A Conceptual Data Modelling Framework Incorporating the Notion of a Thematic Layer

Dr Peter G. Firns[1]

Department of Information Science

University of Otago

August 1994

Abstract

Semantic data models comprise abstractions used, in conceptual database design, to represent real world relationships and aspects of the structure of real world phenomena. Such abstractions have previously been applied to the modelling of spatial concepts, but in the process their semantics are *implicitly* extended. This paper *explicitly* extends the semantics of the entity relationship model, defining two specific types of entity set to enable the notion of a thematic layer to be incorporated in entity relationship schemas. It places this in the context of a conceptual modelling framework to be used in the design of spatially referenced databases.

---

[1] Address correspondence to: Dr P.G. Firns, Lecturer, Department of Information Science, University of Otago, P.O. Box 56, Dunedin, New Zealand. Fax: +64 3 479 8311 Email: pfirns@commerce.otago.ac.nz

1 Introduction

This paper examines the notion of a thematic map layer from a conceptual data modelling perspective, then develops a technique for representing thematic map layers in conceptual schemas. The technique is based on standard entity relationship modelling notation and relies on extended semantics defined for two particular types of entity sets. By comparison to previous work in this area (eg. Calkins & Marble, 1987; Goh, 1988; Armstrong & Densham, 1989; Bedard & Paquette, 1989; Worboys et al, 1990; Laurini, 1991), no attempt is made to model the perceived complexity of spatial data and spatial relationships using existing semantic data modelling abstractions, the problems of doing this having been documented previously (Firns, 1992). The framework developed does, however, lend itself to subsequent extension to accommodate the modelling of spatial relationships. This is discussed in the latter sections of the paper.

In database design, conceptual schemas are developed as the basis from which logical database schemas are derived. Every database is itself a model of some subset of the real world, and the conceptual schema provides a mapping between reality and the database (Batra & Srinivasan, 1992). A conceptual schema must, therefore, strike a balance between two interrelated, and sometimes conflicting, objectives (Elmasri et al, 1985; Ferguson, 1988; Firns, 1990):

•   to represent, in an understandable manner, real world phenomena and relationships that may exist between them, and

•   to provide the basis for a database structure in which specific instances of the real world phenomena may be represented in the form of data values.

The latter objective is well supported by "record-oriented data models" (Kent, 1979), such as the hierarchical (Tsichritzis & Lochovsky, 1976), the network (Taylor & Frank,1976) and the relational (Codd, 1990) models. Record-oriented data models, more commonly referred to as *logical* data models, suffer from inherent limitations in terms of their ability to represent real-world semantics (Kent, 1979), thus they are seen as inappropriate for conceptual schema design. In response to these limitations, a number of semantic data models comprising formally defined abstraction mechanisms and diagramming conventions have been proposed (Abriel, 1974; Chen, 1976; Hammer & McLeod, 1981; Elmasri et al, 1985; Abiteboul & Hull, 1987; Hull & King, 1987; Peckham & Maryanski, 1988). Hence, semantic data modelling formalisms are increasingly being used in conceptual schema design, forming the basis for the

derivation of record-oriented database schemas, typically in the context of non-spatial databases.

A widely accepted and used semantic data modelling technique is the entity relationship (ER) model, which, in its original form has two major abstraction mechanisms: entities and relationships. Though various modifications and extensions to the original model have been proposed, the concepts of entities and relationships have remained central. The degree of acceptance and the relative stability of the ER model combined with the fact that an ER schema is able to be implemented in database systems based any of the major logical data models, have been factors in its adoption for the research described here.

Given that the ER model is typically applied in the context of non-spatial databases, it is appropriate to identify pertinent characteristics distinguishing spatially referenced databases from non-spatial databases. Spatially referenced databases consist of two notionally, and possibly physically, distinct components: spatial data and descriptive data. In terms of low level data structures, there are a number of options available for the storage of spatial data (Peuquet, 1984; Egenhofer & Herring, 1991), but is beyond the scope of this work to discuss these. Descriptive data (also referred to as *attribute* data, *textual* data or *aspatial* data) is essentially alphanumeric, describing real world phenomena, and is typically stored in a structure defined by one of the logical data models referred to above. The separation of spatial and descriptive data presents problems from a conceptual data modelling perspective.

