GeoInformatica 5:1, 11–32, 2001
© 2001 Kluwer Academic Publishers. Manufactured in The Netherlands.

Making Space for Time: Issues in Space-Time Data Representation

DONNA J. PEUQUET
Department of Geography and The GeoVista Center, 302 Walker Building, The Pennsylvania State University, University Park, Pennsylvania, U.S.A. 16802
E-mail: peuquet@geog.psu.edu

Received November 19, 1999; Revised November 21, 2000; Accepted November 22, 2000

Abstract

Even with much activity over the past decade, including organized efforts on both sides of the Atlantic, the representation of both space and time in digital databases is still problematic and functional space-time systems have not gone beyond the limited prototype stage. Why is this the case? Why did it take twenty years from the first GIS for the for representation and analysis in the temporal, as well as the spatial dimension, to begin?

I explore the answers to these questions by giving a historical overview of the development of space-time representation in the geographic information systems and database communities and a review of the most recent research. Within the context of this perspective, I also question what seems to be a spirit of self-accusation in which the lack of functional space-time systems has been discussed in the literature and in meetings of GIS researchers. I close by offering my own interpretation of current research issues on space-time data models and languages.

Keywords: spatio-temporal representation, temporal GIS, temporal databases

1. Introduction

There is a scene from *The Hobbit* [78] that has haunted me for a long time. In the story, the central character, Bilbo Baggins, found himself wandering lost in a cave. He came upon a strange and slippery creature, Gollum, who offered to show him the way out if he can correctly answer a series of riddles. If not, Bilbo risked a high probability of being eaten by the creature. The last, and most difficult riddle the creature asks goes as follows:

This thing all things devours:
Birds, beasts, trees, flowers;
Gnaws iron, bites steel;
Grinds hard stones to meal;
Slays king, ruins town,
And beats high mountain down.

What is it?

The answer is already given away by the topic of this paper; time. There are several

issues incorporated into this story and the riddle that are central to the task of representing time, and space, in a digital context. For a broad range of professionals who study environmental, social and economic issues, the term spatial data refers to data where each entity represented has associated with it a location or series of locations on, in or above the surface of the earth. Space-time data, of course, also includes a temporal coordinate. The representation of time and of space are tightly interrelated, both historically and theoretically.

I can in many ways identify with Bilbo Baggins when confronted by Gollum. Any researcher who has grappled with space-time representation issues—from the most theoretical to the most applied—has discovered how difficult a task this is. It often does feel like encountering a strange creature in the dark and then having it pose a series of riddles. The representation of space and the representation of time are each difficult problems on their own. What progress has been made and what remains to be done in the more difficult problem of space-time representation? In order to try to answer these questions, I will first give a historical sketch, reviewing some of the main issues and events. I will then briefly review the most recent research within this context, and close with some of my own views on what this all means for the future.

2. From space to space-time representation

2.1. Historical context

For millennia, spatial data at geographic scales have been most commonly represented in the form of maps. The first known map that was identifiable as such was scratched on a piece of mammoth bone found at an ice age camp site on the river Dneper [40]. It depicts a human community of 15,000 years ago. So, for a very long time, maps have been used as fairly convenient, efficient and effective representations of our environment or portions of it. Nevertheless, the limitations of the available media (bone, stone, clay, paper, etc.) dictated a static representation.

The first representation of geographic space for use with a computer is generally attributed to Waldo Tobler, who created what he called ''computer maps'' in the late 1950s [77]. The first computer maps were usually in gridded, or raster format, simply because the only computer output device commonly available at the time was the line printer. Thus, the map was quantized into a matrix of cells, so that the map could be printed a row at a time in order to build up the entire image. The earliest geographic information systems either utilized this same grid representational format, or the points, lines and polygons on the map were copied, line for line and point for point using what became known as the vector format. The U.S. Census Bureau developed the DIME (dual independent map encoding) format in 1968 in order to provide automated address encoding for the 1970 Census so that mail out/back could replace the field enumerators. This data model was a major advancement in that for the first time, topology was explicitly incorporated as an integral component of a vector representation.

The topological vector format became the spatial data model of choice during the 1970s.

Spatial data handling systems could generally be described as being monolithic in their database design, being either vector-, or more infrequently, raster-based. As an example of the degree of this separation in the use of these two data models, ESRI marketed two separate GISs (PIOS and GRID), which were vector-based and grid-based, respectively. The vector model was preferred for digital representation and manipulation of geographic data from a conceptual standpoint, expressly because of the direct correspondence to the symbology of points, lines and polygons portrayed on maps. This was the representational paradigm that everyone—researchers, developers and users—were familiar with. This, however, also perpetuated the representation of static space in geographic databases.

It was also during the 1970s, however, that Relational model of RDBMS had first become recognized as a potential alternative for spatial data representation [87], but functional GIS and DBMS systems remained separate, as did development in these respective areas.

In the fall of 1977 the Harvard Laboratory for Computer Graphics and Spatial Analysis conducted the first International Advanced Symposium on Topological Data Structures for Geographic Information Systems. The purpose of that symposium was to define the then current state-of-the-art in computer data structures as applied to GIS, discuss the results of recent research involving the design and implementation and to discuss how to effectively deal with the increasing amount and, more importantly, the increasingly wide variety of digital geographic data that various U.S. government agencies were building or about to build. There was also an initial probing into the relationship between vector, grid and relational data models; that the selection of data model should be conditioned by what one conceives are the properties of the phenomena being represented. Nevertheless, such ''philosophical'' discussions were rare.

2.2. *Initial work on representing time in databases*

The 1980s saw work on the addition of temporal capabilities in database management systems. This was seen as an important capability for transaction systems, such as banking and medical systems where details of past history, as well as the preservation of the individual changes through time, is of critical importance. The representation of past, as well as present, information became possible through rapidly declining storage costs and increasing capacity, and the success of DBMS technology in handling non-temporal databases. Most of the work in temporal DBMS focused upon extension of the Relational model, [7], [74], [75], although there was also some work utilizing hierarchical [68] object-oriented [2], knowledge-based [17] and semantic [83] data modeling approaches. These efforts within the DBMS community were summarized in Roddick and Patrick [67] and in Tansel [76]. Chomicki and Revesz [13] note that there were over 12 different extensions to the relational model for (non-spatial) temporal data discussed in Tansel [76].

There was also a significant amount of work in temporal reasoning following Allen's initial insights [4], [5], [27], [55]. Work on a standardized temporal database query language resulted in TSQL2 [72]. Within the DBMS community, the effort to establish the concepts developed within the context of this language as a new part to SQL3, named

SQL/Temporal, as an international standard continue to progress. Cheng and Gadia [12] made an early effort to extend SQL and SQL-based language for temporal queries to include the spatial as well as temporal dimensions.

Around this same time, gaining a better understanding of the effects of human activities on the natural environment at global, regional and local scales was viewed with increasing urgency. The construction of computerized databases began as a coordinated and shared effort in order to study phenomena such as climate dynamics, ocean dynamics and global warming, encroachments on animal habitats, and the effects of water pollution on aquatic life. Urban planners came to depend on computer databases in managing urban growth and its impacts. Researchers in these realms also shifted their attention from observation and inventory toward integrated and normative approaches that seek to model and predict space-time processes. The availability of remotely-sensed satellite data as well as the exponential accumulation of other geographic observations through successive times enabled researchers to undertake detailed empirical studies of complex spatiotemporal processes at multiple geographic scales.

