The Derivation of Thematic Map Layers from Entity-Relationship Data Models

Peter G. Firns[1]

Department of Information Science

University of Otago

November 1993

Abstract

Semantic data models comprise formally defined abstractions for representing real world relationships and aspects of the structure of real world phenomena so as to aid database design. While previous research in spatial database design has shown that semantic data models are amenable to explicitly representing some spatial concepts, this paper shows that semantic data models may usefully be applied to the design of spatial databases even without explicitly representing spatial concepts. Specifically, an entity-relationship model comprising only "is-associated-with" relationships is used as the basis from which to define thematic layers for a layer based spatial database.

---

[1] Address correspondence to: Mr 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

1 Introduction

In the context of database design, a conceptual data model[2] is a model of real-world semantics from which the structure of a database may be derived. Every database is a model of some subset of the real world and the conceptual data model provides a mapping between reality and the database (Batra & Srinivasan, 1992). A conceptual data model must therefore, strike a balance between two interrelated, and sometimes conflicting, objectives (Elmasri et al, 1985; Ferguson, 1988; Firns, 1990):

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

- to develop the basis for database structures in which specific instances of the real world phenomena may be represented in the form of data values.

The latter objective has been 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, however, suffer from inherent limitations in terms of their ability to represent real-world semantics (Kent, 1979), thus they are seen as inappropriate for conceptual design. In response to these limitations, much research has been carried out in semantic data modelling, a number of semantic data models comprising formally defined abstraction mechanisms and diagramming conventions having been proposed (Abriel, 1974; Chen, 1976; Hammer & McLeod, 1981; Tsichritzis & Lochovsky, 1982; Elmasri et al, 1985; Abiteboul & Hull, 1987; Hull & King, 1987; Peckham & Maryanski, 1988; Batini, 1989; Kangassalo, 1991). Semantic data modelling is increasingly being used in conceptual database design, forming the basis for the derivation of record-oriented database schema, typically in the context of non-spatial databases.

This paper is concerned with the application of semantic data modelling to the design of spatially referenced databases. By comparison to previous work in this area, (e.g. Calkins & Marble, 1987; Goh, 1988; Armstrong & Densham, 1989; Worboys et al, 1990; Laurini, 1991), relatively simple abstraction mechanisms are used to develop a data model for a spatially referenced database,

---

[2] The term data model lacks a rigorous definition and, in the context of this paper, has at least two related, but quite different, meanings: on one hand, the term refers to a model of a specific application environment; on the other hand, it may also refer to the modelling formalisms used to develop such a model. Thus, a data model for a city council's property database may be developed using the entity-relationship [data] model. In general, the intended meaning is readily derived from the context in which the term is used.

illustrating the point that the perceived complexity of spatially referenced databases does not necessarily have to be built into conceptual data models. The structure of the paper is as follows: the nature of spatially referenced databases is discussed in the next section. Previous work in the area of semantic data modelling for the design of spatially referenced databases is then examined in the context of the data modelling objectives stated above. Following this, the data modelling conventions and notation used in the remainder of the paper are elucidated. A case study illustrating the application of these to the design of a spatially referenced database is then presented followed by discussion and conclusions.

2 Spatially Referenced Databases: A Data Modelling Perspective

In order to establish the validity or otherwise of applying semantic data modelling techniques to the design of spatially referenced databases, it is appropriate to first identify the pertinent characteristics distinguishing spatially referenced databases from non-spatial databases. Two fundamental features distinguishing spatial information systems (SIS) from non-spatial information systems are the capabilities:

- to relate otherwise disparate data sets on a spatial basis, for example by performing overlay operations, and

- to perform spatial analyses based on topological relationships represented, either directly or indirectly, in spatial data.

Both these capabilities are dependent upon the existence of the appropriate data in an appropriate format, for example in a spatially referenced database. A spatially referenced database is, for the purpose of this paper, defined as a database with the following characteristics:

- it incorporates textual data of the type typically stored and managed within, for example, relational database management systems

- some non-spatial relationships may be represented in the database structure

- it incorporates spatial data

- some spatial relationships may be represented in the database structure

- it incorporates some form of referencing between spatial and non-spatial data.

3

Spatially referenced databases therefore, consist of two notionally, and possibly physically, distinct components: spatial data and descriptive data. Spatial data represents location of, and spatial relationships between, real world phenomena. There are two major forms of spatial data representation: vector and raster. Vector data models represent the world as geometric structures derived from (inter alia) point, line and polygon primitives, whereas raster data models represent real world phenomena as values assigned to cells in a tessellation. In terms of low level data structures, there are a number of options available for each (Peuquet, 1984; Egenhofer & Herring, 1991). Descriptive data (alternatively known as attribute data, textual data or aspatial data) is essentially alphanumeric, describing real world phenomena. In terms of physical database architecture, three alternatives have emerged (Bracken & Webster, 1989; Maguire et al, 1992):

- file processing systems

- hybrid systems or dual database systems

- extended database management systems.

Under file processing architectures, data sets (spatial and descriptive) are stored in files and accessed for processing by spatial analysis software directly through the operating system. Hybrid architectures, on the other hand, utilise conventional database management systems to store and manage descriptive data while spatial data is stored in files and managed by SIS software interacting with the operating system. Extended database management system architectures store both spatial and descriptive data in a conventional database management system augmented by SIS software to provide appropriate spatial analysis functionality. Of the three architectures, hybrid approaches are the most common, with the so-called "geo-relational model", in which a relational database management system is adopted for the storage of descriptive data, probably the most widely used (Morehouse, 1985; Bracken & Webster, 1989; Healey, 1991; Maguire et al, 1992). Regardless of the architecture being used, two issues unique to the SIS context present problems from a 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. It is readily observed though, that the physical separation of spatial and descriptive data creates an additional level of complexity for SIS application programmers. This problem may however, be negated by addressing the second problem, described below.

The second, perhaps more subtle, 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 the database, would seem to imply that each object should be represented twice. It is contended that this situation could be significantly alleviated if there existed a general conceptual data model which integrated spatial and descriptive data. In the absence of such a model, however, careful application and analysis of standard conceptual data modelling techniques can provide some assistance in the design of spatially referenced databases and in subsequent application programming.

