (1984).

3. Object classes

By definition the objects we wish to reason about are located within the object space as defined above. Three types of object are identified:

Fixed objects:
Objects whose location is fixed within the object space. Consequently their nature (e.g. shape, size, contiguity etc.) is known.

Free objects:
Objects whose shape is known but not their location.

Shapeless objects:

An object of which nothing is known (other than it exists somewhere in the object space).

Using the scripting language, objects are defined in terms of classes and instances of those classes. A class, identified by a class name, defines the object's type, and (in the case of a free object) its shape. The shape is defined in terms of a set of addresses which must incorporate the zero origin address (negative addresses may be used). Note that a fixed object's shape is defined by its location and therefore fixed object class declarations need not include a shape definition (locations are defined as part of instance declarations as described in Section 4 below). Some high level syntax used for the definition of classes is presented in Table 1.

```
<CLASSDECS>   : <CLASSDEC>
              | <CLASSDEC> <CLASSDECS>
              ;

<CLASSDEC>    : 'class(' CLASSNAME ', fixed).'
              | 'class(' CLASSNAME ', free,' <SHAPE> ').'
              | 'class(' CLASSNAME ', shapeless).'
              ;

<SHAPE>       : 'shape([' <SETTA> '])' ;
```

Table 1: *Class declaration syntax*

4. Object instances

When declaring instances of classes, details of attributes that may be associated with spatial objects may be included according to the type of the object - fixed, free or shapeless. The most significant is the object's location or its location space. This is again defined in terms of a set of addresses. The distinction between locations and location spaces is that the first is associated with fixed objects, and describes a sub-space within the object space where a fixed object **is** located; while the second is associated with free or shapeless objects and describes the sub-space within the object space within which such objects **may** be located.

4.1. Free Object Instances

With respect to free objects there exists a special relationship between the shape of such objects and the associated location space. The number of possible locations for the object is defined (by default) as the number of different ways the entire shape can be fitted into the location space without rotation. Which of these locations is acceptable is then determined by the nature of the constraints associated with the spatial problem, which is in turn dependent on the nature of the application. Note that if no locations for a free object can be generated the object cannot be physically realised and consequently it can be said that no solution to the given spatial problem exists. Given a particular application it is possible to influence the generation of locations as follows:

1) Sub-shape: We may include (in the instance declaration) some tesseral sub-shape definition which is a sub-set of the shape defined in the associated class, and then determine candidate locations for the free object by attempting to fit only this sub-shape into the given location space.
2) Rotation: Given a particular application it may be appropriate to permit the shape to be rotated so that additional locations may be generated.
3) Inclusion: We may wish to insist that certain fixed groups of addresses form part of any candidate location for a fixed object. Typically such inclusion groups of addresses are associated with the locations of a fixed object.
4) Exclusion: We may wish to exclude certain groups of addresses form part of any candidate location.
5) Contiguity: Where we have included and/or excluded some set of addresses we may wish to insist that the result is only acceptable if the location set of addresses is in some way contiguous. In addition we may require it to be contiguous on a particular set of addresses.
6) Size: Similarly, where we have included and/or excluded some set of addresses, in may also be the case that we wish to specify some size qualification whereby the resulting locations are only acceptable if they are less than, equal to or greater than some given size (number of addresses).

4.2. Shapeless Object Instances

Shapeless objects have no known location or shape, only a location space in which they can be said to exist. Where such objects appear in a spatial problem scenario the aim is usually to refine, through an appropriately specified set of constraint *mappings* (see below), the location space associated with such objects. Although no knowledge is available concerning the shape, certain other attributes may be known which can be used to validate the result of any refinement of the location space that might take place as a consequence of the satisfaction of mapping constraints. Currently the following may be specified with respect to shapeless objects:

1) Size: For any resulting location/shape to be acceptable it may be required to be less than, greater than or equal to some magnitude (number of addresses).
2) Contiguity: For a result to be acceptable it may be necessary that the set of addresses describing an ultimate location is in some sense contiguous. In addition there may be a requirement that the contiguity holds over a particular set of addresses.

Using the scripting language instances are defined using syntax of the form indicated in Table 2. Note that, from the above, not all attributes are appropriate to all object types (this is not evident from the simplified syntax fragment given in Table 2).

