A Relational Algebra Approach to Problems of Spatio-temporal Representation

Malcolm McQueen, Samuel Mann

Commerce Department, Otago Polytechnic

Dunedin, New Zealand

Phone: +64 3 479-6196 Fax: +64 3 474-1395

Email: malcolm@bit.tekotago.otago.ac.nz, smann@tekotago.otago.ac.nz

Presented at the 10th Colloquium of the Spatial Information Research Centre, University of Otago, New Zealand, 16-19 November, 1998

Abstract

This paper develops an approach for the representation of natural world phenomena within a relational structure. Relational algebra (RA) as the basis for database has been enormously influential and successful in the utilisation of computers in the world of business applications. This success has led to attempts to utilise relational algebra for application of computers to problems of the natural world with varying degrees of success, relational algebra underlies the database of most GIS. In current applications it does not deal well with many features of the natural world. These issues include the representation of features such as plains within valleys, relative representations, the representation of context knowledge, multiple instantiations of elements. The difficulties of addressing these issues within the RA database approach have lead many authors to suggest alternative architectures and relational algebra has become an unfashionable tool in this area. We take a contrary view and instead propose a method that uses the features of the natural world to build relational data structures from both prior and discovered knowledge of the natural world. This is developed from a consideration of the application of relational algebra to the business model. In this case the database issues discussed are rare and fall under the ambit of nulls. In the business case the model is constructed such that nulls (and the like) are considered as exceptions and dealt with on an individual basis. The key point in that the structure is sufficient as to render this problem a minor one. We propose a system that makes use of the structure of the database to include nulls and multiple representations to develop a structure that will support representation of the natural world within a relational structure. This paper takes the form of theoretical discussion with real illustrations and some practical testing.

Keywords and phrases: relational algebra, GIS, spatial, natural world, landscape ecology.

1 Introduction

Relational algebra, as the theoretical basis for database development, has been enormously influential and successful in the utilisation of computers for the world of business applications. This success has lead to attempts to utilise this theory for application of computers to problems of the natural world. This paper examines the use of the relational approach in the natural world. We stress that in investigating of the use of the relational approach we are not arguing that relational is preferable (ie over object approaches), rather that we simply wish to examine the limits of the relational model in the natural world.

2 Relational model

Codd first described the relational model for databases in 1970. As a model for the description of business application has proved remarkably resilient. Perhaps the greatest exponent of the relational approach is Date (1995). Date gives two definitions for the relational database (p 52):

1. the data is perceived by the user as tables (and nothing but tables)
2. operators at the user's disposal generate new tables from old (closure), these include SELECT, PROJECT and JOIN.

This definition includes "perceived by the user"; this then refers to the logical structure:

> The user of a non-relational system, by contrast, sees other data structures, either instead of or in addition to the tables of a relational system. Those other structures, in turn, require other operators to manipulate them (Date 1995 p22).

Figure2: High country area to be represented in a database

Common to most of Date's writing is the SPJ database. This Suppliers-Parts-Project database forms the basis for much of the discussion in the pre-eminent text: Date 1995 (p68). This database may be described by the data definition given in Figure 1.

```
CREATE DOMAIN S#  CHAR(5);
CREATE DOMAIN NAME  CHAR(20);
CREATE DOMAIN STATUS NUMERIC(5) ;
CREATE DOMAIN CITY  CHAR(15)
CREATE DOMAIN P#  CHAR(6)
CREATE DOMAIN COLOR CHAR(6)
CREATE DOMAIN WEIGHT NUMERIC(5);
CREATE DOMAIN QTY NUMERIC(9);
CREATE BASE RELATION S
     (S#        DOMAIN ( S# ) ,
     SNAME      DOMAIN ( NAME ) ,
     STATUS     DOMAIN ( STATUS ) ,
     CITY       DOMAIN ( CITY ) )
     PRIMARY KEY  ( S# ) ;
CREATE BASE RELATION P
     (P#        DOMAIN ( P# ) ,
     PNAME      DOMAIN ( NAME ) ,
     COLOR      DOMAIN ( COLOR ) ,
     WEIGHT     DOMAIN ( WEIGHT ) ,
     CITY       DOMAIN ( CITY ) )
     PRIMARY KEY  ( P# ) ;
CREATE BASE RELATION SP
     ( S#       DOMAIN ( S# ),
       P#       DOMAIN ( P# ),
     QTY        DOMAIN ( QTY ))
     PRIMARY KEY ( S#, P#  )
     FOREIGN KEY (S#) REFERENCES S
     FOREIGN KEY (P#) REFERENCES P;
```