3 Previous Work

Previous work in this area has focused on the representation of so called spatial entities[3] and on the explicit representation of spatial relationships between entities. One approach, for example, used the IFO semantic data model (Abiteboul & Hull, 1987) to model the composition of British post code areas (Worboys et al, 1990). This approach, illustrated in figure 1, incorporates a polygon object type of which postal areas, postal districts and postal sectors are modelled as sub-types. It is argued here that approaches such as this implicitly incorporate different levels of abstraction in the one model. This argument is based on the premise that the polygon is itself an abstraction, thus the polygon object type is a meta-model of reality. A polygon is in fact a means of depicting an administrative area graphically. Its representation as an object type in a data model is therefore, one step further removed from reality than its sub-types. Thus, the first data modelling objective is not well supported. Also, the fact that this meta-model component is likely to be implemented differently than other components of the model and that the model provides no guidelines as to how that implementation will be achieved, is contrary to the second data modelling objective.

A simple way of overcoming the first of these criticisms would be to rename the polygon object type in figure 1 to, say, postal region. The model could incorporate a functional relationship between this and a distinct polygon object type, the structure of which could also be modelled in

---

3 Many synonyms for the term "spatial entity" are used in the literature, for example geographic entity, geometric entity and spatial object.

IFO (Worboys et al, 1990). This does not however, address the lack of support for the second data modelling objective.

Irrespective of the above, it may be argued that the additional complexity of the model in figure 1, brought about by the inclusion of spatial entities, is not justified by the minimal spatial semantics encapsulated. For example, the fact that a postal area is a spatial grouping of postal districts cannot be determined from the model. Also, polygon could just as well be modelled as an atomic type, thus reducing the model's complexity whilst losing no semantic information.

The incidence of problems such as these is not restricted to the use of semantic data models such as IFO. It has been shown that similar issues can arise in applying the entity-relationship model to the design of spatially referenced databases (Firns, 1992). Underlying these problems are two factors:

- a lack of focus on the two data modelling objectives

- undue concern with the perceived complexity of spatial data.

The first of these is a critical issue in the context of any data modelling undertaking. Failure to meet either of these objectives will undoubtedly diminish the value of any data model. The second is unique to the SIS context. While it may be tempting to incorporate spatial relationships in data models, this paper argues that it is not currently feasible to do so. It is contended that further research is required to extend semantic data modelling theory to accommodate spatial relationships and spatial structures in conceptual data models. It is therefore prudent to use existing data modelling techniques for modelling only those aspects of reality for which they were originally defined. This does not however, preclude the ability to derive aspects of spatial database structures from conceptual data models. The remainder of the paper illustrates this point.
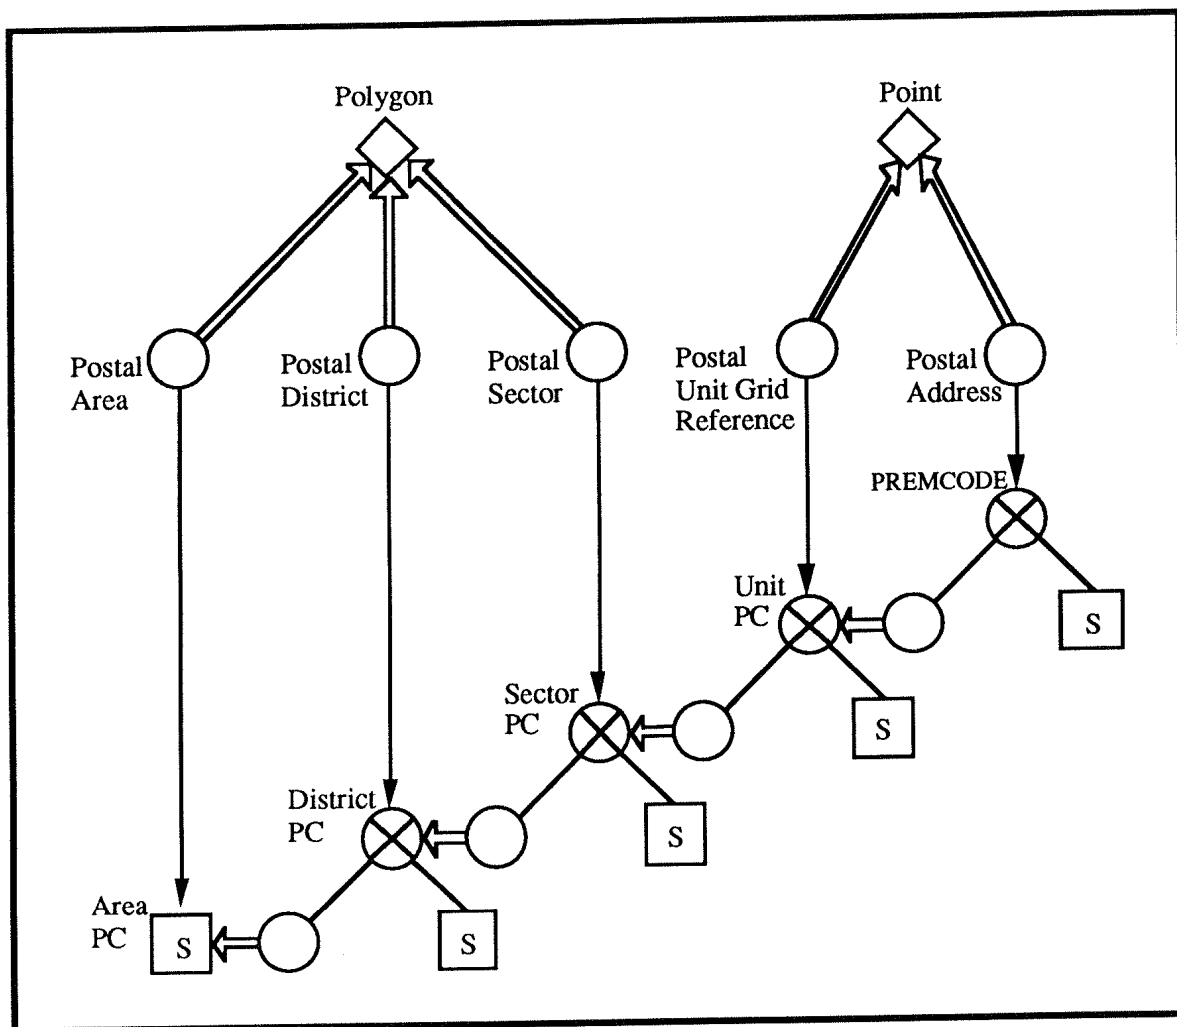
Figure 1: IFO Data Model Incorporating Spatial Entities (Worboys et al 1990)