First, linkages between the spatial and descriptive components of a spatially referenced database is a complex issue for which a generally accepted model has yet to emerge. Developers of SIS packages have found ways to implement such linkages and the results have been satisfactory as evidenced by a large number of SIS implementations supporting a diverse range of applications. At the conceptual modelling level, however, this is an issue which has not been addressed. A related issue is the fact that objects with both spatial and descriptive attributes of interest must be represented in two distinct forms in a spatially referenced database. This leads to problems in conceptual data modelling due to the conflicting nature of the two data modelling objectives. The first objective, that of developing a model of reality, implies that each object should be modelled only once. On the other hand, the second objective, that of modelling databases, could imply that each object should be modelled twice. These problems could be significantly alleviated if there existed a semantic data model which integrated

spatial and descriptive data. This paper, while not formally defining such a model, establishes a basic framework within which to do so.

2 The Thematic Layer as a Conceptual Model

Spatially distributed phenomena may be represented using the notion of a thematic map layer, whereby different types of phenomena are depicted on separate maps which may cover the same spatial extent. Historically, the thematic layer has played an important role in the structure of, and the use of, spatial information systems. The ARC/INFO software for example, uses the *coverage* as its basic unit of spatial data storage, and a coverage is "a digital version of single map sheet layer" (Peuquet & Marble, 1990:95). The GeoVision software (Westwood & Brinkman, 1988) uses layers, and other vector based systems adopt the concept of a thematic layer under various nomenclature[2]. Raster based systems by their very nature use the concept of a thematic layer as the basis for spatial analysis - given that a cell in a tessellation can have only one value assigned to it at any instance in time, then it is necessary to have multiple tessellations, each of which is, in effect, a distinct thematic map layer, where two or more themes covering the same spatial extent are required for spatial analysis.

The concept of cartographic modelling, "a general methodology for the analysis and synthesis of geographic data" (Tomlin, 1990; Tomlin 1991), also relies heavily on the concept of a thematic map layer as the basic means of organising data. Tomlin makes the following comments regarding "map layers" as defined in the cartographic modelling methodology: "Each of a map layer's locations, ..., is characterised in terms of exactly one attribute. In these terms, map layers are similar to what are variously called *themes, overlays, coverages, maps*, and *data elements* as well as *layers* in cartography; *images, data planes*, or *picture functions* in image processing; and *variables* in statistics" (Tomlin, 1990:6). It is reasonable to deduce that Tomlin's meaning by the term attribute was not intended to be the same as in the context of the entity-relationship model, wherein attributes are properties or characteristics of entity sets. The meaning and implied usage of "attribute" from the context of entity relationship modelling is retained here. Thus, Tomlin is formalising the concept of a single class of spatially distributed phenomenon being considered in isolation from all others as described earlier in this section.

---

[2]The two cited systems, ARC/INFO and GeoVision, were among the early geographic information system tools to achieve wide-spread commercial success, hence the significance attached to their particular use of layers in this context.

For the remainder of the paper, the qualifier 'map' is dropped, and the term 'thematic layer' is defined as:

... a *conceptual* model by which it is possible to represent a spatial extent as comprising multiple, independent, spatially distributed themes. An instance of a thematic layer corresponds to a single theme, but may comprise different types of phenomena which are, in some way, *directly* related to each other.

Notable in the above definition is the use of the term *conceptual* model. This implicitly points to an important difference between Tomlin's map layer and that of a thematic layer defined here. In the context of Tomlin's cartographic modelling methodology, the map layer more closely resembles a *logical* model as opposed to a conceptual model, the difference being that a logical model defines the way in which data will be structured in databases at the implementation level rather than representing the semantics of real world phenomena described by data. A map layer in Tomlin's cartographic modelling is defined by a statement in the data definition component of a high-level computational language. The notion of a thematic layer, as defined above, is an abstraction applicable to conceptual database design, and is independent of whether the intended implementation environment is layer based or non-layer based. The point is that a (conceptual) thematic layer will provide a framework within which to define relationships which will be required to be built into database structures regardless of whether the database is built upon the layer paradigm. While some work has been carried out in specifying database query language elements to process thematic map layers (eg. Scholl & Voisard, 1990), nothing has yet been done in the area of conceptual database design.