5. Constraints

Given a spatial application, and using the proposed system, the relationships that are desired to exist between objects are expressed in terms of constraints on the possible locations for those objects. Constraints comprise two operands and an infix operator linking the two operands. Each operand, in turn, comprises one or more object locations

```
<INSTANCEDECS>   : <INSTANCEDEC>
                 | <INSTANCEDEC> <INSTANCEDECS>
                 ;

<INSTANCEDEC>    : 'instance(' <ID> `,' CLASSNAME `).' ;
                 | 'instance(' <ID> `,' CLASSNAME `,' <LOCATION> `).'
                 | 'instance(' <ID> `,' CLASSNAME `,' <LOCATION> `,'
                               <ATTRIBUTES> `).'
                 ;

<ATTRIBUTES>     : <ATTRIBUTE>
                 | <ATTRIBUTE> `,'<ATTRIBUTES>
                 ;

<ATTRIBUTE>      : `subShape(' <SETTA> `)'
                 | `rotation()'
                 | `include(' <ID> `)'
                 | `exclude(' <ID> `)'
                 | `size(' <OPERATOR> N `)'
                 | `contiguous(' `)'

<ID>             : OBJECT
                 | <SETTA>
                 ;

<LOCATION>       : `[' <SETTA> ']' ;
```

Table 2: Instance declaration syntax

(typically identified by an object name). Satisfaction of a constraint will result in one or more pairs of compatible object locations according to the nature of the constraint. For example given two objects, `X` and `Y`, each of which has three locations associated with it, then there are nine possible pairs of locations of which one or more may satisfy the constraint. Thus, as a result of satisfying a constraint, the maximum number of solutions that may result is equivalent to the size of the Cartesian product of the two groups of locations. In the following section, where the operation of the system is discussed, the efficiency concerns associated with the generation of a large number of solutions are addressed. Note that if no pair of locations satisfies the constraint the constraint is considered to have failed and consequently no solution can be generated for the given problem in its current form.

Three categories of constraint are supported:

1) Filters: Constraints that consider each pair of locations associated with a constraint and only "pass" those that satisfy the constraint operator. Such constraints can only operate between fixed and/or free objects where definite locations for objects either exist or can be generated. Currently the system supports five filters: `equals`, `subset`, `superset`, `disjoint` and `intersects`. The expressiveness of the last can be increased by specifying the nature of the intersection to be less than, greater than or equal to some magnitude.
2) Mappings: Constraints that operate between shapeless objects and fixed or free objects and serve to prune the location space associated with the shapeless object by mapping some operation onto the location space with respect to the location(s) associated with the given fixed or free objects. The system supports two mapping operators

`complement` and `intersects`. Again in the case of the intersection operator the nature of the intersection can be specified. On satisfaction of a mapping constraint there may be contiguity and/or size requirements associated with the referenced shapeless object in which case the resulting space must be validated with respect to these requirements before satisfaction of the constraint can be said to be complete.

3) Group filters: Similar to (1) except that the prefix or postfix operand comprises locations associated with more than one object.

Using constraints of the above form many of the standard temporal (e.g. Allen 1983, Dechter et al. 1991, Ladkin 1992) and topological (e.g. Egenhofer 1994, Cohn et al. 1995, Hernández 1991) relations encountered in the spatio-temporal literature can be specified. The expressiveness of the above constraint system can be significantly increased by incorporating the notion of offsets. Offsets are sets of addresses that are applied to locations, associated with the operands for a constraint, prior to attempting to satisfy that constraint. In this manner we can identify spaces associated with locations (for example we can identify (say) the area to the south-west of some 2-D location) or we can expand locations either uniformly or in some direction. Offsets may be applied to all the addresses associated with a location or only to the *origin* for the shape definition. The latter is defined as the corner address, of the minimum surrounding parallelepiped for the shape, nearest the origin of the object space.

Although the expressiveness of constraints is greatly increased using offsets, they cannot easily be used to identify shapes defined by uniformly "shrinking" a location. To this end a *shrink* option is also provided. In addition offsets may be further influenced by insisting that certain groups of addresses are included or excluded.