4 ER Modelling Formalisms and Notation

Semantic data modelling abstractions are used in conceptual database design to model structures and relationship types which occur in reality and which are to be represented in database schema. In the context of non-spatial information, the derivation of database schema and integrity constraints from semantic data models is well understood (Abiteboul & Hull, 1987). Accordingly, databases can be modelled and structured to support the production of information that can validly be assumed to be consistent and to have an acceptable level of integrity, if not "correctness". An important principle of conceptual database design is that models are developed independent of the database management system in which they will eventually be implemented. This means that data modellers are primarily concerned with developing appropriate representations of reality. A conceptual data model is subsequently transformed to a logical model, for example the relational

model, which is specific to a particular class of database management system. The database design process is depicted in figure 2.

One aspect in which semantic models differ, is in the degree of semantic richness supported. Entity-relationship (ER) modelling is a conceptual data modelling technique, originally proposed by Chen (1976), which, in its original form has just two abstraction mechanisms: entities and relationships. The ER model, though widely accepted and used, is the subject of on-going research and various modifications and extensions to the original model have been proposed. The concepts of entities and relationships have however, remained central. This, combined with the fact that an ER model is able to be implemented in any of the major database models, has gained ER modelling wide acceptance as an analytical and design technique. The ER model suits the purpose of this paper because, while it overcomes the lack of semantics provided by record-oriented models, it is sufficiently simple in concept so as not to over-shadow issues related to the modelling of spatial concepts.

An entity is some phenomenon, about which, data is to be stored. An entity may be an object, an event, a place, a person, an activity or any other phenomenon of interest. An entity-set is a set of similar entities (e.g. the set of all properties) with similar characteristics or attributes of interest. A relationship is a logical, meaningful connection between entity-sets. Relationships can be specified in terms of the following properties:

- whether participation in the relationship by instances of each of the participating entity-sets is optional or mandatory - this characteristic is termed relationship participation or simply participation

- the cardinality of the relationship which indicates the number of such relationships in which instances of each of the participating entity-sets may appear - relationship cardinality can be one to one (1:1), one to many (1:n) or many to many (n:m).
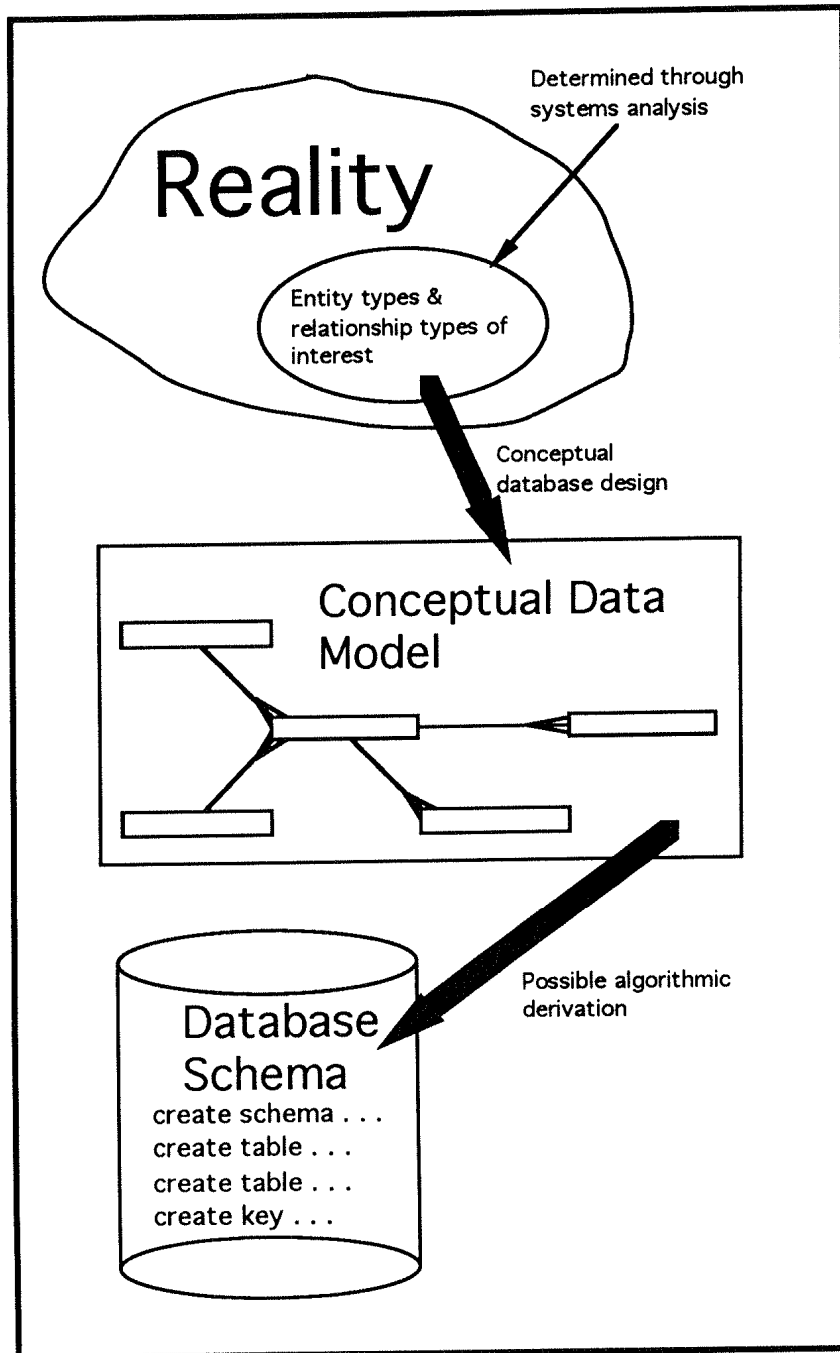
8