In order to advance such empirical studies, a simple and pragmatic solution was used. Temporal data could be stored in existing GIS using a ''snapshot'' data model. This type of representation consists simply of a differing use of one of the two traditional types of spatial data representation used in existing GIS; a sequence of spatially-registered grids. But instead of a single gridded file representing a complete thematic map layer as in a static (i.e., non-temporal) spatial database, each gridded image represents a ''world state'' relative to a given thematic domain, storing a complete image, or ''snapshot,'' at a known point in time. The snapshot model for representing spatio-temporal data can thus be viewed as a 3-D (when representing planar space) or (as in the case of topographic data) a 4-D space-time cube. Each cell within a separate snapshot contains the value for the corresponding location at that time. Although it introduces a large degree of redundancy in the data when the phenomenon of interest is not constantly changing everywhere, this approach is conceptually straightforward, and the entire ''world state'' at any stored space-time location can easily be retrieved. This ''snapshot'' approach also reflects how data from commonly available sources are gathered (e.g., satellite imagery and census data). The snapshot approach therefore was recognized early-on as a convenient, although often highly redundant means of spatio-temporal representation that is location-based.

Langran's groundbreaking work on the extension of GIS to explicitly represent and analyze temporal, as well as spatial, data in the late 1980s [51] was prompted by some of the same motivation that spurred the initial work in Temporal DBMS. Certainly, the feasibility by this time to store the much larger volumes of data necessitated by the addition of the temporal dimension was one motivation. The work and successes in the Temporal DBMS realm was another motivation for her work. But by far, the most important motivation generally for the inclusion of time in geographic databases was, and still is, the need for such a tool to aid in the development of a greater understanding of geographic processes, through time as well as space, and the need to effectively access the increasing amount of available space-time data.

By the beginning of the 1990s, there were a number of efforts to build geographic databases and prototype systems specifically intended for space-time data. Some of these

utilized extended DBMS [5], [29], [52], [64]. A number of approaches for space-time representation proposed within the GIS community utilized direct extensions of vector, as well as raster, spatial models [41], [47], [50]. Increasing attention was given to the use of the object-oriented approach as a means of extending conceptually vector or raster approaches [42], [65], [84], [88].

These have provided pragmatic solutions for some application areas, but was seen to fall short of the longer-term need. Extending the Relational model has been shown to be an effective solution for temporal (as well as bitemporal) data when the temporal dimension is conceptually linear in form. Extending the Relational model to include space as well as time has led to implementational forms that are both complex and voluminous. Extending the vector model as used in GIS suffers from the same shortcomings. Extending the raster model in the form of temporal ''snapshot'' sequences maintains the simplicity of the raster model, but at the expense of a considerable amount of redundant data storage. Object-oriented techniques have most often been used to simply implement conceptual raster and vector models, most commonly through the use of a relational database system. This falls short of use the true power that object-oriented modeling can provide for space-time data representation. Experience so-far has shown that simply extending a spatial data model to include temporal data, or vice versa, will result in inflexible and inefficient representations for space-time data [58], [69].

Because of the growing difficulties with performance and flexibility of traditional representation methods in existing GIS, the need for increased attention on a theoretical and abstract level became recognized. The COSIT (International Conference on Spatial Information Theory) series of conferences, which has been held bi-annually since 1993 evidences such recognition.

3. Recent research

For most of the work on space-time database representation historically, efforts progressed separately and in parallel within the DBMS and GIS research communities. It has only been within the last few years that there has been much interaction or synergy. This was perhaps inspired by funded interdisciplinary and multi-institutional projects to establish new linkages among researchers, particularly CHOROCHRONOS Project in Europe and the NCGIA Varenius Project in North America. Part of this change has also precipitated a number of recent review papers (the current paper included) that are intended to help provide an integrated view derived from both communities and to delineate current research issues [1], [20], [69].

As already delineated in the papers just cited, evolving research on space-time representation has focused on a number of specific areas, including; (a) the ontology of space and time, (b) development of efficient and robust space-time database models and languages, (c) inexactness and scaling issues, (d) graphical user interfaces and query optimization, and (e) indexing techniques for space-time databases. Given the need for a more complete theoretical foundation, the remainder of this paper will focus on the areas

more centrally related to foundational issues, particularly the first three just listed, rather than issues related more directly to implementation.

3.1. Ontology of space and time

The ontology of space-time deals with the nature of space, of time, and the unique interactions of space-time. Investigation at this highly abstract level is essential in order to develop a common conceptual framework. Issues specific to time; of linear vs. cyclic views of time, multiple times (world time, valid time, user-defined time), continuous or discrete change, branching time and alternative timelines, etc., seem to be recognized and reasonably-well understood in a database context.

A growing number of researchers in both the DBMS and GIS communities have come to the realization that a general, application-independent, solution that allows an optimal combination of simplicity, flexibility and efficiency will require rethinking at an abstract level, and new types of implementational data models and associated query languages. This approach needs to be based upon a uniform ontological framework and requires a multi-representational approach. Peuquet [58], and Yuan [90] had proposed similar solutions from differing perspectives early-on in this process. Peuquet suggested a general ontological framework, and used that to derive a Triad (what/when/where) conceptual model. Yuan suggested a what/when/where framework, deriving hers from an application-driven perspective.

The nature of concepts that are inherently both spatial and temporal (e.g., motion) and the interactions of the spatial and temporal dimensions need to be better understood in order to develop truly effective and robust space-time data structures, query languages and user interfaces. Part of the fundamental difficulty that needs to be overcome, as discussed by Parent et al. [56], is that there is an increasing degree of mismatch between the nature of the data as seen conceptually from an application-oriented view, and implementation-oriented views used in both the DBMS and GIS communities.

The means to a better ''fit'' is in part being provided by the object-oriented approach, which supports flexible constructs such as abstract data types (ADT's). However, there remains a lack of a common abstract theoretical framework that explicitly describes the fundamental characteristics information in a context-independent way from which to draw application- and user-oriented views. A focusing of attention in the research community to these issues is evidenced by the focus of some recent meetings, such as the Ontology of Fields meeting held in Bar Harbor, Maine, sponsored by the Varenious Project in the U.S. in 1998 and the GIScience 2000 meeting held in Savannah, Georgia in October, 2000.

3.2. Space-time database models and languages

Within the database community, there has been a significant effort in database models and operators for applications involving moving objects [21], [22], [38]. Current temporal

DBMS can handle discrete change, but continuous change can only be implied from one recorded state to the next. ''Traditional'' temporal DBMS applications (e.g., banking) most often needed to deal only with discrete change (transactions). However, most geographic spatio-temporal applications involve continuous change of objects in space. One commonly cited exception to this is the case of cadastral systems, where property boundary changes are recorded in discrete transactions.

There are two basic types of spatial change relating to objects; they can (a) change location in space, i.e., move, and they can (b) change in spatial extent by growing or shrinking and by changing shape. Depending on the application, change can consist of movement of point objects, which have no notable areal extent. This might include applications such as the tracking of delivery vans, taxi cabs, or military vehicles. Movement-only change conceptually also includes the movement of rigid areas through space.