Figure 1: Data definition for the SPJ database (S, P, SP shown, from Date 1995)

The questions that may be posed of this structure range from the trivial to the complex For example to get the colour and city for "non Paris" parts with weight greater than ten requires the following SQL (Date 1995 p226):

```
SELECT P.COLOR, P.CITY
FROM P WHERE P.CITY <> 'Paris'
  AND P.WEIGHT > '10;
```

3 Application of the relational model to the natural world

Given the success of the relational model in the business world it is hardly surprising that attempts have been made to extend the application beyond the business world. One such application, the subject of this paper, is the natural world. What though, is meant by 'natural world'? Take the landscape shown in Figure 2. This may be represented in a database, most usually in a relational database taking the form of a vector system where points, line and polygons are assigned unique identifiers (aka primary keys).

Such databases, with a focus on spatial representation and analysis are named Geographic Information Systems (GIS, synonymous with SIS: Spatial). Openshaw (1991) described the many features making up the toolbox that is considered GIS: overlay, multiscale, Boolean operations, spatial and textual relationships and data management, analysis and protection. Most GIS are, in effect, a relational database (with the addition of spatial operators). Each entity has attributes that may or may not include foreign key relationships to other tables (type of road and so on). The usual form of interaction with the user is through some form of structured query language (Egenhofer and Herring 1993). These databases are heavily used in automatic management/ facilities management (AM/FM) and in natural world management. Morehouse (1992) described how the ARC/INFO GIS falls under such a definition, while Batty (1992) describes how GIS would perform better if they focused even more on the relational concepts. Firns (1994) described an extension to the generic entity relationship approach to include spatial relationships (spatially extended entity relationship SEER). Not all GIS are relational, however, Herring (1992) described the TIGRIS system, which used an object-oriented approach.

For the sake of argument imagine a database (or GIS) that represents the landscape shown in Figure 2. Let's define the area of interest by posing questions we would like to ask:

1. Will the sheep preferentially graze the north facing spur in the foreground?
2. What impact will a warm winter have on the river?
3. How will this area respond to burning?
4. What impact does grazing at this altitude have on downstream water quality and quantity?
5. What is the relationship between soil percentage carbon levels and economic performance of the sheep station?
6. Is the water in this waterhole sufficient to get the stock back to the yards?
7. How does this tarn affect the competitive dynamics of the vegetation on the hillside?

These questions may seem diverse, but they are the sorts of questions that a professional in natural resource management might pose of a natural world database. The questions are clearly complex and some common themes may be seen. Note also that, so far, we have entirely overlooked the issues of data capture. How should such a database be populated? There are multiple representations, for example, different questions require the flat area to be considered both a plain and a valley. Entities may be classified more than once, the representation of a thing versus its use may be quite different. There may be entities that are defined only by their comparison to other entities; higher up, drier areas. There are some ideas for which the entity while initially simple may become complex. Take, for example, the river. The river may change course, but we would still call it the same river, it has inherent properties, upstream/ downstream, and changes in attributes both spatially and temporally. Spatial information must be stored to permit overlay and distance operators. These questions also require topological information (this polygon is next to that polygon).

We must also represent dynamics (eg. vegetation community trends, water balances) and be able to support the generation of simulations. Wholly different databases may be needed to support information such as the financial information of the farm alongside soil information gathered from sample points over the farm (probably spasmodically temporally and spatially). The questions we ask of the database may also be considered to be somewhat naive (Egenhofer and Mark 1995) and not really conform to SQL standards.

Burrough and Frank (1995 p 105) argued that "there is a large gap between the richness of the ways in which people can perceive and model spatial and temporal phenomena and the conceptual foundations of most commercial geographical information systems". Does this mean we have to admit that the questions we wish to pose of the hypothetical natural world database are simply "wicked" (McNamee *et al.* 1986 p 395), being particularly intransigent in terms of understanding and computing? Lowes and Walker (1995 p 5) write of environmental data being "invariably partial and inadequate" yet decision-makers are faced with "overwhelming data of limited utility". Do the problems discussed above and the incomplete information set preclude the natural world from the advantages of computing, in particular the benefits of the relational approach? Landauer (1996 p 6) described phase one of computing as "computers can do anything that can be reduced to numerical or logical operations", could it be that this does not include the natural world?