Figure 2: Database Design Process

While a variety of ER notations have been proposed, most are extensions of either that proposed by Chen (1976) or that proposed by Finkelstein (1989). There are others, for example the French notation called Individual Formalism (Bedard & Paquette, 1989). The notation adopted here is based on the Finkelstein notation, and is illustrated in figure 3. Entity-sets are represented by rectangular boxes and relationships by lines connecting the entity-sets. The sub-set notation

actually denotes interdependent 1:1 relationships between the super-type and each of the sub-sets. These relationships are mutually exclusive for instances of the super-type and mandatory for each of the sub-sets. Two aspects of relationships warrant comment. First, relationships are always binary. This is inherent in the choice of notation and is useful in that it simplifies the derivation of database schema. Second, relationships are almost devoid of semantics, each relationship indicating only that the two entity-sets are associated with each other. This may be seen as a limitation, but it is an advantage here because a notation rich in semantics could detract from the objective of examining issues related to the incorporation of spatial concepts in data models.

(a) A and B have a relationship of cardinality 1:1. The relationship is optional for A as indicated by the circle (i.e. an instance of A may exist without an associated instance of B.) The relationship is mandatory for B - the vertical bar indicating this.

(b) A and B have a 1:n relationship, the *crow's foot* at the B end of the relationship denoting that there may be many instances of B for each instance of A. The relationship is mandatory for A & optional for B. A m:n relationship is denoted by *crow's feet* at each end.

(c) B, C & D are sub-sets of A. The notation denotes the following conditions: each instance of A must belong to one and only one of the sub-sets. Instances of each of the sub-sets must also belong to A.
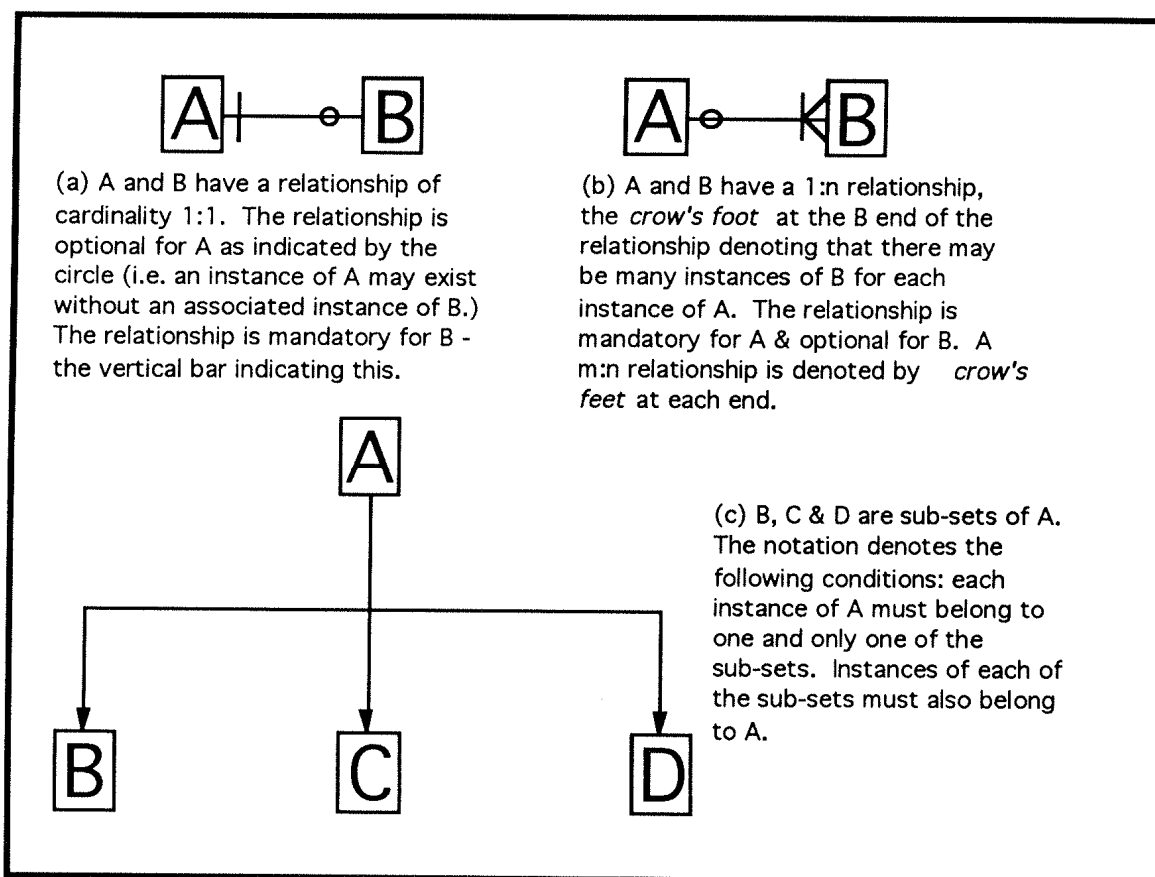
Figure 3: ER Diagramming Notation

Figure 4 illustrates the different ways in which a pair of 1:n relationships can connect three given entity-sets. Of interest for the purpose of examining the model to be developed in the next section 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 4. Figure 4(i) corresponds to a many to many relationship between entity-sets A and C, B being an associative entity-set. In Figure 4(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