Using the scripting language constraints are specified as indicated in Table 3.

```
<CONSTRAINTDEC>      : <CONSTRAINT>
                     | <CONSTRAINT> <CONSTRAINTDEC>
                     ;

<CONSTRAINT>         : 'constraint(' <OPERAND> ',' OPERATOR ','
                                   <OPERAND> ').' ;

<OPERAND>            : <OBJECT>
                     | <OBJECT> ',' <OFFSETS>
                     | <SETTA>
                     | <SETTA> ',' <OFFSETS>
                     ;

<OFFSETS>            : <OFFSET>
                     | <OFFSET> <OFFSET_ATTRIBUTE>
                     ;

<OFFSET_ATTRIBUTE>   : 'include(' <ID> ')'
                     | 'exclude(' <ID> ')'
                     ;
```

Table 3: Constraint declaration syntax.

6. Operation

Completed scripts are passed to the reasoning system. This incorporates a lexical analyser and a parser which translates the script into a set of object and constraint data structures. The constraints are then used to verify locations for existing objects, discriminate between candidate locations for free objects and to "prune" the location space associated with shapeless objects. The result is stored in a solution tree which is constructed and destructed dynamically as the solution process progresses. Nodes in the tree represent branches in the solution process where a constraint can be satisfied in more than one manner. This typically occurs when a constraint incorporates a free object that has more than one candidate location associated with it. Wherever a node occurs in the tree the solution process continues in a depth first manner. If a given spatial problem has only one solution the resulting solution tree will comprise only a single node.

The aim of the constraint satisfaction strategy used is to limit the growth of the solution tree. This is achieved through:

1) A constraint selection process which, on each iteration, seeks to minimise the number of candidate solutions that may result on satisfaction of the constraint.
2) A merging mechanism which, should more than one solution be generated, attempts to identify solutions that can co-exist within the restrictions of the given constraint and consequently reduce the number of branches that might otherwise be generated.

The constraint selection process commences by attempting to identify constraints that reference a fixed object or a free object which has only one candidate location associated with it. These constraints will, after appropriate merging of solutions, produce only a single result and will therefore not generate branches in the solution tree. During the processing of these "single solution" constraints it is likely that the number of candidate locations associated with affected free objects will be decreased and a significant amount of pruning of location spaces associated with free objects will have been undertaken. In addition, given a particular scenario, single solution constraints are the most critical - if such constraints cannot be satisfied no solution to the problem can be generated. The constraint selection strategy used thus also aides the early identification of fruitless branches in the solution tree. Once all single solution constraints have been addressed the remaining constraints will reference free and shapeless objects. Selections are then made according to the number of candidate locations associated with free objects. A free object which has (for example) two candidate solutions can, in the worst case, result in only two branches in the solution tree; while a free object which has (say) ten branches associated with it will, in the worst case, result in ten branches in the solution tree.

The research team have conducted experiments into the application of alternative mechanisms whereby fruitless branches within the tree can be identified early in the search process. Examples include dependency based backtracking and *a priori* pruning techniques such as forward checking and "look ahead" (see Mackworth 1977 and Jaffar et al. 1986 for further discussion). However, it was found that the computational overheads involved were such that no advantage would be gained from the adoption of such techniques. The current, heuristically guided, pruning mechanism therefore seems to be the most appropriate. Similar approaches have been used in many well established problem solving systems, for example ALICE (Lauriere 1978). Whatever the case, the design of the SPARTA system is such that the constraint satisfaction mechanism is independent of the representation, so that the system is not tied to any particular Constraint Satisfaction Problem

(CSP) solving approach.

7. Example

The mechanism outlined here has been applied to a number of applications including Geographic Information Systems (GIS) (Coenen et al. 1994, Beattie et al. 1995, Coenen et al. 1996b), environmental impact assessment (Beattie et al. 1996), time-tabling and scheduling (Coenen et al. 1995), and shape fitting (Coenen et al. 1997). In this paper we illustrate the operation of the technique by applying it to the classic AI 8-queens problem. Although this is a 2-D problem its solution, using the SPARTA system, clearly demonstrates the operation of the system while at the same time indicating its versatility. The technique can equally well be extended to consider a 3-D or 4-D N-queens problem, however, space limitations dictate the exclusion of such an example from this paper.