Because of the many time-critical commercial and military applications, such as those just mentioned, it is this first type of spatial change relating to objects that has received the most attention. Since the complexity of representing points in space-time is not as great as that for representing objects, which can themselves change in spatial configuration, a number of specialized moving-objects database (MOD) systems, not utilizing relational DBMS, are already available commercially. The Omnitracs system developed by Qualcomm for the transportation industry is one example [63]. Although this gets around the immediate representational focus of current temporal DBMS on representing discrete change, this does not get around the fundamental issue that computers are discrete devices.

With geographic-scale applications, it is usually the case that for non-point objects the spatial extent is changing at the same time as the location (e.g., storm systems, vegetation areas, land use). Based on an overall framework proposed in Güting et al. [38], Erwig et al. [22] detail a conceptual-level view (or in the authors' terms, an abstract-level view) based on objects with observed or projected coordinates (x, y, t), as attribute sets, in 3-D space. This is used as the basis of defining abstract data types with associated functions that can in turn can be implemented as data blades, cartridges, etc. for extensible DBMSs.

This has distinct parallels in overall approach to earlier work by Worboys [89], who describes a conceptual model of space-time objects as ''spatio-bitemporal complexes.'' (Worboy's model is bitemporal in the sense that both valid and database time are represented.) However, the earlier work of Worboys focuses upon the representation of discrete changes in the topology of a network of points, lines and areas within a relational DBMS framework. This at a conceptual level represents an extension of the topological vector format popular in GIS from planar space to (2-D) space + two temporal dimensions.

In contrast to representing topological elements, which in themselves do not represent items that people commonly utilize conceptually, Erwig focuses upon the representation of spatial objects, as real-world ''things'' (e.g., roads, rivers, lakes), as well as their interrelationships. These objects, in turn, may be wholistically viewed as point, linear or areal in nature as they continuously move and change in configuration through space-time. Thus, in the representation discussed by Erwig, movement is explicitly stored as trajectory vectors in 3-D space for coherent objects.

In subsequent work building upon the abstract model given in Erwig et al. [22], Florizzi et al. [25] describe an implementation-level representation for continuously moving objects. This is a high-level specification of data structures that can be implemented in a temporal DBMS. Because continuous time must still be represented discretely on an implementation level, they utilize a set of temporal units, called slices, that represent the temporal intervals for each data type. This allows for the temporal granularity and individual temporal intervals to be individually selected for each data type.

Within the GIS community, work on space-time data models for geographic phenomena has always been viewed from within a spatial context and emphasizing spatial change, and utilize either an object-based (as is currently the more frequent view in research) or field-based view. Certainly, this is in large part due to the standard and traditional object/field dichotomy utilized in GIS data modeling. This dichotomy, however, evolved for static representation of geographic objects—before the issue of representing change and dynamics had arisen. The object-based and field-based views are best organized for performing object-based and location-based queries, respectively. This dichotomy ignores the third element of what/where/when in also providing for a representational perspective that is time-based, needed in order to explicitly represent evolution of entities and locations, and the interrelationships of these through time. Certainly, temporal information can be captured within field-based and object-based representations, but extracting specific temporal interrelationships, such as relative temporal coexistence of specific entities or the relative temporal configuration of various events, becomes computationally taxing.

Such a time-based conceptual-level model was proposed by Peuquet and Duan [59] in ESTDM (the event-based spatiotemporal data model). This model records the timestamp for any change and associated details describing each specific change in temporal order. The sequence of timestamps can thus be viewed as a timeline, or a compacted temporal (1-D) vector corresponding to the forward progression of time. New events that occur as time progresses are added to the end of the list. The time line is compacted in that temporal locations where no change occurs are not recorded. It is assumed that each event list and associated changes relate to a single thematic domain (e.g., land use or population), although these stored events can relate to a change occurring at either a location or set of locations or to an object or set of objects. A similar model, with associated query operators, was described by Claramunt and Thériault [15], [16].

More recently, Hornsby and Egenhofer [43] describe a visual language for the explicit description of change relating to objects, called the change description language. This work focuses on the appearance/disappearance of entities and particularly of the identity transitions from one object to another that are possible. Identity is seen as a means of tracking and querying the existence of specific objects and types of objects independent of specific attribute values. Bonhomme and her colleagues [8] extended the Lvis language, originally designed for visual query of spatial data [6], to provide spatio-temporal visual query capabilities. These visual queries are implemented by translation into an SQL language with spatio-temporal extensions.

Raza and Kainz [66] propose a conceptual data model that also represents an extension to Worboy's model, but utilizes the Object-Oriented design methodology to conceptually

separate out space, time and attributes, and consolidate them in a modular fashion as distinct elements.

In addition to these abstract models, a number of researchers have also recently proposed a number of modeling approaches—structured methodologies for deriving models that are tailored to specific applications. Such efforts include that of Parent and her colleagues [56] who propose a spatio-temporal modeling technique called MADS (modeling of application data with spatio-temporal features) for bridging the application-view–implementation-view gap. Tryfona and Jensen [80], [81] propose an extension of the E-R technique, which they call STER (spatio-temporal ER). Their technique utilizes a set of very basic ontological principles dealing with the unique aspects of the space-time domain; that properties and interrelationships of objects can be defined on the basis of either discrete or continuous change, and for either of the two basic types of spatial change as already mentioned (i.e., movement, change of size/shape, and change of non-spatial properties such as temperature). They also offer an extension of the object modeling technique (OMT) utilizing the same approach [82].

A number of other recent efforts have utilized the constraint paradigm for developing conceptual-level space-time representations and languages. One example is DEDALE [37]. The basic idea of this approach is to represent collections of points in n-dimensional space using constraints expressed in a first-order mathematical language. It is viewed as an extension of the relational model in that it allows the generalized representation of potentially infinite collections of tuples [46], [48]. Chomicki and Revesz [13] have also proposed constraint databases as a common language for facilitating interoperability of spatio-temporal databases.

3.3. Inexactness and scaling issues

There is inherent inexactness built into all spatial, temporal and spatio-temporal databases. This derives from multiple factors. One is the problem of the artificial discretization of what are most often continuous phenomena required for computer representation. In the case of spatial data, this was compounded by the cartographic tradition of drawing crisp lines to define geographical entities on maps. In the early days of GIS, the development of the vector data model as one of the standard representational schemes was derived directly from the idea of (crisp) lines on maps. However, this continuity undoubtably also facilitated the development of spatial databases by means of digitizing maps, as sources of archival geographic data. Another factor is that specific data classifications and object definitions used within any given spatial, temporal or spatio-temporal databases must be defined within an application-dependent semantic framework [10], [54].

Developing means of handling the inherent inexactness in spatial databases has long been an active area of research within the GIS community [24]. Handling inexactness in temporal databases attracted much less attention because, as previously stated, this area had developed historically focusing upon application contexts that involve discrete, transaction, applications (e.g., banking and inventory). It has only been very recently that

any work on inexactness involving the combined effects of uncertainty in the spatial and temporal dimensions has appeared.