4 Business model

This paper has so far focussed on the identification of a number of difficulties in applying the relational approach to the natural world. Most, if not all, of these concerns may be summarised as a difficulty with ill-defined systems. This is not new. The very nature of ecological sustainability is almost undefinable. Schaller (1993 p 93) argued that "as a destination, sustainability is like truth and justice - concepts not readily captured in concise definitions". Standard texts on the design of information systems (Senn, 1989), however, take the approach that somebody, or a group of somebodies (Sallis *et al.* 1995), can understand and define the system of interest. Unfortunately, such an informed client is unlikely in the field of environmental management. Furthermore, even if the concepts could be better defined, the natural world features to be represented may remain vague and indefinite.

There are two quite distinct approaches for dealing with uncertainty in computing. The first is to attempt to remove uncertainty through problem definition. Aside from the previously discussed problem of the unlikelihood of this happening in environmental management (Mann 1995), two other problems are apparent. As Wilson (1994) pointed out, after strict problem definition the problem is limited to whatever the developer perceived as a problem, which results in a forced description of it. Wilson sees this as part of the "unrelenting pressure to translate vaguely stated socio-cultural [strategic] problems into clearly stated business [operational] ones" (p 295). A second difficulty with strict definition is the inability of such systems to cope with unexpected developments. Grossman (1994) discussed how predictions of air-pollution completely failed as a result of the unification of Germany. They could not of course be expected to predict this, but had no mechanisms for dealing with this uncertainty.

The second approach for dealing with uncertainty and ill-defined systems is to build a capacity for managing these phenomena within the computer system. Hay (1998) argued that the problem with the tools used as part of the relational approach lies with how people use the models, not the underlying complexity of what is being represented. This is not entirely foreign to the business relational model.

Hay argued that we must remember that the database does not attempt to model the real world, the database provides a model of the knowledge of the real world. This is the difference between the SPJ database and the hypothetical database that we assigned the task of describing Figure 2. The SPJ database describes a well-understood model, the natural world database we seem to be expecting to describe the real world.

Date's SPJ database describes a fictitious business model: supplier, parts and projects. It has clearly defined entities, relationships and domains. Date (1995 p70) acknowledges that the "database is of course extremely simple, much simpler than any real database is likely to be in practice most real databases will involve many more entities and relationships than this one does". Date argued though, that the basic constructs of the SPJ are sufficient to demonstrate the constructs of much larger, "real" databases. Similar databases used in teaching such as bibliographic (author - author_book - book - book_subject - subject), sales (customer - order - lineitem - product) and education (student - student_subject - subject) are used because they behave. They can be clearly defined with reference to our model of the real world, made easier by the entities all being human constructs to start with.

This is not to say that the databases used to represent these models of the real world are perfect. Date (1991) highlights a number of difficulties with the SPJ database. These difficulties which include domains, nulls and counter-intuitive results are briefly described below.

Date described the use of domains as "one of the most fundamental components of the relational model" (p28). A domain is a conceptual pool of values from which columns (in one of more tables) draw their actual values. This enables us to compare S.CITY and P.CITY. Whether the two meanings of city are relevant and logically comparable is ignored but it doesn't matter, we are happy to accept that the model the database represents has already considered that. The use of domains also does not help with the understanding of why S.S# = P.P# is not a legal expression despite S.S# = 'X4' and P.P# = 'X4' both being legal. The business model is sufficiently understood that we can ignore this paradox (Date 1991 p 43). Date goes on to discuss "cross-domain coercion" (p 48) that would allow comparison across domains but this is treated as an exception to the norm: "the system could either reject the query...or, - provided that the user is appropriately authorised - translate...".

Date (1991) discussed the use of "nulls", in short, "that some of the information is missing" (p219). This is not a value, rather that there is no value. Nulls have been interpreted in many ways; value not applicable, value unknown, value does not exist, value undefined, value not valid, value not supplied, and value is undefined. Date argued that each of these are quite separate and should be represented and manipulated in different ways (he expands the argument to say that with interaction between types of null there are an infinite number of type of null!). The usual practice (and that performed in the SPJ database) of defaulting nulls to a single value does not account for this variation. So why doesn't it matter? Again, the problem is treated as an exception; the business model is set up such that it doesn't matter.