Another difference between Tomlin's map layer and that defined above, is that the latter allows for more than one type of phenomena to be represented in a single layer, with the qualification that these must be directly related to each other. This last point implies both that the phenomena in a thematic layer must be related to some common theme, *and* that there will be relationships, spatial and/or non-spatial, of a structural or functional nature amongst instances of different types of phenomena in a thematic layer. That is, it would be insufficient, in a properly developed conceptual schema, to include in the same thematic layer, phenomena that are related to each other only by virtue of spatial coincidence - that is that they have common locational characteristics. This approach is not inconsistent with the general concept of a thematic map comprising more than one theme (Hodgkiss, 1981). The following section presents a straight-

forward technique by which thematic layers are able to be represented in entity relationship schemas.

3 Modelling Thematic Layers in Entity Relationship Schemas

The thematic layer has the potential to play a central role in conceptual schema development for the design of spatially referenced databases. This section aims to show how the thematic layer complements abstraction mechanisms traditionally associated with entity relationship modelling. It first elucidates an analogy between a commonly occurring structure in ER schemas, that of *parallel linkages* (Kennedy, 1993), and the concept of multiple thematic layers. It then shows how this analogy enables the derivation of thematic layers from ER models incorporating appropriately defined entity-sets related by parallel linkages.

3.1 An ER Structure Analogous to a Thematic Layer

In an entity relationship schema, a pair of one to many relationships can connect a group of three entity sets in one of three different structures as illustrated in figure 1. Of interest for the purpose of this paper is the way in which entity-sets A and C can be said to be related to each other in each of the three models in figure 1, given that neither of the two relationships in any of the three models *directly* connect entity sets A and C. Figure 1(i) corresponds to a many to many relationship between entity-sets A and C, B being an associative entity-set. In figure 1(ii), there is an implicit 1:n relationship between A and C. That is, given an instance of C, there will be at most one instance of A. On the other hand, in figure 1(iii), there is no direct relationship between A and C implied. Instances of A and C in figure 1(iii) can only be related to each other by virtue of the fact that they are both related to the same instance of B. Any relationship between instances of A and instances of C is therefore purely coincidental. This is quite different than figure 1(i) in which instances of B incorporate pairings of instances of A and C. These observations have been examined formally from a data retrieval perspective (Kennedy, 1993).

The thesis of this paper rests upon an analogy drawn between entity-sets connected by parallel linkages and different thematic layers related to the same spatial extent. Recall that the definition of a thematic layer from the preceding section included the following key points:

- Thematic layers represent a spatial extent as comprising multiple, *independent*, spatially distributed themes.

- An instance of a thematic layer may comprise different types of phenomena, but these must be *directly* related to each other.

Both these points can be considered in the context of an analogy between thematic layers and parallel linkages. First, the *spatial extent* referred to in the first point can be considered analogous to the entity-set B in figure 1(iii). This is because each thematic layer covering that spatial extent is independent of all other thematic layers covering the same spatial extent. The spatial extent, therefore is the only common factor linking these layers, in the same way that instances of entity-set B are the only factors linking instances of entity-sets A and C. Put another way, thematic layers are independent of each other in all respects other than that they may be referenced to a common location. Similarly, instances of entity-sets A and C are independent of each other in all respects other than that they may be related to a common instance of entity-set B.
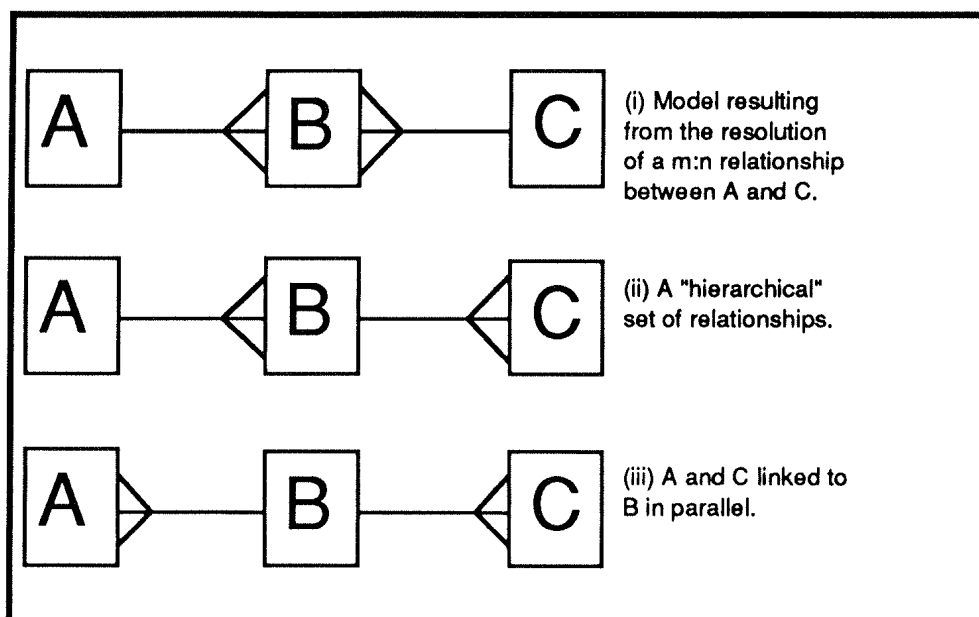