There are two basic kinds of uncertainty that are often distinguished in spatiotemporal databases and applications; fuzziness and uncertainty. Fuzziness concerns inherent imprecision. This includes situations where boundaries between things or classes of things are indistinct, for example the spatial boundary between areas of dominant vegetation types. Boundaries can also be fuzzy in the temporal dimension. In the example of vegetation types, there is no sharply-defined instant in time under usual circumstances when an area changes from one vegetation type to another. This tends to be a gradual change. Of course, the sharpness of such distinctions also depends upon scale. There is also often fuzziness in categorization. The distinction between what may be classified as ''urban'' vs. ''rural,'' for example, depends upon context and is often a personal judgement call. In contrast to fuzziness, uncertainty concerns the lack of information; something that is not exactly known. For example, the exact location of a specific truck at a given moment (e.g., somewhere between New York and Philadelphia). These types of inexactness also have combined spatial, temporal and attribute effects.

The popular snapshot representation introduces temporal, as well as spatial uncertainty into the data by nature of the representation itself. The actual changes that occurred at locations between recorded temporal locations are not explicitly stored. These can only be derived by comparing the pixel value differences between successive snapshots. For discrete change, the exact time that the change occurred is not recorded. For continuous change, the trajectory of values between two snapshots, or the exact time that some change began, is not known. It may also be that some changes are completely missed because of the particular temporal resolution used. Even when change is recorded in an asynchronous manner and represented correspondingly, for example in a timeline, changes that actually occur may not be recorded.

This problem is an extension in the temporal dimension of a classical problem within geography and other fields that deal with spatial data known as the modifiable areal unit problem (MAUP) [26], [35]. This problem has two aspects; scale and zonation. The scale aspect of the MAUP refers to the variation in results obtained based on the size of the size of the elemental unit. As recorded values are aggregated over larger and fewer units, variation is also lost. The zonation aspect of the MAUP refers to the variation in results obtained depending upon how the units are overlaid onto the space at the same scale. A slight shift can produce very different results. The MAUP problem exists for both discrete and gradual change over space, time, or space-time.

A variety of stochastic models have traditionally been applied to controlling or estimating the inexactness in spatial databases. However, these assume randomness in the variation of recorded values, which can limit their representational power. Methods that have been developed using this approach, as well as others, are described in Goodchild and Gopal [31].

For bringing specific knowledge about the behavior of the phenomenon represented to bear, the probability vector approach has most commonly been used for representing uncertainty in spatial boundaries [32], [53]. In this case, the steepness of the transition gradient from one area to another can be represented by the length and sequence of

probability values contained within the probability vector. As noted by [32], a similar approach has also been used in GIS for dealing with the MAUP in assigning values to spatial pixels containing mixed classifications [73], [79]. For example, a single pixel in a land use map may be assigned a vector of values designating that 60% of the area in the pixel is residential, 30% is commercial and 10% is parkland. In temporal databases, the probability vector approach was also used by Dyreson and Snodgrass [19] in an extended tuple-timestamped relational model to note potential times when the event might have occurred.

Another approach utilizes fuzzy set theory [91]. The fuzzy set representation provides a means to express degree of membership rather than probability of membership in a particular class, or set. This approach has been used in dealing with classification of specific locations or specific objects [9], [11], [54], [85]. Wang [86] also described a means of extending the relational model for this purpose for representing fuzzy category membership. The use of rough sets, another extension of classical set theory, has been recently proposed [3]. Rough sets are represented by the upper and lower approximation of an uncertain set. In using rough set theory, locations or objects can be designated as either being definitely within a given class (between the upper and lower bounds for the class), definitely outside, or possibly inside the class. A rule-based approach for handling uncertainty in spatial databases has also been proposed [23].

With all of the various approaches proposed, particularly within the realm of spatial databases, there are still no truly workable means developed so-far for handling uncertainty in either spatial or temporal databases [44]. There have been very few researchers who have addressed the problem of handling uncertainty in spatio-temporal databases. The work of Shibasaki [70] represents an early attempt and describes some possible approaches at a general conceptual level. Work in this area has picked up only in the past couple of years. Pfoser and Jensen [60] describe the use of interpolation for moving point objects. Pfoser and Tryfona [61] explore the general conceptual differences in the use of fuzzy set theory and probability theory in the spatio-temporal domain.

On a level of practical implementation, linear interpolation has been the predominant approach used for dealing with inexactness in existing space-time databases. Recently, Dragicevic and Marceau [18] describe a fuzzy set approach that can be used to generate snapshots at times intermediate to those explicitly recorded within the database if additional information about the overall nature of change is available from other sources. Using the likely trajectories of gradually progressing from one class to another (e.g., from forest to agricultural to urban), the degree of membership in a specific class at a particular space-time location can be calculated on the basis of fuzzy set membership.

4. Are we simply taking our time?

4.1. Needs vs. capabilities

The scientific study of processes through time as well as over space is neither new or unique. Perhaps the best-known efforts within the field of geography that made explicit use

of time as a variable in the study of spatial processes are Hägerstrand's models of diffusion and time geography [39], [57], [62]. Diffusion models focus on the overall pattern of specific natural or cultural phenomena as change spreads through space over the passage of time. The ''theory of diffusion'' has been applied to a diverse range topics, including agricultural innovation, the spread of political unrest, and the spread of AIDS [33], [34], [57].

Moreover, the concept of time is intuitive. We use it constantly in our daily lives. The passage of time is normally understood via changes we perceive occurring to objects in space—their transformations over time and their movements in relation to one another. Inclusion of the temporal element in the data is required, in everyday life as well as in scientific study, in order to derive cause-and-effect relationships and ultimately an understanding of the nature and structure of the various elements in the world around us [36].

Even with much activity over the past decade, including organized efforts on both sides of the Atlantic, the representation of both space and time in digital databases is still problematic, and functional space-time systems have not gone beyond the limited prototype stage. Why is this the case? Why did it take twenty years from the first GIS for the representation and analysis in the temporal, as well as the spatial dimension, to begin?

4.2. The issues

For both of these questions, there are two simple answers, arising to fundamental research issues still outstanding. I have already alluded to both of these. One is that we had been stuck in a traditional, static cartographic paradigm. Indeed, spatial representations on a conceptual level in both the GIS and the DBMS realms are for the most part direct interpretations of their cartographic progenitors, either as points, lines and polygons, or as cells in a Cartesian grid. Such an approach is now seen to have intrinsic limitations for representing space-time dynamics. Another reason is that historically there has been a general emphasis on the short-term and implementation-oriented solution [30]. I would now like to investigate these reasons a little more by posing two obvious follow-on questions: First, should we have expected things to have evolved any differently? Second, would we be further-along today if a top-down approach had been followed from the beginning?

The use historically of the traditional, static cartographic paradigm and an implementation-oriented approach are readily seen in retrospect. However, I question what seems to be a spirit of self-accusation in which they have been discussed in the literature and in meetings of GIS researchers. The implication that has gone along with citing these reasons is that both the GIS and the DBMS research communities have been short-sighted. We have gone from vectors and rasters in the 1960s and 1970s to entities and fields in the 1990s. This reflects a shift in emphasis from practice to theory instead of starting with a theoretically- and conceptually-informed framework from the beginning. But could we have started with entities and fields in the 1960s and 1970s?

First, there is the obvious issue of the maturity of computing and data modeling

techniques. It is undeniable that hardware capacity and software tools and techniques have been limited and primitive relative to the tasks to be accomplished. It could be argued that this alone would force an emphasis on implementation. As discussed previously, the early developments in geographic databases were driven by public agencies, who had immediate and specific needs for computer-based geographic databases. On the other hand, it was only through the substantial budgets afforded to these efforts that made initial advancements in geographic databases possible. This included the generation of new theoretical (and hardware) advancements when necessary to accomplish a specific task. The integration of topology in the development of DIME by the U.S. Census Bureau is a prime example.