Date also described areas where the SPJ model appears to be counter-intuitive, at least initially. He described how asking "Show me all the suppliers in London" and "Show me all the suppliers not in London" would not necessarily produce the total of all suppliers (1991 p 234). He argued that while the two states "location is London" and "location is not London" are indeed mutually exclusive and exhaust the full range of possibilities in the real world, the "database does not contain the real world - it only contains knowledge about the real world". A supplier whose location is not known will not result from either query. Date concludes that this counter-intuitive example arises from "a simple confusion over levels: the user is thinking at the real-world level, but the system is thinking at the level of its knowledge concerning the real world". The authors of this paper argue that the reason the SPJ database gets away with this potential confusion is that the underlying business model means the knowledge and the real world are never far apart and so, where the differences occur, they are sufficiently infrequent that they can be individually dealt wit as an exception.

In summary, then, SPJ the database provides a model of the system. This provides the user with knowledge of the real world. This knowledge is different from the real world, these differences can be seen in the problems with cross-domain comparisons, nulls and counter-intuitive results. The differences however are minor as the SPJ database is supported by a business model that very closely maps 'the real world'.

Pascoe and Penny (1994) aimed to provide a translation service between various levels of schema in a GIS. They described three levels of abstraction in describing a database, the abstraction being a collection of ideas about data to be represented. The conceptual abstraction is the "properties and relationships of the real world phenomena that are to be digitally represented, in terms that are natural for the users of that data" (p 277). The implementation abstraction is the "properties and relationships...[in a way suitable]...for implementation in a DBMS or GIS". The physical is the "data structures and algorithms for traversing those structures". The authors of this paper argue that the SPJ database works because there is a robust and well-developed model between the database and the 'real world'. This has been missing from attempts to create databases to represent the natural world.

5 A Model for the Natural World

Naveh and Lieberman (1994) described a systems-based theory of landscape ecological processes (SLE). Because of the complexity of interweaving biological, physical and cultural systems in the landscape, an interdisciplinary approach is encouraged in which landscape ecology provides a theoretical framework and information source. Landscape ecology helps guide the understanding of the Total Human Ecosystem (THE) as an ecological system that includes natural and human-made systems.

Beyond this conceptual linking of man-and-the-environment, SLE provides two kinds of information. First, landscape ecology describes the structure of the physical and biological environment at a scale that is practical for humans. Second, landscape ecology describes dynamic processes in time and space, and explores the way in which structures shape processes.

Naveh and Lieberman described how the conceptual framework for this approach is derived from three closely connected scientific theories:

General systems theory (GST): A holistic scientific theory and philosophy of the hierarchical order of nature as open systems with increasing complexity and organisation and with living systems and ecological systems as their special biosystem.

Biocybernetics: The theory of cybernetic regulation of biosystems, enabling their self-stablisation and self-organisation through deviation countering (negative) and deviation-amplifying (positive) feedback couplings.

Ecosystemology: The theory of a transdisciplinary ecosystem concept with the Total Human Ecosystem (THE) as the highest level of ecological integration and with the ecosphere as its concrete space-time-defined global landscape bread.

Naveh and Lieberman defined the approach largely through a large set of definitions that comprise the three components listed above. Of key importance throughout, is the focus on the system, "a set of elements in certain state...connected by relations...which together make up the structure of the system...because of relations...the system is always more than the sum of its elements (a whole)...emergent properties from system's behaviour of its elements as a whole...basic holistic axiom" (p 26). They describe the structure and coupling of black boxes of elements of systems, the structure of a system is the "set of relations that combines all elements with all their isomorphic relations" (p 27). Feedback couplings govern these relations and processes. Also important are hierarchical notions, derived from the holistic axiom where the whole is greater than sum of parts (because of relationships), there are therefore, emergent properties at hierarchical levels (p 50).

Most of the terms used in Naveh and Lieberman's development of SLE match terms used in relational database literature. The question arises, would inserting a model based on the SLE as the conceptual abstraction provide the structure needed to make the relational approach suitable for natural world databases?