Figure 1: Generic Models of two 1:n Relationships & three Entity-sets

An example of a set of entities conforming to the generic structure defined by figure 1(iii) will further illustrate the point. The example is deliberately chosen for its *lack of spatial connotations*. This emphasises the fact that an analogy is being drawn between a structure typically occurring in aspatial data models and a structure inherently associated with spatially referenced data. Figure 2 is an entity-relationship model representing an application in which the management of a residential college requires to record data

pertaining to STUDENTs and their interests. Students' sporting and musical interests are to be recorded separately because the attributes of the two are quite different: for musical interests, the attributes include instrument played, proficiency level, type of music (jazz, rock, classical etc.) and others; for sporting interests, only the name of the sport and preferred position is required. This application results in a model with parallel linkages between the STUDENT entity-set and SPORTING-INTERESTS and MUSICAL-INTERESTS. Analysis of this model will indicate that some instances of the two latter entity-sets will be related to the same instance of STUDENT. Records in the database concerning the facts that a particular student plays the piano and that the same student also plays rugby are not directly related to each other. They are however, both related to the same student record and inextricably so. To include in the conceptual data model a relationship between the entity-sets SPORTING-INTERESTS and MUSICAL-INTERESTS would result in problems from a database maintenance perspective because the linkage would be spurious and result in the unnecessary storage of redundant data. This is analogous to the connection between two thematic layers covering the same spatial extent.
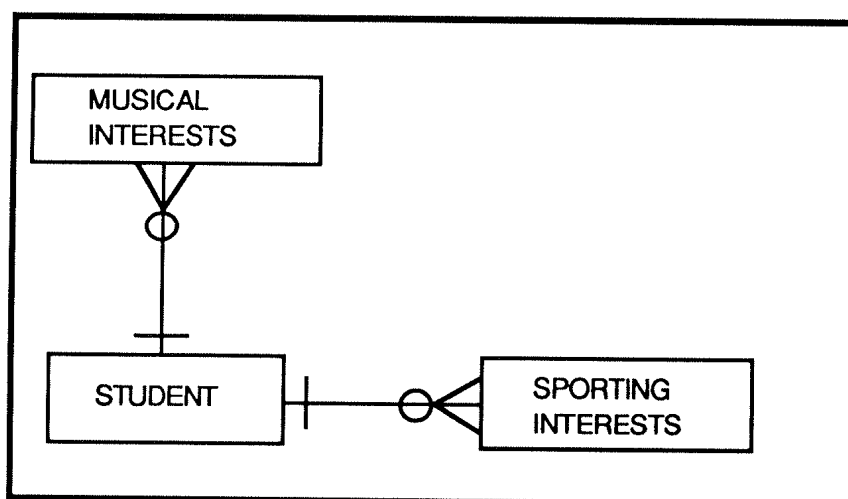

Figure 2: Entity Relationship Model of Student Sporting and Musical Interests

Extending the analogy, this discussion indicates that the representation of inter-layer relationships in a conceptual schema is not necessary or even desirable. Assuming that thematic layers have been appropriately defined, then only relationships within layers (ie intra-layer relationships) should be represented in a conceptual schema. Intuitively, this is so because inter-layer spatial relationships will not be represented within a database structure. An important point to note is that the fact that relationships are not represented in a conceptual schema or are not incorporated in a database structure, does not preclude subsequent analysis of those relationships taking place (this is true for the

student interests case as well as the multiple layered spatially referenced database case). Inter-layer relationships in a spatial database application would be derived and analysed using overlay operations, subsequent to the implementation of the database, in the usual way. Another way of viewing this issue is to consider the representation of topology in a vector based spatial database. One advantage of using multiple layers in such an environment is that the system does not have to maint in topology between every type of object in the database - topology is only maintained within layers, not between layers. This feature of layer based spatial databases is implicit in the analogy with parallel linkages. The parallel linkage notation is adopted in the next sub-section to explicitly represent the existence of multiple thematic layers covering the same spatial extent.