Perhaps more importantly, ''jumping-in'' and using the existing and familiar paradigm for getting started in the development of any new field is a very rational approach. It allows things to be tried out in a new context as quickly as possible in order to see what works and what doesn't, as well as to allow previously unseen problems to reveal themselves. The automobile was originally called the ''horseless carriage'' and used a chassis very similar to that of the buggy—with the same wheels and suspension. It is also the usual progression in science in acquiring knowledge in any new area to begin at the lowest level, trying things out empirically, gathering observational data. From this some generalities can be made in an inductive manner. The last stage is to be able to fit past experiences into an overarching theoretical context. Leaps tend to not be made until something is noticed that no longer works or doesn't fit. The development of helicentrism after the discovery of retrograde motion is a case in point.

There were articles on the larger theoretical context, some on a very philosophical level, very early-on [14], [71] but these were only occasional contributions of individuals that didn't gain much attention at the time. The initial work on the integration of time into geographic databases by Langran and others was over ten years ago. These were in large part spurred by the need to explicitly look at dynamics over space, but we are still at the stage of using the snapshot approach as a means of ''fooling'' current systems. Proposed new representations for space-time data remain simple extensions of static, spatially-based approaches that often have major shortcomings. Why is it taking so long to develop effective space-time representations? Again, we seem to be learning by trying and observing.

The incorporation of time as an extension of raster and vector data models—and on a more abstract level, that time is just another dimension—is seen by many as the obvious solution for representation of spatial change. Moreover, such an approach appears to be supported by the mathematical tradition of the Minkowski framework in which time is represented as an additional dimension or axis in a 4-D space-time visual geometric, that is, x, y, z, t, with t representing time in a hypercube coordinate volume space [28]. The combined treatment of space and time also seems supported by natural language (at least in the case of English), in which time tends to be expressed in spatial terms [49], e.g., in the phrases: ''We are *close* to the end of the year,'' or ''There is going to be trouble *down the road*.''

It is not surprising, therefore, that the problem of representing spatiotemporal data in GIS have been regarded as one of implementation. Nevertheless, we are discovering

limitations with this approach, and it is these limitations now seen at a very high level of abstraction, with a better understanding gained through experience, that we are adding to our efforts to derive a united theory for space-time representation. The view of time as equivalent to space really only began with Minkowski and by Einstein. Time had previously been considered separate from space, and this upon some reflection makes sense. Although space and time share many characteristics, there are also a number of critical differences. For example, in real-world experience, it is possible to move forward and backward (as well a sideways) in space. It is only possible to progress forward in time in terms of experience (e.g., September 1, 1999 will only happen once). Historical recollection is the only sense in which a location in time can be returned to. Time and space also require different measurement schemes. It usually does not make sense to measure time in meters or feet.

5. Foundation for a theoretical framework for spatio-temporal modeling

There is a fundamental theoretical framework that has begun to emerge over the past ten years or so that can serve as a robust basis for moving forward; the notion of the discrete vs. the continuous view. These can be briefly defined as follows:

- In the ''discrete view,'' (or entity-based view), distinct entities, such as a lake, a road or a parcel of land are the basis of the representation. Their spatial and temporal extents are denoted as attributes attached to these entities.
- In the ''continuous view'' (or field-based view), the basis of the representation is space and/or time. Individual objects are denoted as attributes attached to a given location in space-time. Using land ownership information as an example, the particular parcel number would be an attribute of the entire space it occupies, with locations denoted in some continuous coordinate field.

For both of these views, there may also be other types of attributes (e.g., a lake may have additional information describing concentrations of specific pollutants, etc.). Individual entities or locations—as the fundamental element of these the views, respectively, may have various interrelationships stored as attributes. These may include is-a and part-of relationships for the object-based view that in turn form conceptual taxonomic and partonomic hierarchies, respectively. For the field-based view, part-of relationships would form a scale hierarchy. Given these characteristics, object-oriented data modeling techniques seem particularly well-suited for specific implementations of this representational framework.

Moreover, just as the discrete view can be applied to entities in space (static entities) or in space-time (dynamic entities), it can also be applied to events, as entities that have a temporal but perhaps no spatial extent. Examples of such purely temporal events would be a bankruptcy or an election. Events that occur in space and time would include an earthquake or a storm. Whether the temporal (or spatial) extent of any object is a point or

some interval is dependant upon the temporal (or spatial) scale being used to record the data.

These two views are in fact duals of each other: The same spatio-temporal phenomenon can be described as either discrete or continuous. Perception is indeed a very personal matter, and preference for one view or the other is often dependent on the particular task at hand. Therefore data models that are used by humans (i.e., conceptual models) should provide the choice between the two views, as dual models within the database.

It is important to note that the physical model used by the system for implementation can be different from the one presented through the users' interface. For instance, a system supporting a continuous view would need to break space and time into discrete units, for the simple fact that computers store information in discrete units. These can still be presented graphically as continuous fields.

5.1. The discrete view

In this view, a spatial value or set of values is associated with each spatial entity, which describes its spatial extent. Similarly, a temporal value or set of temporal values, usually called the lifespan of the entity, is associated with each temporal entity which describes its temporal extent. A set of associated specific operations are also defined for these objects, some of which are associated with its spatio-temporal configuration type, like point, area, instant or time interval.

In a discrete view, spatio-temporal entities are perceived as linked via spatio-temporal relationships. For instance, roads go through towns, two cities are 50 km from each other, one harvest yield in a given agricultural district succeeds a previous one. Spatial relationships needed for database operations are the topological and metric relationships now well-known thanks to thirty years of experience with GIS. Likewise, a robust system of ''topological'' temporal relationships for time intervals were defined by Allen [4] and subsequently refined [45]. Temporal metrics are a simplified case of spatial metrics (distance is the only type of temporal metric, since time is normally assumed to be one-dimensional). In order to support these applications, data models should provide the concept of spatial/temporal relationships as operators.

Data models in this view can be either simple, incorporating the basic spatial/temporal types: point, line and area, instant and time-interval. More sophisticated models may include:

- Support for an extensible, derived hierarchy of various spatial/temporal entity categories (e.g., city, county, state). These derived categories can incorporate complex types made up of a heterogeneous set of configurational types, like lines and areas.
- Support for structural constructs—object class, attribute and even relationship—to have a spatial extent/temporal lifespan, thus achieving orthogonality between the structure and the space/time dimensions.

5.2. *The continuous view*

In this view, each continuous space-time extent is characterized by the domain and range of its defining function. Often, the domain is the whole spatial and temporal extent described by the statio-temporal database. It can also be any contiguous subarea, and/or time interval selected from within the database. The range is the entire set of values for some measured attribute (e.g., temperature, population density) associated with the given space-time domain. There can be any number of such attributes and corresponding attribute ranges. A set of associated functions are defined for a given attribute and may be specific to the nature of the attribute. Spatial and temporal relations are valid as operators, as in the discrete view.

Similar to the discrete view, data models can be simple, incorporating a set of attributes that can conceptually be treated as flat ''layers.'' More sophisticated models may include:

- Support for a multi-scale spatial and/or temporal hierarchy. These derived hierarchies can represent regular subdivisions (e.g., area coatdress, actress, etc.).
- Support for an irregular subdivision of space and/or time that are dictated by a particular application context (e.g., population density or elevation, implemented via TIN's) or distribution of specific clusters of attribute values (e.g., R-trees).
- Support for structural constructs where entities have identity as sets of space-time locations.