1 0

other hand, in Figure 4(iii), there is no direct relationship between A and C implied. Instances of A and C in figure 4(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 4(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 map layers related to the same geographic area. This is elucidated in the following sections.
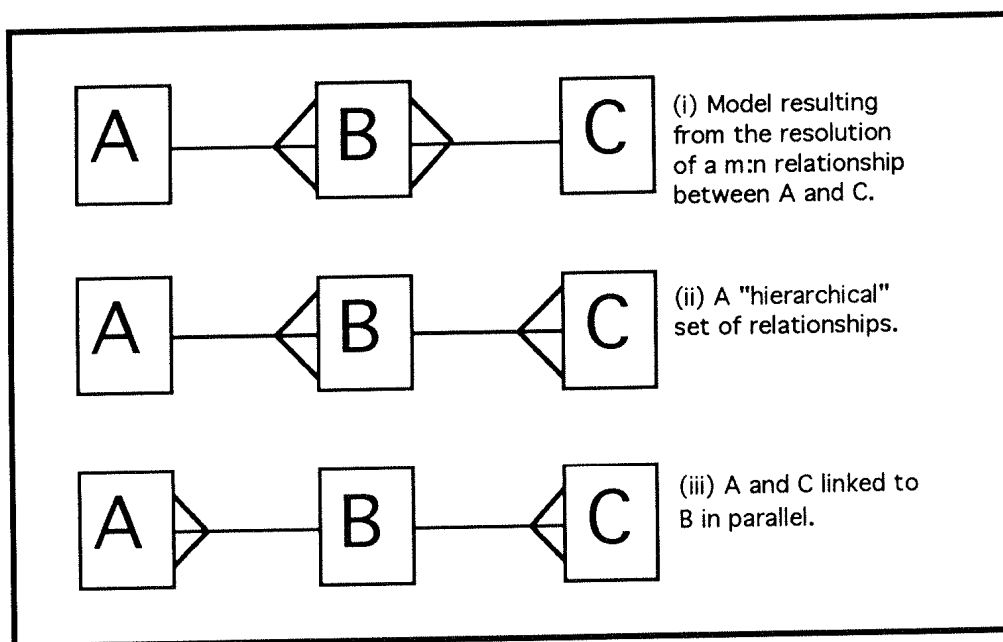

Figure 4: Generic Models of two 1:n Relationships & three Entity-sets

5 Case Study: Conceptual Modelling for a Pipe Database

This section presents a case study, first elucidating the requirements for a database pertaining to assets comprised in a city's drainage and sewerage networks, then developing an ER model for that database. The case study is based on an actual project in which the data model was required to be sufficiently flexible to be implemented as either a spatially referenced database or a non-spatial database (Firns & Read, 1989). A major objective of the project was to provide the basis for a data resource for estimating maintenance and replacement costs and for prioritising maintenance and replacement work. It was also stipulated that the pipe network data model needed to be considered in the context of a corporate (city-wide) data model. Possible relationships to entity-sets not directly associated with the pipe networks would therefore have to be considered.

5.1 Background

Engineers responsible for the pipe networks were faced with a situation where the aging network was deteriorating at such a rate that then current maintenance and replacement programmers were unable to keep up. Part of the problem was a lack of objective information meaning that key areas for maintenance were not being targeted. It was believed that with appropriate information, the resources allocated to maintenance and replacement could be applied more effectively. Initially, a data model was to be constructed for the application. For each pipe, the following data was required:

- length
- material
- diameter
- year laid
- invert-level
- repairs (date and type of each repair)
- existing defects (date observed, type, comment)
- pipe type (foul sewer, stormwater sewer, combined sewer).

The following data pertaining to manholes was also required:

- size (diameter of opening)
- material
- construction date
- condition (existing defects)
- repairs (date, type of each repair).

Other appurtenances such as mudtanks, pumps & valves had also to be represented within the database structure. The following were the relevant details:

- appurtenance type
- installation date
- condition (existing defects)
- repairs (date, type of each repair).

Environmental data was also of interest in that many environmental factors could affect either the condition of pipes or the cost of maintenance. Specific trees, for example, were of interest because their roots could damage pipes and other underground appurtenances. Also the location and nature of significant users of the networks were of interest for maintenance scheduling, or for determining likely sources of contaminants. Specifically, the following locationally referenced data was considered important:

- land use zoning
- significant users
- land stability
- level relative to sea-level
- traffic volume
- bus routes
- location of significant trees
- location of structures likely to be affected by excavations.

The possibility that the database would be implemented in a GIS had to be considered in determining ways to integrate the pipe network data with other locationally referenced data (e.g. road network data). The model developed in the following sub-section takes into account this possibility, while still being readily implementable as a relational schema.