3.2   Proposed Technique for Modelling Thematic Layers

Two special types of entity sets with extended semantics, a LOCATION[3] entity set and any number of *locational-reference* entity sets, are defined as the basis to model reality as comprising a number of thematic layers. The relationships between LOCATION and all locational-reference entity sets are the basis for the representation of inter-layer connections. Intra-layer relationships are derived from the semantics of sub-types of locational-reference entity sets and relationships between these and the usual entity sets of the entity relationship model. The semantics of LOCATION and locational-reference entity sets are developed in the remainder of this section.

The role of the LOCATION entity set is to provide the basis for defining aspects of the underlying spatial coordinate system which will be used to implement models and also to support the representation of certain types of spatial meta-data such as scale. In principle, LOCATION could be represented in a relational database schema with a composite key comprising, for example, x, y coordinate pairs, x, y, z coordinates or x, y, z & t coordinates. Alternatively, it could refer to a cell in a tessellation. In practice, the means by which LOCATION is implemented will depend upon the spatial data structure underlying the spatial data management system being used. An instance of LOCATION is simply a reference to a place or locality at an arbitrarily defined resolution.

Given the existence of a LOCATION entity set, then any entity set, say LOCATIONAL-REFERENCE[I], linked to LOCATION in a 1:1 or 1:n relationship such

---

[3]Upper case letters are used here to denote reference to a specific, named entity set.

that a given instance of LOCATIONAL-REFERENCE^I will be related to only one instance of LOCATION, will be only indirectly related to any other such entity set, LOCATIONAL-REFERENCE^J. This situation is illustrated in figure 3, from which it may be observed that LOCATIONAL-REFERENCE^I, LOCATIONAL-REFERENCE^J and LOCATIONAL-REFERENCE^N are linked in parallel via LOCATION. In principle, this means that each may correspond to the existence of a single 'thematic layer', in the conceptual modelling sense of the term defined earlier. A locational-reference entity set is any entity set which participates in a 1:1 or 1:n relationship, with LOCATION, such that any instance of the entity set can be related to only one instance of LOCATION. Furthermore, a locational-reference entity set, by definition, cannot participate in a relationship with LOCATION, such that for an instance of the locational-reference entity set, there can be more than one instance of LOCATION. Formally, a relationship, defined by these two conditions may be represented as a subset of the cartesian product of LOCATION and LOCATIONAL-REFERENCE^I, or as a function such that LOCATION is functionally dependent upon LOCATIONAL-REFERENCE^I.

Complementary to the semantics of both locational-reference entity sets and their relationship to LOCATION, is a constraint restricting any entity set other than a locational-reference entity set from participating in relationships with LOCATION. Effectively, this means that any instance of any entity set participating in a relationship with LOCATION can be related to only one instance of LOCATION. An instance of LOCATION can be related to any number of instances of any number of locational-reference entity sets.
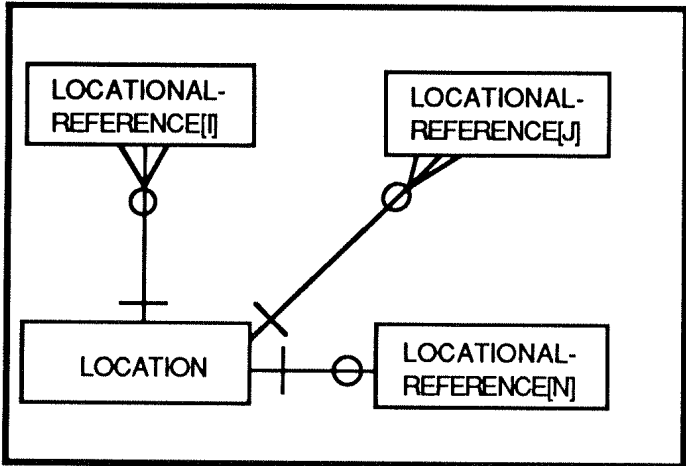

Figure 3: LOCATION and Related Locational-reference Entity Sets

4 A Conceptual Modelling Framework for Spatial Semantics

The framework proposed here will form the basis for extending the entity relationship model to incorporate the semantics of spatial relationships distinctively from non-spatial relationships *and* to distinguish between semantically different spatial relationships. A basic premise upon which the framework rests is that spatial relationships between objects in different thematic layers will not be represented within the structure of databases, an issue already discussed in the previous section. The modelling abstractions to support the concepts discussed are not fully developed in this paper, hence the term 'modelling framework' is used rather than 'model'. There are two spatial elements to the proposed conceptual modelling framework: the modelling of spatial characteristics of real world phenomena; and the modelling of spatial relationships between world phenomena. Each of these is briefly discussed below.