5.3. *Integrated implementations*

Most spatial applications need both discrete and continuous views of space and time: some information is ''naturally'' perceived as spatial entities while other as continuous fields. For instance, weather forecast programs use discrete areas for counties, discrete points for cities and continuous spatial fields for temperature and rainfall. And indeed, most current commercially available GIS software provides this capability. In the same way, temporal applications need both discrete and continuous views: some temporal entities are ''naturally'' perceived as discrete and separate ''events'' (e.g., storms, volcanic eruptions, elections, etc.). Others are perceived as time functions, such as temperature and precipitation.

Moreover, discrete and continuous views for statio-temporal databases need to be functionally interrelated, supporting a common language and a common interface without requiring to transform one view into the other. Many current GIS do not support such a common interface, yet such an integration would allow combinations of spatial and temporal relationships, as operators defined within a unified query language, to form complex queries that can operate on discrete and continuous views simultaneously. For example; find all houses within 100 meters of the river built after the last flood and on land less than 500 meters elevation.

6. Toward the future

None of the above discussion is meant to imply that a unified representational theory of space-time, finely-attuned to the digital database context will solve all representational issues, and that we can all go off to build fully functional space-time databases and software systems as a straightforward engineering exercise. There is still much to do, at all levels; theoretical, conceptual and implementational, as I will expand upon below.

Before expanding upon what we still need to tackle, I go back to my earlier questions: Should we have expected things to have evolved any differently?—probably not. This renders my second question of whether we would now be further along moot. Looking at this same question relative to the future, however, we should simply keep in mind that we need to be conscious of working on all of these levels and at how they interrelate.

Back to our hero and his encounter with Gollum in the cave. Like Bilbo Baggins, researchers in space-time representation started out feeling their way in the dark. We then came across the unexpected in our explorations. It is the unexpected, found early-on through these initial explorations, that provides both: (1) The realization that the problem is harder than we thought, as Bilbo discovered when challenged by the unexpected riddles, and (2) a guide to the way out.

We have already learned much from our explorations in space-time data representation. Here are four broad, but key, theoretical areas, from my perspective, that we still need to tackle. Some are recently recognized, some have already been worked on for a significant period of time:

- *Development of a theory of space-time representation*: Although progress is being made within GIS and DBMS communities and closer cooperation and coordination between them, our approach needs to become more broadly interdisciplinary. Existing theories in a range of fields, from philosophy and physics to cognitive psychology need to be examined within the context of digital database representation and use. The need here is to allow the building of representations that go beyond traditional approaches. Semantically-driven representations and query languages are needed that seem ''natural'' to the human user, and that at the same time utilize the medium of computing to best advantage. This includes continuing development of an algebra of space-time relations.

  Specific questions: What are the interrelationships between continuous and discrete space and time? When are linear and when are cyclical representations oftime appropriate? How do we handle the persistence of facts/objects in a changing history? How do we deal with the potential of multiple histories or multiple potential futures? These mostly are problems relating to temporal logics, as well as representation of stored information.
- *Development of strategies for handling inexact and complex objects and relationships*: Geographic entities tend to have very complex interrelationships, and the definition of individual entities and relationships tends to be imprecise and context-dependent. For example, fundamental spatial and temporal relationships include near/far and before/

after, respectively. With current representational approaches, we have been imposing an artificial level of certainty and precision on the data.

Specific questions: How do we represent fuzzy and inexact concepts in a computing context? How can we develop space-time query languages that can appropriately handle inexactness?

- *Development of a means to deal with missing data*. Temporal databases will often represent historical data, where it is impossible to fill in the gaps of missing information from some previous time. This is particularly important when data must be collected retrospectively via secondary evidence. Applications abound where this is standard procedure (e.g., reconstructing cultural histories and evolution of land use through archaeological digs, reconstructing climate histories through tree ring and geological evidence).

  Specific questions: How do we accommodate incomplete temporal information? What kinds of interpolation procedures would be appropriate for various situations to fill-in missing data?
- *Development of a means to deal with multiple times and alternative histories*. It is not always as simple as valid and database time. There can be many others. An archaeological application is a good example. In the database for an archaeological dig, there can be stored; the date an artifact was found, the date the artifact was entered into the database, the date attributed to the object as the time of its fabrication, and the date attributed to the stratum in which the artifact was found. In this and many other contexts dealing with space-time information, simulation is an important analytical tool. The results of multiple simulations, as alternate histories or future-histories, need to be recorded.

  Specific questions: How do multiple times, as separate yet related pieces of information interrelate? How should they be represented? Compared?

References

1. T. Abraham and J.F. Roddick. ''Survey of Spatio-temporal databases,'' *Geoinformatica*, Vol. 3:61–69, 1999.
2. H. Afsarmanesh, D. McLeod, and A. Parker. ''An extensible Object-oriented approach to databases for VLSI/CAD,'' in *Proceedings of the Eleventh International Conference on Very Large Data Bases*, Stockholm, 1985.
3. O. Ahlqvist, J. Keukelaar, and K. Oukbir. ''Rough classification and accuracy assessment,'' *International Journal of Geographical Information Science*, Vol. 14:475–496, 2000.
4. J.F. Allen. ''Maintaining knowledge about temporal intervals,'' *Communications of the ACM*, Vol. 26:832–843, 1983.
5. K.K. Al-Taha. Temporal Reasoning in Cadastral Systems. Department of Surveying Engineering, University of Maine, 1992.
6. M.-A. Aufaure-Portier and C. Bonhomme. ''A high level visual language for spatial data management,'' in *Proceedings of Visual '99, 3rd International Conference on Visual Information Systems*, Amsterdam, 1999.
7. F. Barbic and B. Percini. ''Time Modeling in Office Information Systems,'' in *Proceedings of SIGMOD '85*, Austin, TX, 1985.
8. C. Bonhomme, C. Trépied, M.A. Aufaure, and R. Laurini ''A visual language for querying spatio-temporal