The aim of the 8-queens problem is to place 8 queens on a standard chess board so that no queen threatens any other queen (A complete discussion of the problem can be found in Bratko 1986). The object space is thus of the form presented in Figure 2. A script to achieve this is presented in Table 4. The script comprises four parts:

```
space(8,8)

class(queens,free).

instance(q1,queens,[0..1792]).
instance(q2,queens,[1..1793]).
instance(q3,queens,[2..1794]).
instance(q4,queens,[3..1795]).
instance(q5,queens,[4..1796]).
instance(q6,queens,[5..1797]).
instance(q7,queens,[6..1798]).
instance(q8,queens,[7..1799]).

constraint([q1,q2,q3,q4,q5,q6,q7,q8], offset(all,[-1799, -1792..1792,
        -1785, -1542, -1530, -1285, -1275, -1028, -1020, -771, -765, -5
        -510, -257, -255, -007..0007, 0255, 0257, 0510, 0514, 0765,
        0771, 1020, 1028, 1275, 1285, 1530, 1542, 1785, 1799]), disjoin
```

Table 4: 8-Queens script

1) *Space declaration* where the object space in question is specified.
2) *Class declaration* in which the classes of interests are defined, in the above case there is only one class `queens`.
3) *Instance declarations* where instance or classes are described. In this case we have eight instances of the class `queens`, each of which has a location space associated with it describing a column on the board (we know that only one queen can exist per board column).

4) *Constraint declaration* in which the relations we require to exist between objects is defined. In our case we of course require the queens to be disjoint, but disjoint in a certain manner. This is defined by the offset included in the script which describes the horizontal, vertical and diagonal lines of attack associated with queens in the game of chess. Consequently the given constraint states that, to be valid, the location for a queen must be disjoint from the location of any other queen extended by the application of the defined offset.

Table 4 also illustrates one of the many short cuts, abbreviated forms and defaults, supported by the scripting language. The script includes a *compound* constraint which states that each instance in the given list of operands must be disjoint from any other instance in the list of operands (with the given offset applied). Thus the compound constraint actually translates to 28 individual constraints defining, in this case, each possible pairing of queens.

Once the script has been successfully parsed the resolution process commences with the selection of a constraint. At this stage, in the given example, each constraint has an equal weighting - each instance has 8 possible candidate locations associated with it - the first constraint encountered is therefore selected (`q1 disjoint from q2`). Satisfaction of this constraint then results in 8 branches in the solution tree each representing a possible configuration which satisfies the constraint. More specifically, each branch defines a single location for `q1` and five or six possible locations for `q2` (the remaining two or three possible locations having been eliminated at this stage as a result of satisfaction of the constraint).

The following 6 constraints selected, in the example, all reference `q1` (because for each branch this has only a single candidate location associated with it). Satisfaction of each these constraints can only produce one result (after merging). Thus once the first seven constraints have been satisfied the solution tree still comprises a single root node with eight branches originating from it, each culminating in a leaf node representing a possible location for `q1` and five or six possible valid locations for each of the remaining queens. Further selection of constraints, in the given scenario, is then made according to the minimum number of candidate locations associated with any instance in the scenario, i.e. a queen with only five possible locations. In this manner the solution process continues with the tree expanding up to some 200 branches at depth five where many fruitless branches are detected (dead-end nodes). Consequently the width of the tree begins to decrease, as the solution process continues, down to 150 branches at a depth of 6 where many solution configurations are reached (leaf nodes). All possibilities are considered by depth seven (28 branches) where the remaining solutions are discovered. In this manner all 92 possible solutions to the eight queens problem are obtained.

8. Conclusions

A qualitative spatio-temporal reasoning system has been described founded on a tesseral representation and linearisation of space. The mechanism offers the following significant advantages:

1) It is universally applicable regardless of the number of dimensions under consideration.

2) It is suited to a wide range of applications.

3) It is conceptually simple and computationally effective.

These advantages are gained primarily as a consequence of the particular tesseral representation on which the mechanism is founded which allows for computationally efficient manipulation of sets of addresses to support the resolution process. An additional advantage of the representation is that is compatible with many existing raster image and GIS formats.

The system has been tested against a great many spatial problems, a number of which were referred to in this paper. This testing has demonstrated that the approach can be applied to scenarios involving any number of dimensions without incurring any corresponding computational overhead as the number of dimensions increases. Current work is focused on noise pollution modelling and assessment in the city of London, and marine electronic chart interaction.

REFERENCES

1) J.F. Allen (1983). Maintaining knowledge about temporal intervals. Communications of the ACM, 26(11), pp832-843.