5.2 The Data Model

Analysis of the information requirements described above led to the development of the normalised (Date,1990) ER model in figure 5. Data traditionally stored on reticulation sheets could be stored in a database conforming to this model and be extracted and combined with other data resources to produce planning information. The space available here precludes the presentation of all the attributes defined for the entity-sets shown, however primary key attributes for a possible relational implementation of the model are given in Table 1.
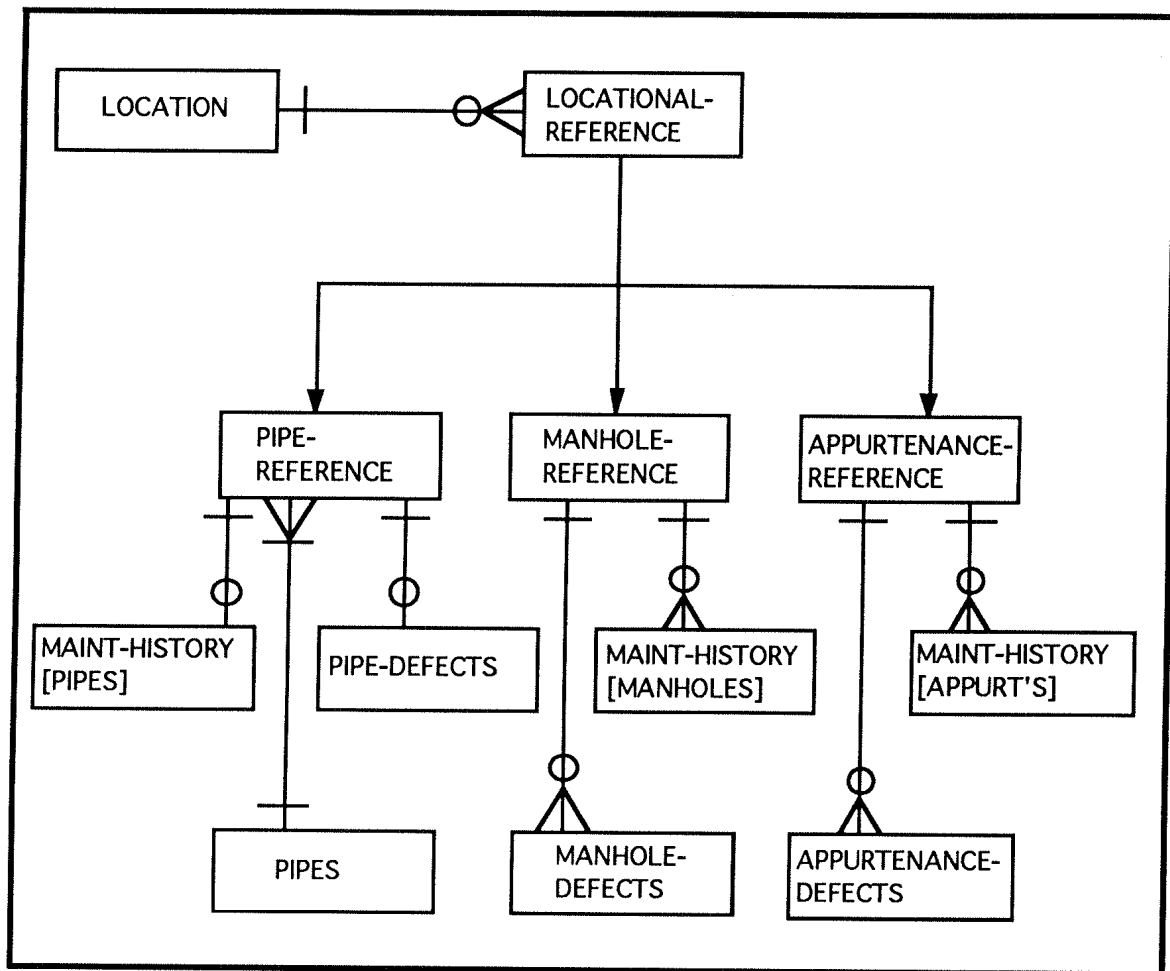
Figure 5: ER Data Model for Pipe Network

The pertinent aspect of figure 5 is the way in which locational data has been modelled. There are two entity-sets in the model associated with location. Firstly, the LOCATION entity-set with a concatenated primary key consisting of X and Y coordinates (as shown in Table 1). Other attributes of LOCATION include:

- street address.
- land use zone.
- details of any significant (industrial) users.
- traffic volume.
- whether the location is on a bus route.
- whether there are any large trees or structures in the immediate vicinity.

The attributes of LOCATION listed above are only those considered relevant to forward planning for preventive maintenance. These attributes would potentially affect the direct cost of carrying out maintenance or give some indication of the degree of disruption likely to occur if extensive excavations took place. Analysis of the entire City's information requirements would doubtless indicate many more locational attributes relevant to other divisions of the City.

ENTITY	PRIMARY KEY
Location	X/Y Coordinate pairs
Locational-reference	Ref-no.
Pipe-reference	Ref-no.
Appurtenance-reference	Ref-no.
Manhole-reference	Ref-no.

Table 1

The second entity-set associated with location is LOCATIONAL-REFERENCE, with primary key "ref-no". Instances of this entity-set would be generated to represent features of the network - a new instance being created to represent any observation or characteristic of the pipe network. Each of these would belong to one of the sub-sets, each of which also has the primary key "ref-no". The non-primary key attributes of LOCATIONAL-REFERENCE are listed below:

- X coordinate
- Y coordinate
- reticulation sheet no
- ref type

The relationship between LOCATIONAL-REFERENCE and LOCATION would, in a relational database, represented by the inclusion of X and Y coordinates (the primary key of LOCATION) as attributes of LOCATIONAL-REFERENCE. Reference to reticulation sheets was necessary as these would continue to be used as working documents in the foreseeable future. The attribute "ref-type" is used to indicate the subsets to which instances of LOCATIONAL-REFERENCE belong (i.e. each is an appurtenance reference, a pipe reference or a manhole reference).

As shown in Figure 5, there may be many locational references at a given location. This is because there may be more than one observation about the network at a location. For example, at a

manhole a locational reference may be created to represent the manhole, and locational references may be created at the same point to represent each of the pipes that enter that manhole. This would result in three or more instances of LOCATIONAL-REFERENCE with the same set of X,Y coordinates.

It is necessary to create at least two instances of LOCATIONAL-REFERENCE to represent a pipe in the database. Additional locational references would be generated to represent points at which repairs were carried out along a pipe. New references would also be generated to record data such as invert level at intermediate points along a pipe.

6 Discussion of the Asset Management System Data Model

The data model in figure 5 is adapted from a model developed for Dunedin City Council (Firns, 1989) with some detail omitted for the sake of clarity. Certain aspects of the model are particularly pertinent in the SIS environment, but first it is important to realise that this model is directly implementable in, for example, a relational database management system. Each entity-set and each relationship, together with all associated attributes could readily be implemented as a table in a relational database in the same way as any other data model developed using the same notation. This is because the standard modelling conventions have been used without augmentation to represent spatial concepts.

Consider now the entity-sets LOCATION, LOCATIONAL-REFERENCE, PIPE-REFERENCE and APPURTENANCE-REFERENCE. These have been defined in the context of the application being modelled and without special regard for the fact that the model may be implemented as a spatially referenced database. The LOCATIONAL-REFERENCE entity-set could however, correspond to the concept of a node in a topological data structure. Here, a LOCATIONAL-REFERENCE is defined as "a reference to some characteristic of the pipe network at a location". Because there may be many characteristics of interest at any given location, there is a 1:n relationship between these two entity-sets. The LOCATION entity-set could be defined by any form of location identifier: X,Y coordinate pairs for example, or X, Y, Z coordinates. Alternatively it could be defined simply on the basis of nearest street address. It could also be the means by which a relational database is tied to a spatial referencing system in an SIS environment. In this latter case, the pipe network data would form a layer in a layer based SIS.