databases,'' in *Proceedings of the 7th ACM Symposium on Advances in Geographic Information Systems*, Kansas City, MO, 1999.
9. P.A. Burrough. ''Fuzzy mathematical methods for soil survey and land evaluation,'' *Journal of Soil Science*, Vol. 40:477–492, 1989.
10. P.A. Burrough and A.U. Frank. ''Concepts and paradigms in spatial information: Are current geographical information systems truly generic?,'' *International Journal of Geographical Information Systems*, Vol. 9:101–116, 1995.
11. P.A. Burrough, R.A. MacMillan, and W. van Deursen. ''Fuzzy classification methods for determining land suitability from soil profile observations and topography,'' *Journal of Soil Science*, Vol. 43:193–210, 1992.
12. T.S. Cheng and S.K. Gadia. ''A pattern matching language for spatio-temporal databases,'' in *Proceedings of the Third International Conference on Information and Knowledge Management*, Gaithersburg, MD, 1994.
13. J. Chomicki and P.Z. Revesz. ''Constraint-based interoperability of spatiotemporal databases,'' *Geoinformatica*, Vol. 3:211–243, 1999.
14. N.R. Chrisman. ''Concepts of space as a guide to cartographic data structures,'' in *Proceedings of Harvard Advanced Study Symposium on Topological Data Structures*, Cambridge, MA, 1977.
15. C. Claramunt and M. Thériault. ''Managing Time in GIS: An Event-Oriented Approach,'' in J. Clifford and A. Tuzhilin (Eds.), Recent Advances in Temporal Databases. Springer-Verlag: 1995.
16. C. Claramunt and M. Thériault. ''Toward semantics for modelling spatio-temporal processes within GIS,'' in M. Kraak and M. Molenaar (Eds.), Advances in GIS Research II: Proceedings of of the Seventh International Symposium on Spatial Data Handling. Taylor & Francis: 1996.
17. U. Dayal and J.M. Smith. ''A knowledge-oriented database management system,'' in M.L. Brodie and J. Mylopoulos (Eds.), On Knowledge Base Management Systems. Springer-Verlag: 1986.
18. S. Dragicevic and D.J. Marceau. ''A fuzzy set approach for modelling time in GIS,'' *International Journal of Geographical Information Science*, Vol. 14:225–245, 2000.
19. C.E. Dyreson and R.T. Snodgrass. ''Valid-time indeterminacy,'' in *Proceedings of the 9th International Conference on Data Engineering*, Vienna, 1993.
20. M.J. Egenhofer, J. Glasgow, D. Günther, J.R. Herring, and D.J. Peuquet. ''Progress in computational methods for representing geographical concepts,'' *International Journal of Geographical Information Science*, Vol. 13:775–796, 1999.
21. M. Erwig, R.H. Güting, M. Schneider, and M. Vazirgiannis. ''Abstract and discrete modeling of spatio-temporal data types,'' in *Proceedings of ACM GIS '98*, Washington, D.C., 1998.
22. M. Erwig, R.H. Güting, M. Schneider, and M. Vazirgiannis. ''Spatio-temporal data types: An approach to modeling and querying moving objects in databases,'' *GeoInformatica*, Vol. 3:269–296, 1999.
23. P.F. Fisher. ''Knowledge-based approaches to determining and correcting areas of unreliability in geographic databases,'' in M. Goodchild and S. Gopal (Eds.), Accuracy of Spatial Databases. Taylor & Francis: 1989.
24. P.F. Fisher. ''Models of uncertainty in spatial data,'' in P.A. Longley, M.F. Goodchild, D.J. Maguire, and D.W. Rhind (Eds.), Geographical Information Systems. Wiley: 1999.
25. L. Florizzi, R.H. Güting, E. Nardelli, and M. Schneider. ''A data model and data structures for moving objects databases,'' in *Proceedings of SIGMOD Conference 2000*, Dallas, TX, 2000.
26. A.S. Fotheringham and D.W.S. Wong. ''The modifiable areal unit problem in multivariate statistical analysis,'' *Environment and Planning A*, Vol. 23:1025–1044, 1991.
27. C. Freksa. ''Temporal reasoning based on semi-intervals,'' *Artificial Intelligence*, Vol. 54:199–227, 1992.
28. P.L. Galison. ''Minkowski's space-time: From visual thinking to the absolute world,'' *Historical Studies in the Physical Sciences*, Vol. 10:85–121, 1985.
29. K. Gardels. ''Sequoia 2000 and geographic information: The Guernewood geoprocessor,'' in *Proceedings of the Sixth International Symposium on Spatial Data Handling*, Edinburgh, Scotland, 1994.
30. M.F. Goodchild. ''Geographical information science,'' *International Journal of Geographical Information Systems*, Vol. 6:31–45, 1992.
31. M.F. Goodchild and S. Gopal. Accuracy of Spatial Databases. Taylor & Francis: 1989.

32. M.F. Goodchild, S. Guoqing, and Y. Shiren. ''Development and test of an error model for categorical data,'' *International Journal of Geographical Information Systems*, Vol. 6:87–104, 1992.
33. P. Gould. The Slow Plague: A Geography of the AIDS Pandemic. Blackwell: 1993.
34. P. Gould. ''Why not? The search for spatiotemporal structure,'' *Environment and Planning A*, Vol. 25:48–55, 1993.
35. M. Green and R. Flowerdew. ''New evidence on the modifiable areal unit problem,'' in P. Longley and M. Batty (Eds.), Spatial Analysis Modelling in a GIS Environment. Wiley: 1996.
36. D. Gregory. ''Solid geometry: Notes on the recovery of spatial structure,'' in P. Gould and G. Olsson (Eds.), A Search for Common Ground. Pion: 1982.
37. S. Grumbach, P. Rigeux, and L. Segoufin. ''Spatio-temporal data handling with constraints,'' in *Proceedings of ACM GIS '98*, Washington, D.C., 1998.
38. R.H. Güting, M.H. Böhlen, M. Erwig, C.S. Hansen, N.A. Lorentzos, M. Schneider, and M. Vazirgiannis. ''A foundation for representing and querying moving objects,'' *ACM Transactions on Database Systems*, Vol. 25: 2000.
39. T. Hägerstrand. Innovation Diffusion as a Spatial Process. 1967 ed. The University of Chicago Press: 1967.
40. J.B. Harley and D. Woodward. The History of Cartography. Vol. 1. University of Chicago Press: 1987.
41. N.W.J. Hazelton. Integrating Time, Dynamic Modelling and Geographical Information Systems: Development of Four-Dimensional GIS. Dept. of Surveying and Land Information, The University of Melbourne, 1991.
42. J.R. Herring. ''TIGRIS: A data model for an object-oriented geographic information system,'' *Computers & Geosciences*, Vol. 18:443–452, 1992.
43. K. Hornsby and M.J. Egenhofer. ''Identity-based change: A foundation for spatio-temporal knowledge representation,'' *International Journal of Geographical Information Science*, Vol. 14:207–224, 2000.
44. G.J. Hunter. ''Managing uncertainty in GIS,'' in P.A. Longley, M.F. Goodchild, D.J. Maguire, and D.W. Rhind (Eds.), Geographical Information Systems. Wiley, 1999.
45. C.S. Jensen, et al. ''The TSQL benchmark,'' in *Proceedings of the International Workshop on an Infrastructure for Temporal Databases*, Arlington, TX, 1993.
46. P.C. Kanellakis, G.M. Kuper, and P.Z. Revesz. ''Constraint query languages,'' *Journal of Computer and System Sciences*, Vol. 51:26–52, 1995.
47. J. Kelmelis. Time and Space in Geographic Information: Toward a Four-Dimensional Spatio-Temporal Data Model. The Pennsylvania State University, 1991.
48. G.M. Kuper, L. Libkin, and J. Paredaens. (Eds.), Constraint Databases. Springer Verlag: 2000.
49. G. Lakoff. ''The invariance hypothesis: Is abstract reason based on image-schemas?,'' *Cognitive Linguistics*, Vol. 1:39–74, 1990.
50. G. Langran. Time in Geographic Information Systems. Taylor & Francis: 1992.
51. G. Langran and N.R. Chrisman. ''A framework for temporal geographic information,'' *Cartographica*, Vol. 25:1–14, 1988.
52. P. Lolonis and M. Armstrong. ''Temporal information in spatial decision support systems,'' in *Proceedings of GIS/LIS*, Minneapolis, 1993.
53. D. Mark and F. Csillag. ''The nature of boundaries on 'area-class' maps,'' *Cartographica*, Vol. 26:65–78, 1989.
54. M. Molenaar. An Introduction to the Theory of Spatial Object Modelling for GIS. Taylor & Francis: 1998.
55. A. Montanari and B. Pernici. ''Towards a temporal logic reconstruction of temporal databases,'' in *Proceedings of the International Workshop on an Infrastructure for Temporal Databases*, Arlington, TX, 1993.
56. C. Parent, S. Spaccapietra, and E. Zimányi. ''Spatio-temporal conceptual models: Data structures + space + time,'' in *Proceedings of ACM-GIS'99: the 7th International Symposium on Advances in Geographic Information Systems*, Kansas City, MO, 1999.
57. D. Parkes and N. Thrift. Times, Spaces, and Places. John Wiley & Sons: 1980.
58. D.J. Peuquet. ''It's about time: A conceptual framework for the representation of temporal dynamics in geographic information systems,'' *Annals of the Association of American Geographers*, Vol. 84:441–461, 1994.