4.1   Modelling Spatial Characteristics of Real World Phenomena

It is readily observed that some form of geometric abstraction is required to represent the spatial dimension of objects in databases. As part of conceptual schema design therefore, it would be useful to specify the required dimension of each type of object. It is assumed that the appropriate spatial dimension should be determined on the basis of the nature of the real world phenomena being modelled. In practice however, this is precluded from being adopted as an overriding design principle because spatial dimension will be partly determined on the basis of the model of space supported by the intended implementation tool, and on the limitations of such model. Assuming a two dimensional model of space for example, each object to be represented spatially in a database must be represented as a zero dimensional, a one dimensional or a two dimensional object. Furthermore, given the nature of the model, it can be advantageous to represent some objects as zero or one dimensional objects even though it is known that in reality, any concrete object has at least two and probably three dimensions. Hence, virtually any spatial model requires some form of abstraction when representing real world phenomena.

Notwithstanding these observations, it is contended that the best geometric representation of a real world phenomena is inherently a function of the characteristics of that phenomena. Hence, a conceptual data model oriented to the design of spatially referenced databases should facilitate the representation of *dimension* associated with object types (entity sets in the entity relationship model).

Within the modelling framework proposed here, it is assumed that instances of any entity set directly related to a locational-reference entity set (see figure 3) must be represented spatially in a database. The definition of each such entity set should therefore, include the specification of the desired spatial representation of instances of the entity set. Thus, the basic modelling framework incorporates the notion of a *geometric type attribute* associated with any entity set related to a locational reference entity set or sub-type thereof. A *type attribute* (or class attribute) is an attribute, the value of which is the same for every instance of an entity set (Hull & King, 1987; Tan & Dillon, 1992). This is therefore, seen as an appropriate mechanism to represent dimensions in an extended entity relationship model, assuming that each instance of a locationally referenced entity set will have the same geometric data type. If this assumption does not hold for as particular entity set, then sub-types can be created in the entity relationship schema. Each entity set or sub-type which is locationally referenced can then be assigned a value for the geometric type attribute - typical values for such an attribute being *point, line* or *region.*

An important point to note regarding the above discussion, is that the modelling framework does not specify how spatial data will be structured or stored in databases. At the logical level, some form of spatial data model is required, as is a spatial data structure at the physical level, but reference to the logical and physical levels is deliberately excluded from the framework proposed here. A conceptual data model of the nature being proposed will, it is envisaged, be a useful complement to developments in spatial data types for database systems, an area of increasing interest to a number of researchers (eg. Choi & Luk, 1992; Vijlbrief & van Oosterom, 1992; Güting & Schneider, 1993). The representation of geometric data types at the conceptual level will provide useful input to schema design for database systems incorporating spatial data types or geometric object types.

4.2   Spatial Relationships Between Real World Phenomena

The second main element required to extend the entity relationship model for the design of spatially referenced databases is the ability to represent spatial relationships and the semantics thereof. This is an issue which has not previously been addressed at the conceptual modelling level. Other researchers have attempted to build spatial relationships into conceptual schemas using either the entity relationship model (eg. Laurini, 1991) or some other semantic data model (eg. Worboys et al, 1990), but none of these has attempted to model the spatial semantics of such relationships. In contexts other than that of conceptual data modelling, significant contributions have been made

in defining the nature of spatial relationships - for example by formally defining the possible topological relations between point intervals in one and two dimensions (Pullar & Egenhofer, 1988; Egenhofer & Franzosa, 1991). Spatial relationships in conceptual schemas can, it is contended, be usefully represented distinctively from non-spatial relationships. It is aimed to specify a semantic model incorporating a taxonomy of spatial relationships and related constraints, as exemplified in the following paragraph.

Assuming the notion of a thematic layer as defined earlier, the only spatial relationships required to be formally represented in database structures are intra-layer relationships. Furthermore, given the notion of a geometric type attribute defined above, constraints are placed on the nature of spatial relationships in which instances of pairs of entity sets can participate - for example, given the general notion of a 'contain' relationship, one dimensional objects could not be specified as 'containing' two dimensional entities. By defining a taxonomy of such relationships and associated constraints, abstraction mechanisms corresponding to formally defined topological relationships can be integrated with the abstraction mechanisms traditionally associated with conceptual database design (eg. entities and relationships) to form a semantic data model applicable to the design of spatially referenced databases.