Of particular interest with regard to the LOCATION, are its possible attributes. The list at the end of section 5.1 indicates attributes of LOCATION which would be of interest to a pipe engineer

either because of effects on the pipes themselves or because of effects on the ability of engineers to maintain the pipes. Each of the attributes in that list could form the basis for the definition of a separate layer in a spatially referenced database. Furthermore, many of these attributes could have other associated data not directly connected to the pipe network data. An important consideration then is the problem of how to represent this additional data in the data model of figure 5.

One solution is to create additional sub-sets of LOCATIONAL-REFERENCE representing the additional features, as in figure 6. In this model, new sub-sets are defined for each of the environmental factors, and each of these could be related to other entity-sets not necessarily associated with pipes. An alternative solution is given in figure 7, where, instead of creating sub-sets of LOCATIONAL-REFERENCE for each of the environmental factors, new entity-sets are linked in parallel to the LOCATION entity. It is contended that the solution in figure 7 is in fact the correct solution. Recall from section 3.0 that where a data model has parallel linkages, the entity-sets which are linked in parallel cannot be said to be directly related to each other. Instances of these entity-sets can only be related by virtue of being coincidently connected to the same instance of the entity-set at the centre of the linkage. A fundamental difference between the models in figures 6 & 7, is that in figure 7, it would be useful to explicitly represent topological relationships between instances of the sub-sets of LOCATIONAL-REFERENCE within a database structure. This is not the case with the model in figure 6, where many of the implied topological relationships would be of no consequence. The model in figure 7, therefore provides an unambiguous representation of reality. It also clearly indicates the layers that would be required in a layer oriented spatial database as well as the tables required in a relational database. The LOCATION entity-set could itself be implemented as a table in a relational database. This model therefore supports both data modelling objectives whether or not the implementation is expected to be a spatially referenced database.
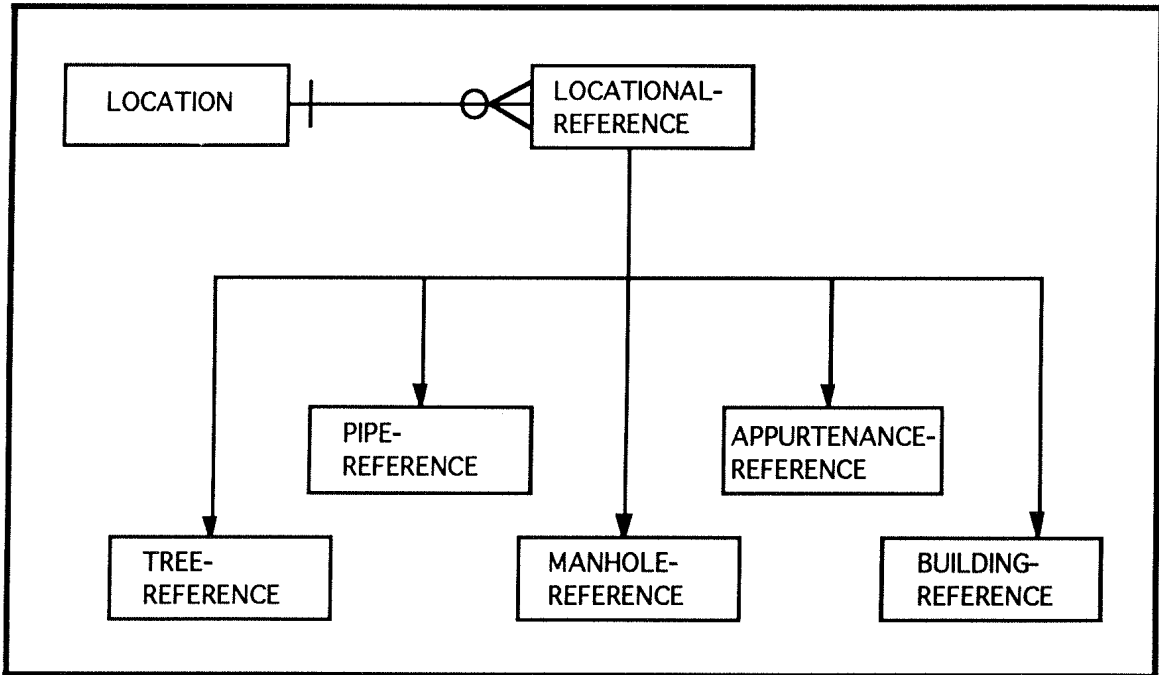
Figure 6: One Approach to Representing Additional Locationally-referenced Data
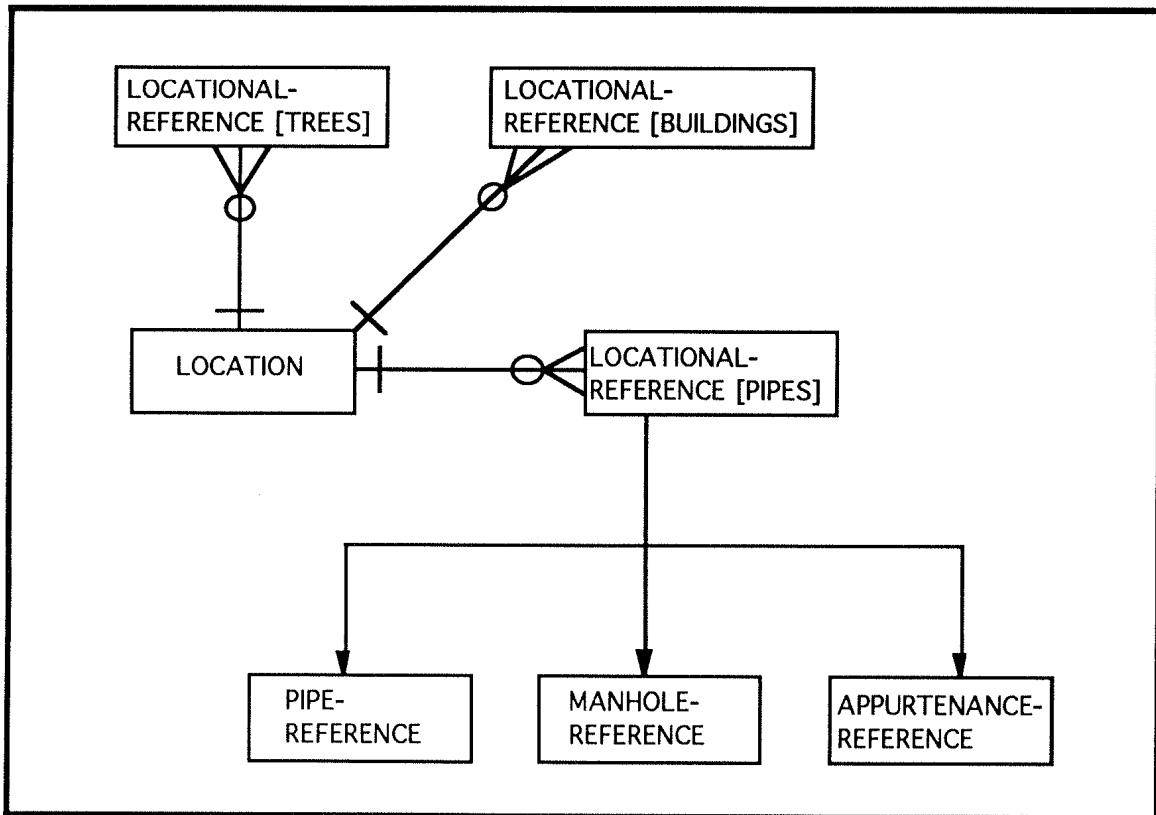in the Pipe ER Model