59. D.J. Peuquet and N. Duan. ''An event-based spatiotemporal data model (ESTDM) for temporal analysis of geographical data,'' *International Journal of Geographical Information Systems*, Vol. 9:7–24, 1995.
60. D. Pfoser and C.S. Jensen. ''Capturing the uncertainty of moving-object representations,'' in *Proceedings of the 6th International Symposium on the Advances in Spatial Databases—SSD '99*, Hong Kong, 1999.
61. D. Pfoser and N. Tryfona. Fuzziness and Uncertainty in Spatiotemporal Applications, Chorochronos project report, 2000.
62. A. Pred. ''The choreography of existence: Comments on Hagerstrand's time geography and its usefulness,'' *Economic Geography*, Vol. 53:207–221, 1977.
63. Qualcomm. OmniTRACS: System Overview, www.qualcomm.com/qwbs/products/omnitracs-overview.html.
64. H.M. Rafaat, Z. Yang, and D. Gauthier. ''Relational spatial topologies for historical geographical information,'' *International Journal of Geographical Information Science*, Vol. 8:163–173, 1994.
65. J. Raper and D. Livingstone. ''Development of a geomorphological spatial model using object-oriented design,'' *International Journal of Geographical Information Systems*, Vol. 9:359–383, 1995.
66. A. Raza and W. Kainz. ''Cell tuple based spatio-temporal data model: An object oriented approach,'' in *Proceedings of ACM-GIS'99: Proceedings of the 7th International Symposium on Advances in Geographic Information Systems*, Kansas City, MO, 1999.
67. J.F. Roddick and J.D. Patrick. ''Temporal semantics in information systems—A survey,'' *Information Systems*, Vol. 17:249–267, 1992.
68. U. Schiel. ''An abstract introduction to the temporal-hierarchic data model,'' in *Proceedings of the Ninth International Conference on Very Large Data Bases*, New York, 1983.
69. T. Sellis. ''Research issues in spatio-temporal database systems,'' in *Proceedings of Advances in Spatial Databases: 6th International Symposium, SSD '99*, Hong Kong, China, 1999.
70. R. Shibasaki. ''Handling spatio-temporal uncertainties of geo-objects for dynamic update of GIS databases from multi-source data,'' in *Proceedings of Advanced Geographic Data Modelling*, Delft, The Netherlands, 1994.
71. D. Sinton. ''The inherent structure of information as a constraint to analysis: Mapped thematic data as a case study,'' in *Proceedings of Harvard Advanced Study Symposium on Topological Data Structures*, Cambridge, MA, 1977.
72. R.T. Snodgrass (Ed.) The TSQL2 Temporal Query Language. Kluwer International Series in Engineering and Computer Science; No. 330. Kluwer: 1995.
73. State of Maryland. ''MAGI: Maryland automated geographic information system,'' in D.J. Peuquet and D.F. Marble (Eds.) Introductory Readings in Geographic Information Systems. Taylor & Francis: 1990.
74. M. Stonebraker. ''Triggers and inference in database systems,'' in M.L. Brodie and J. Mylopoulos (Eds.) On Knowledge Base Management Systems. Springer-Verlag: 1986.
75. R. Studer. ''Modeling time aspects of information systems,'' *Proceedings of the International Conference on Data Engineering*, 1986.
76. A.U. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Seger, and R. Snodgrass (Eds.) Temporal Databases: Theory, Design and Application. Database Systems and Applications. Benjamin/Cummings Publishing Co.: 1993.
77. W. Tobler. ''Automation and Cartography,'' *Geographical Review*, Vol. 49:526–534, 1959.
78. J.R.R. Tolkein. The Hobbit, or There and Back Again. Houghton Mifflin: 1966.
79. R.F. Tomlinson, H.W. Calkins, and D.F. Marble. Computer Handling of Geographical Data: An Examination of Selected Geographical Information Systems, International Geographical Union, Natural Resource Series No. 8, 1976.
80. N. Tryfona and C.S. Jensen. ''Conceptual data modeling for spatiotemporal applications,'' *GeoInformatica*, Vol. 3:245–268, 1999.
81. N. Tryfona and C.S. Jensen. ''Using abstractions for spatio-temporal conceptual modeling,'' in *Proceedings of the 2000 ACM Symposium on Applied Computing*, Como, Italy, 2000.
82. N. Tryfona, D. Pfoser, and T. Hadzilacos. ''Modeling behavior of geographic objects: An experience with the object modeling technique,'' in *Proceedings of the Ninth Conference on Advanced Information Systems and Engineering (CASE '97), Lecture Notes in Computer Science 1250*, Springer-Verlag, 1997.

83. S.D. Urban and L. Delcambre. ''An analysis of the structural, dynamic and temporal aspects of semantic data models,'' in *Proceedings of International Conference on Data Engineering*, New York, 1986.
84. M. Wachowicz and R. Healey. ''Towards temporality in GIS,'' in M.F. Worboys (Eds.), Innovations in GIS. Taylor & Francis: 1994.
85. F. Wang and G.B. Hall. ''Fuzzy representation of geographical boundaries in GIS,'' *International Journal of Geographical Information Systems*, Vol. 10:573–590, 1996.
86. F. Wang, G.B. Hall, and Subaryono. ''Fuzzy information representation and processing in conventional GIS software: Database design and application,'' *International Journal of Geographical Information Systems*, Vol. 4:261–283, 1990.
87. W. Weber. ''Three types of map data structures, their ands and nots and a possible or,'' in *Proceedings of Harvard Advanced Study Symposium on Topological Data Structures*, Cambridge, MA, 1977.
88. M.F. Worboys. ''Object oriented approaches to geo-referenced information,'' *International Journal of Geographic Information Systems*, Vol. 8:385–389, 1994.
89. M.F. Worboys. ''A unified model for spatial and temporal information,'' *The Computer Journal*, Vol. 37:26–33, 1994.
90. M. Yuan. ''Wildfire conceptual modeling for building GIS space-time models,'' in *Proceedings of GIS/LIS '94*, Phoenix, 1994.
91. L.A. Zadeh. ''Fuzzy sets,'' *Information and Control*, Vol. 8:338–353, 1965.

Donna Peuquet is currently Professor in the Department of Geography, The Pennsylvania State University. She holds degrees from the University of Cincinnati and the State University of New York at Buffalo. Her principal research interests are spatio-temporal data models, spatial cognition, and data mining.