5   Conclusions and Research Direction

It was stated in the introduction that this paper would address problems associated with the separation of spatial and descriptive data in spatially referenced databases, proposing a framework within which to solve these problems. The proposed modelling framework is based on the notion of a thematic layer. The modelling of thematic layers is based on an analogy between the notion of a thematic layer and that of a parallel linkage, a commonly occurring structure in entity-relationship models. Central to the modelling of thematic layers, is an acceptance of the definitions of the entity-sets LOCATION and LOCATIONAL-REFERENCE as discussed in section 3.2. Thematic layers provide an objective means by which to separate groups of entities so that spatial relationships desired to be built into database structures will occur only within layers. Complemented by the notion of a geometric type attribute, this provides the basis for a framework within which to rigorously represent the semantics of spatial relationships of interest when designing spatially referenced databases. This framework addresses the issues referred to because it will lead to the development of a conceptual data model in which spatial and non-spatial relationships will be able to modelled in an integrated fashion independent of the type of system in which databases will be implemented.

Ongoing research in this area aims to further refine the definition of the geometric type attribute discussed in the preceding section, and to propose appropriate notation for representing in data models the geometric data type of entity-sets. Another avenue of research being pursued involves examining the potential use of semantically richer data modelling abstraction mechanisms to represent spatial relationships (eg. relevant topological connections) between entity-sets within thematic layers. It is intended that this research will facilitate the formal definition of an extended entity relationship model to incorporate spatial concepts. Diagramming notation based on such an extended model will then form the basis for a new methodical approach to the design of spatially referenced databases.

An issue not discussed at length in the main body of the paper, but mentioned here because of its inherent association with a conceptual data modelling framework, is that of spatial meta-data pertaining to scale, positional accuracy and lineage. Such data could, in principle, be represented by type attributes either associated with LOCATION or with locational-reference entity sets. This is to be the subject of further research into the specification of a data dictionary for spatial data management.

References

Abiteboul, S. & Hull, R. [1987] "IFO: A Formal Semantic Database Model", *ACM Transactions on Database Systems*, 12 (4), December, 525-565.

Abriel, J.R. [1974] "Data Semantics". In Klimbie, J.W. & Koffeman, K.L. (eds), "Database Management". *Proceedings of the IFIP Working Conference on Database Management* (Amsterdam: North-Holland), pp.1-60.

Armstrong, M.P. & Densham, P.J. [1990] "Database Organisation Strategies for Spatial Decision Support Systems", *International Journal of Geographical Information Systems*, 4 (1), 3-20.

Batra, D. & Srinivasan, A. [1992] "A Review and Analysis of the Usability of Data Management Environments", *International Journal of Man-Machine Studies*, 36, 395-417.

Bedard, Y. & Paquette, F. [1989] "Extending Entity/Relationship Formalism for Spatial Information Systems". In *Proceedings: Auto-Carto 9* (Baltimore: ASPRS & ACSM), pp.819-827.

Calkins, H.W. & Marble, D.F. [1987] "The Transition to Automated Production Cartography: Design of The Master Cartographic Database", *The American Cartographer*, 14 (2), 105-119.

Chen, P.P. [1976] "The Entity-Relationship Model - Towards a Unified View of Data", *ACM Transactions on Database Systems*, 1 (1), March, 9-36.

Choi, A & Luk, W.S. [1992] "Using an Object-Oriented Database System to Construct a Spatial Database Kernel for GIS Applications", *Computer Systems Science and Engineering*, 7 (2), April, 100-121.

Codd, E.F. [1990] *The Relational Model for Database Management (Version 2)* (Reading, Massachusetts: Addison-Wesley).

Egenhofer, M.J. & Franzosa, R.D. [1991] "Point-set Topological Relations", *International Journal of Geographical Information Systems*, 5 (2), 161-174.

Egenhofer, M.J. & Herring, J.R. [1991] "High Level Spatial Data Structures for GIS". In *Maguire et al, 1991*, Vol. 1, 227-37.

Elmasri, R., Weeldreyer, J. & Hevner, A. [1985] "The Category Concept: An Extension to the Entity-Relationship Model", *Data & Knowledge Engineering*, 1, 75-116.