There are three major consequences of any decision to consciously develop the SLE as the conceptual abstraction from which (it is hoped) relational databases may emanate. The first is a question of query, it may not be possible to answer questions such as those relating to Figure 2 in SQL. The query by users should be in the conceptual level "in terms that are natural for the users of the data" (Pascoe and Penny 1995 p 277). This suggests that the query should be by some form of natural language (Rashid *et al.* 1998) or schematic interface (Mann 1996, Albrecht *et al.* 1997). Given Date's definition of a relational database including "anything that the user perceives as tables and nothing but tables" (Date 1995 p 52), it is not clear whether placing this non-relational layer on top of a relational database renders the database non-relational.

The second question is of analysis. The complexity of model suggested by the SLE is too demanding for a single desktop computer (Mann 1997). Maxwell and Costanza (1995) and Westervelt *et al.* (1995) described responses to this problem in terms of closely linked clusters or supercomputers while Marr *et al.* (1998) describe a generic model for performing such analysis over a distributed network.

This leaves the representation; can this SLE model of the natural world and its related questions be supported by a relational approach? The authors believe that acknowledging the use of the SLE model will provide the structure that will allow the successful application of relational constructs to the natural world. Here we illustrate the approach taken.

6 Testing

The natural world problems of ill-defined and missing data and the necessity of multiple representations have their counterparts in the problems of nulls for the SPJ database. In particular the problem of Nulls for primary key and foreign key attributes. Solving this in a systematic way (as is required for the SLE) rather than treating individual cases as exceptions gives insight as to how the problem of representing the natural world may be solved. We consider two cases.

Case 1

A table has four existing entries. Suppose a database entry is to be made, one that has a missing foreign key value. We cannot directly make this entry without violating the referential integrity of the database. However, suppose we also know that the missing foreign key refers to one of the existing entries in the table to which it refers. We now have five possible models of the real world. This is our knowledge of the real world. It can be done by making a phantom entry in the primary key and supplying its primary key value to the missing foreign key. This phantom entry refers to one of the existing entries and, on the basis of no further information, we can consider the probability of its association with any particular one to be 1/N where N is the number of real (non-phantom) entries in the primary key table. The information on the probabilities of association is held in an association table. Queries of the database requiring a join on the affected tables can now represent our state of knowledge of the real world by also going through association table, picking up the probability of the answer to the real world on the way. The following example illustrates the action of the system.

1. The SPJ database has been loaded into the system.
2. A new record is added to the SPJ table with a missing value for the foreign key S#. The other foreign key values are:

   P# = P1 referring to the existing record in the P table - P1 |Nut |Red | 12|London

   J# = J1 referring to the existing record in the J table - J1 |Sorter |Paris
3. An option is used to indicate that it is not known whether the corresponding parent record exists.
4. The following SQL statement operates in the 'ordinary' fashion since the conditions do not involve the newly added supplier (the SPJ_AUG table is an intermediate table constructed from the association table and the SPJ table).

   — get all suppliers who supply Cogs to jobs in Athens

```
SELECT S.SNO,SNAME,ASSOC,QTY
FROM S,P,J,SPJ_AUG
WHERE S.SNO=SPJ_AUG.SNO
   AND P.PNO=SPJ_AUG.PNO
   AND J.JNO=SPJ_AUG.JNO
   AND P.PNAME='COG'
   AND J.CITY='ATHENS';
```

returns:

SNO	SNAME	ASSOC	QTY
S5	ADAMS	1	500

5. On the other hand, where the conditions do involve the newly added supplier, the linkage of this supplier to the others, creates extra output. The value of assoc may be interpreted as the probability that any of these extra records is real.

   – get all suppliers who supply Nuts to jobs in Paris

```
SELECT S.SNO,SNAME,ASSOC,QTY
FROM S,P,J,SPJ_AUG
WHERE S.SNO=SPJ_AUG.SNO
   AND P.PNO=SPJ_AUG.PNO
   AND J.JNO=SPJ_AUG.JNO
   AND P.PNAME='NUT'
   AND J.CITY='PARIS';
```

returns:

sno	Sname	assoc	qty
D_26		0.17	99
S1	Smith	1	200
S1	Smith	0.17	99
S2	Jones	0.17	99
S3	Blake	0.17	99
S4	Clark	0.17	99
S5	Adams	0.17	99

There are many possible variants on this situation. We might not know that the entity referred to by the missing foreign key actually already exists in our database. Or that this entity instance is one of a restricted set of existing entity instances. In all cases we can incorporate our knowledge of the real world into our database by appropriate probability values in the association table.