Figure 7: Alternative Approach to Representing Additional Locationally-referenced Data
in the Pipe ER Model

1 8

7 Conclusions and Research Direction

This paper has shown that data modelling concepts initially developed for the design of non-spatial databases are readily applicable to the design of aspects of spatially referenced databases. The models in figures 5 & 7 show that by using standard data modelling techniques, it is possible to derive the map layers required to support an application. Derivation of the map layers is based on an analogy between the notion of a map layer and that of a parallel linkage, a commonly occurring structure in entity-relationship models. Central to this result, however, is an acceptance of the definitions of the entity-sets LOCATION and LOCATIONAL-REFERENCE as discussed in section 6. The result implies that focussing attention on the perceived complexity of spatial data when developing entity-relationship models for spatial databases may not be necessary. The modelling of reality with a focus on the requirements of the application and a good understanding of the data modelling techniques remains paramount in importance. It also shows that data modelling notation rich in semantics is not necessarily a prerequisite to being able to model the perceived complexities of spatially referenced databases. An understanding of the implications of alternative models of reality is however, essential.

Ongoing research in this area is examining the potential use of semantically richer data modelling abstraction mechanisms to represent spatial relationships (e.g. relevant topological connections) between entity-sets within specific layers defined by parallel linkages to a LOCATION entity-set. The use of such abstraction mechanisms in this way may form the basis for defining spatial integrity constraints. Another avenue of research being pursued by the author is to propose a notation for representing in data models the spatial data type of an entity-set (e.g. point, line or polygon). This would reduce the potential for semantic overloading of entity-sets such as LOCATIONAL-REFERENCE and its sub-sets in the model developed here.

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 & Koffeman, 1974, 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.

Batini, C. (ed) [1989] Proceedings: 7th International Conference on The Entity-Relationship Approach, Rome, 1988 : North-Holland, Amsterdam.

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, April, 819-827.

Bracken, I. & Webster, C. [1989] "Towards a Typology of Geographical Information Systems", International Journal of Geographical Information Systems, 3 (2), 137-152.

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.

Codd, E.F. [1990] The Relational Model for Database Management (Version 2), Addison-Wesley, Reading, Massachusetts, 538p.

Date, C.J. [1990] An Introduction to Database Systems, Vol. 1 (5th ed), Addison-Wesley, Reading, Massachusetts, 854p.

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.

Finkelstein, C. [1989] An Introduction to Information Engineering: From Strategic Planning to Information Systems, Addison-Wesley, Sydney, 393p.

20

Firns, P.G. [1989] "Modelling Locational and Descriptive Data for a Pipe Network Information System". In Proceedings: URPIS 17, the 17th Annual Conference of the Australasian Urban and regional Information Systems Association, November 1989. Perth, Western Australia : Australasian Urban and Regional Information Systems Association, 359-368.

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.

Firns, P.G. & Read N.S.L. [1989] "A Data Model for A Pipe Network Information System". In Proceedings: Inaugural Colloquium of the Spatial Information Research Centre, November 1989. Dunedin : University of Otago, 52-65.

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

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

Healey, R.G. [1991] "Database Management Systems". In Maguire et al, 1991, Vol. 1, 251-67.

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

Kangassalo, H. (ed), [1991] "Entity-Relationship: The Core of Conceptual Modelling", Proceedings: 9th International Conference on The Entity-Relationship Approach, Lausanne, Switzerland, October, 1990 : North-Holland, Amsterdam.

Kennedy, G.J. [1993] "A Systematic Approach to the Specification and Evaluation of an Information Systems Development Methodology", PhD Dissertation, University of Otago, Dunedin, New Zealand, 155p.

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

Klimbie, J.W. & Koffeman, K.L. [1974] "Database Management". Proceedings of the IFIP Working Conference on Database Management, Corsica, France, April, North-Holland.

Laurini, R. [1991] "Introduction to E.R. Modelling for Urban Data Management". In Kangassalo, 1991, 475-485.

Maguire, D.J., Goodchild, M.F. & Rhind D.W. (eds.) [1991] Geographical Information Systems: Principles and Applications, Longman, London, vol. 1: 649p., vol. 2: 447p.

Maguire, D., Stickler, G. & Browning, G [1992] "Handling Complex Objects in Geo-relational GIS". In Proceedings: 5th International Symposium on Spatial Data Handling, Charleston, August 1992, IGU Commission on GIS, pp 652-661.

Morehouse, S. [1985] "ARC/INFO: A Geo-relational Model for Spatial Information". Proceedings: Auto-Carto 7, Falls Church, ASPRS, pp 388-397.

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.

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

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

Tsichritzis, D.C. & Lochovsky, F.H. [1982] Data Models, Prentice-Hall, Englewood Cliffs, New Jersey, 381p.

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.