Ferguson, J.R. [1988] "Data Modelling and Systems Analysis", Paper Presented at *19th Colleges of Advanced Education Computing Conference*, Australia.

Firns, P.G. [1990] "Determining a Useful Balance Between Understandability and Rigour in Data Modelling", *New Zealand Journal of Computing*, 2 (1), 13-21.

Firns, P.G. [1992] "Data Modelling for Spatial Information Systems: The Notion of a Spatial Entity", *New Zealand Journal of Computing*, 4 (1), 17-28.

Goh, P. [1988] "Spatial Database Updates - A Formal Treatment Using Entity-Relationship Diagrams", *The Australian Surveyor*, 34 (3), 307-313.

Güting, R.H. & Schneider, M. [1993] "Realms: A Foundation for Spatial Data Types in Database Systems". In Abel, D. & Ooi, B.C. (eds), "Advances in Spatial Databases", Lecture Notes in Computer Science No. 692 (Springer), pp.14-35.

Hammer, M. & McLeod, D. [1981] "Database Description with SDM: A Semantic Database Model", *ACM Transactions on Database Systems*, 6 (3), 351-386.

Hodgkiss, A.G. [1981] *Understanding Maps: A Systematic History of their Use and Development* (Kent: Dawson).

Hull, R. & King, R. [1987] "Semantic Database Modeling: Survey, Applications, and Research Issues", *ACM Computing Surveys*, 19 (3), September, 201-260.

Kennedy, G.J. [1993] "A Systematic Approach to the Specification and Evaluation of an Information Systems Development Methodology", *PhD Dissertation* (Dunedin, New Zealand: University of Otago).

Kent, W. [1979] "Limitations of Record-Based Information Models", *ACM Transactions on Database Systems*, 4 (1), 107-131.

Laurini, R. [1991] "Introduction to E.R. Modelling for Urban Data Management". In Kangassalo, H. (ed), "Entity-Relationship: The Core of Conceptual Modelling", *Proceedings: 9th International Conference on The Entity-Relationship Approach*, (Amsterdam: North-Holland), pp.475-485.

Maguire, D.J., Goodchild, M.F. & Rhind D.W. (eds.) [1991] *Geographical Information Systems: Principles and Applications*, vol. 1 & 2 (London: Longman).

Peckham, J & Maryanski, F. [1988] "Semantic Data models", *ACM Computing Surveys*, 20 (3), September, 153-189.

Peuquet, D.J. [1984] "A Conceptual Framework and Comparison of Spatial Data Models", *Cartographica*, 21, 66-113.

Peuquet, D.J. & Marble, D.F. (eds), [1990] *Introductory Readings in Geographic Information Systems* (London: Taylor & Francis).

Pullar, D.V. & Egenhofer, M.J. [1988] "Toward Formal Definitions of Topological Relations Among Spatial Objects". In *Proceedings: 3rd International Symposium on Spatial Data Handling* (Sydney: International Geographical Union), 225-241.

Tan, P.L. & Dillon, T.S. [1992] "Incorporating Semantics, Relationships and Exceptions into Object-Oriented Data Models", *Computer Systems Science and Engineering*, 7 (2), April, 122-135.

Taylor, R. & Frank, R. [1976] "CODASYL Database Management Systems", *ACM Computing Surveys*, 8 (1), 67-104.

Tomlin, C.D. [1990] *Geographic Information Systems and Cartographic Modelling*. Englewood Cliffs (New Jersey: Prentice-Hall).

Tomlin, C.D. [1991] "Cartographic Modelling". In *Maguire et al, 1991*, Vol. 1, 361-374.

Tsichritzis, D. & Lochovsky, F. [1976] "Hierarchical Database Management", *ACM Computing Surveys*, 8 (1), 105-24.

Vijlbrief, T. & van Oosterom, P. [1992] "The GEO++ System: An Extensible GIS". In *Proceedings: 5th International Symposium on Spatial Data Handling* (Charleston: International Geographical Union), Vol. 1, 40-50.

Westwood, K. & Brinkman, J. [1988] "Toward the Successful Integration of Relational and Quadtree Structures in a Geographic Information System". In *Proceedings: URISA '88*, The 1988 Conference of the Urban and Regional Information Systems Association.

Worboys, M.F., Hearnshaw, H.M. & Maguire, D.J. [1990] "Object-Oriented Data Modelling for Spatial Databases", *International Journal of Geographical Information Systems*, 4 (4), 369-383.