Case 2

Another case occurs when we wish to make an entry to the database but do not know which entity-type to instantiate. Perhaps the entry is fragmentary. We wish to record the data but have a choice under which entity-type to instantiate it. In this case we instantiate it under entity types that it fits and again record this multiple representation in the association table. Thus our database comes to represent a class of possible models of the real world corresponding to our knowledge.

This case may be simply extended to include multiple representations that stay in that state. This is the case for representing natural world phenomena such as plain and valley, or representations of hedgerows.

The above descriptions clearly obey the tenants of relational algebra. In the present implementation the associative table is a rather unwieldy creature and further development may force us to move it outside the relational structure. The individual models, however, will still be relational.

A potential criticism of the use of the relational model in this case may be related to performance, in particular the storage and processing requirements of the association table. If the association of every instance with every other instance recorded in the database has to be to be stored, then the association table grows as N^2 (N the number of instances in the database). However, in most practical cases the number of values that have to be recorded will be much smaller. First, it is only necessary to record non-zero values; second, the association values are likely to be non-zero only between instances of the same type (or a small number of the total number of types); and, third, if the instances in the database are distributed among a fairly large number of types the extra storage grows as the square of the number of instances in the tables as opposed to the entire database. GIS systems, in comparison to traditional business systems, tend to be more similar to this last case, so the overhead of the association table is a relatively less serious problem than might first be expected. A further issue is the difficulty of dealing with multi-dimensional data in the relational model. This impacts the current approach as we are dealing with multiple representations over time. This is the focus of future research.

7 Conclusion: Application to the natural world

We believe that many of the problems of applying relational algebra for representing natural world databases described in the first part of this paper can be traced to missing values of primary and foreign keys. We have described how examining the natural world in terms of a conceptual model (the SLE) lead us to examine the way in which the SPJ database is based on the business model. This allowed us to describe a systematic way of treating these problems for the SPJ database within the framework of relational algebra. This has resulted in the SPJ database representing our knowledge of the real world. We believe that such an approach can be applied to a relational database for representing our knowledge of the natural world.

Acknowledgements

The authors would like to thank Peter Brook and George Benwell for their discussions on this topic.

References

Albrecht, J. H., Jung, S. and Mann, S. (1997). VGIS: a GIS shell for the conceptual design of environmental models. Innovations in GIS 4. Z. Kemp. London, Taylor and Francis**:** 154-165.

Batty, P. (1992). "Exploiting relational database technology in a GIS." Computers and Geosciences **18**(4): 453-462.

Burrough, P. A. and Frank, A. U. (1995). "Concepts and paradigms in spatial information: are current geographical information systems truly generic." International Journal of Geographic Information Systems **9**(2): 101-116.

Date, C. J. (1991). Relational Database writings. Reading, Massachusetts, Addison-Wesley.528

Date, C. J. (1995). An introduction to database systems. Reading, Massachusetts, Addison-Wesley.839

Egenhofer, M. and Herring, J. (1993). Querying a GIS. Human factors in GIS. D. Medyckyj-Scott and H. Hearnshaw. London, Belhaven Press**:** 266 124-136.

Egenhofer, M. J. and Mark, D. M. (1995). Naive Geography. Spatial Information Theory: A Theoretical Basis for GIS,. A. U. Frank and W. Kuhn. Berlin, Springer-Verlag**:** 1-15.

Firns, P. G. (1994). Extending entity relationship modelling formalisms for the design of spatially referenced databases. In Proceedings of the Sixth Annual Colloquium of the Spatial Information Research Centre., Dunedin, New Zealand,, University of Otago. 289-297

Grossman, W. D. (1994). "Socio-economic ecological models: criteria for evaluation of state-of-the-art models shown on four case studies." Ecological Modelling **75/76**: 21-36.

Hay, D. C. (1998). "Making data models readable." Information Systems Management **15**(1): 21-33.

Herring, J. R. (1992). "TIGRIS: a data model for an object-oriented geographic information system." Computers and Geosciences **18**(4): 443-452.