2) B. Beattie, F.P. Coenen, T.J.M. Bench-Capon, B.M. Diaz and M.J.R. Shave (1995). Spatial Reasoning for GIS using a Tesseral Data Representation. In Revell, N. and Tjoa, A.M. a(Eds), Database and Expert Systems Applications, (Proceedings DEXA'95), Lecture Notes in Computer Science 978, Springer Verlag, pp207-216.

3) B. Beattie, F.P. Coenen, Hough, A., T.J.M. Bench-Capon, Diaz, B.M. and M.J.R. Shave (1996). Spatial Reasoning for Environmental Impact Assessment. Third International Conference on GIS and Environmental Modelling, Santa Barbara: National Centre for Geographic Information and Analysis, WWW and CD.

4) I. Bratko (1986). PROLOG programming for Artificial Intelligence. Addison-Wesley.

5) F.P. Coenen, B. Beattie, B.M. Diaz, T.J.M. Bench-Capon and M.J.R. Shave (1994). A Temporal Calculus for GIS using Tesseral Addressing. In Bramer, M.A. and Macintosh, A.L. (Eds), Research and Development in Expert Systems XI, Proceedings of ES'94, pp261-273.

6) F.P. Coenen, B. Beattie, T.J.M. Bench-Capon, Shave, M.J.R and B.M. Diaz (1995). Spatial Reasoning for Timetabling: The TIMETABLER system. Proceedings of the 1st International Conference on the Practice and Theory of Automated Timetabling (ICPTAT'95), Napier University, Edinburgh, pp57-68.

7) F.P. Coenen, B. Beattie, T.J.M. Bench-Capon, Shave, M.J.R and B.M. Diaz (1996a). An Ontology for Linear Spatial Reasoning. In Wagner, R.R. and Thoma, H. (Eds), Database and Expert Systems Applications, (Proceedings DEXA'96), Lecture Notes in Computer Science 1134, Springer Verlag, pp718-727.

8) F.P. Coenen, B. Beattie, T.J.M. Bench-Capon, B.M. Diaz and Shave. M.J.R. (1996b). Spatial Reasoning for Geographic Information Systems. Proceedings 1st International Conference on GeoComputation, School of Geography, University of Leeds, Vol 1, pp121-131.

9) F.P. Coenen, B. Beattie, T.J.M. Bench-Capon, B.M. Diaz and Shave. M.J.R. (1997). A tesseral Approach to N-Dimensional Spatial Reasoning. In Hameurlain, A. and Tjoa, A.M. (Eds), Database and Expert Systems Applications, (Proceedings DEXA'97), Lecture Notes in Computer Science 1308, Springer Verlag, pp633-642.

10) A.G. Cohn, J.M. Gooday, B. Bennet and N.M. Gotts (1995). A logical approach to representing and reasoning about space, Artificial Intelligence Review, Vol 9, pp255-259.

11) R. Dechter, I. Meiri and J. Pearl (1991). Temporal constraint networks. Artificial Intelligence, Vol 49, pp61-95.

12) B.M. Diaz and S.B.M. Bell (1986). Spatial data processing using tesseral methods. Natural Environment Research Council, Swindon, England.

13) M.J. Egenhofer (1994). Deriving the composition of binary topological relations. Journal of Visual Languages and Computing, Vol 5, pp133-149.

14) D. Hernández (1991). Relative representation of spatial knowledge: the 2-D case. In D.M. Mark and A.U. Frank, Cognitive and Linguistic Aspects of Geographic Space, Kluwer, Dordrecht, Netherlands, pp373-385.

15) J. Jaffar, J-L. Lasses and M.J. Maher (1986). A logic programming language scheme. In D. DeGroot and G. Linderstrom (Eds), Logic Programming, Relations, Functions and equations, Prentice-hall, Englewood Cliffs, NJ, USA.

16) P. Ladkin (1992). Effective solution of qualitative interval constraint problems. Artificial Intelligence, Vol 52, pp105-124.

17) J-L. Lauriere (1978). A language and a program for stating and solving combinatorial problems. Artificial Intelligence, 10(1), pp29-127.

18) A.K. Mackworth (1977). Consistency in networks of relations. AI Journal 8(1), pp99-118.

19) H. Samet (1984). The Quadtree and Related Data Structures. ACM Comput. Survey 6(2), pp187-260.