Landauer, T. K. (1996). The trouble with computers: usefulness, usability and productivity. Cambridge Massachusetts, Bradford MIT.425

Lowes, D. and Walker, D. (1995). Environmental Management: Opportunities and Challenges in the Application of AI. Proceedings of the Workshop on AI and the Environment: Eighth Australasian Joint Conference on Artificial Intelligence, C. McDonald, Australian Defence Force Academy, Canberra, University College, University of New South Wales. 1-11

Mann, S. (1996). Environmental decisions with spatial process modelling. 1st International Conference on GeoComputation, R. J. Abrahart, University of Leeds. <http://www.bit.tekotago.ac.nz/~sam>

Mann, S. (1997). The GIS revolution, are we ready for Geocomputation? Second International Conference on GeoComputation and 9th Colloquium of the Spatial Information Research Centre, R. T. Pascoe, Dunedin, New Zealand, University of Otago. 31-39

<http://divcom.otago.ac.nz:800/conferences/geocomp97/>

Mann, S., Benwell, G. L. and Smith, C. B. (1995). Integrating Modelling and Simulation into a Problem Solving Paradigm for Improved Regional and Environmental Decision Making. International Congress on Modelling and Simulation, P. Binning, H. Bridgeman and B. Williams, University of Newcastle, NSW. 4:216-220

Marr, A. J., MacDonell, S., Benwell, G. L. and Mann, S. (1997). Spatial Process Modelling and Interoperability. International Conference on Interoperating Geographic Information Systems, Santa Barbara 3-4 Dec, NCGIA. 127-141

Marr, A. J., Pascoe, R. T., Benwell, G. L. and Mann, S. (1998). "Development of a Generic System for Modelling Spatial Processes." Journal of Computers, Environment, and Urban Systems. 22(1): 57-69

Maxwell, T. and Costanza , R. (1995). “Distributed Modular Spatial Ecosystem Modelling.” International Journal of Computer Simulation: Special Issue on Advanced Simulation Methodologies **5**(3): 247-262 <http://kabir.umd.edu/SMP/MVD/SME2.html.>

McNamee, P. J., Bunnell, P. and Sonntag, N. C. (1986**).** Dealing with wicked problems: a case study of models and resource management in southeast Alaska. Wildlife 2000: modelling habitat relationships of terrestrial vertebrates. J. Verner, M. L. Morrison and C. J. Ralph. Wisconsin, UWP**:** 395-399.

Morehouse, S. (1992**).** “The ARC/INFO Geographic Information System.” Computers and Geosciences **18**(4): 435-441.

Naveh, Z. and Lieberman, A. S. (1994). Landscape Ecology: Theory and Application. New York, Springer-Verlag.360

Openshaw, S. (1991). Developing appropriate spatial analysis methods for GIS. Geographical Information Systems; Vol 1: Principles. D. J. Maguire, M. F. Goodchild and D. W. Rhind. UK, Longman Scientific and Technical**:** 389-402.

Pascoe, R. T. and Penny, J. P. (1995**).** “Constructing interfaces between (and within) Geographical Information Systems.” International Journal of Geographical Information Systems **9**(3): 275-291.

Rashid, A. B. M., Egenhoffer, M. J. and Mark, D. M. (1998). “Natural language spatial relations between linear and areal objects: the topolgy and metric of English language terms.” International Journal of Geographical Information Systems **12**(3): 215-245.

Sallis, P. J., Tate, G. and MacDonell, S. G. (1995). Sotware Engineering. Sydney, Addison-Wesley.214

Schaller, N. (1993**).** “The concept of agricultural sustainability.” Agriculture, Ecosystems and Environment **46**: 89-97.

Senn, J. A. (1989). Analysis and design of information systems. New York, McGraw-Hill.853

Westervelt, J. D., Hannon, B., Sheikh, P., Cheng, A., Rubin, C., Dvorak, B., *et al.* (1995) Dynamic, Spatial, Ecological Modeling: A Demonstrated Simulation of the Sage Grouse Habitat at the Yakima Training Center, Washington, (USACERL Technical Report 95/16): US Army Corps of Engineers’ Construction Engineering Research Laboratories.

Wilson, F. A. (1994). “Computer support for strategic organisational decision-making.” Journal of Strategic Information Systems **3**(4): 289-298.