Barrera, Renato, Andrew U. Frank, and Khaled Al-Taha. "Temporal Relations in Geographic Information Systems: A Workshop at the University of Maine." Orono: NCGIA, 1991.

Temporal Relations in Geographic Information Systems: A Workshop at the University of Maine, Orono, October 12-13, 1990

April 1991

Renato Barrera, Andrew Frank, and Khaled Al-Taha

NCGIA
Department of Surveying Engineering
University of Maine
Orono, ME 04469

National Center for Geographic Information and Analysis

Report 91-4

INTRODUCTION

As the ability to acquire data and process information constantly improves, the time interval between the occurrence of some event and its description as mapped information is decreasing rapidly. Where once a map was an amalgam of data covering decades, even centuries (Blakemore and Harley 1980 p. 56), today a map may be updated every few hours, as are weather maps. This ability to update maps on a frequent basis brings with it a greater need to evaluate the temporal accuracy of geographic information. For now it is still up to the user to search out temporal information, which is seldom presented as part of the map or explicitly incorporated into a database.

This paper proposes a model of time that can be useful in an analysis of temporal accuracy of geographic information. The nature of time is discussed, with brief attention to its abstraction in some models, then the notion of multiple temporal dimensions is presented. When a basis has been laid out for a multidimensional model, temporal data is plotted into a two dimensional coordinate system, and the resultant structure analyzed. Finally, the two dimensional model is expanded into a multidimensional model with and indeterminate number of dimensions.

THE ABSTRACTION OF TIME

Nothing exists in isolation, neither time nor space nor all the stuff of the universe. Time is not a separable phenomenon, but exists irrevocably as an element of the tripartite relationship of mass, space and time. More recently, in the words of Albert Einstein, "the special theory of relativity ... discards absolute time and space and makes them in every instance relative to moving systems. ... now we know that time and space are not the vessel for the universe, but could not exist at all if there were no contents, namely, no sun, earth and other celestial bodies." (Lorentz 1920 p. 12).

However, in the normal course of our activities, we heartily embrace the apparently contradictory view. We think of time as a thing with a life of its own. We all, of course, have an intuitive understanding of what time is. Most of us see this "psychological time" as something flowing inexorably from the past, through the present and into the unknown future (Hawking 1988 p. 145). Scholarly and scientific disciplines use it as a primary base, history and paleontology among them. This intuitive understanding is itself a model, a purposeful abstraction. It is not an assertion of reality opposed to Dr. Einstein's, above.

Abstraction is useful; it makes certain understandings possible, that, in its absence, would be impossible. Certain details (accurate and full correspondence with reality) are temporarily suspended, in order to clarify some part of reality to the extent that we can better understand it. Specifically, we ignore the details of time as relationship, and formulate concepts of time as a thing. We are wise insofar as we maintain this notion of time as an useful abstraction, with some of its details suspended temporarily, and we are foolish insofar as we simplify it by eliminating those details altogether.

MODELS OF TIME

The simplest temporal model is a time line. Langran and Chrisman have used this as a module to build their descriptive model with multiple time lines (Figure 1). It models the database as a collection of separate time lines for each object, which are parallel with the time lines of each other object and their collective map, and which share a common temporal metric (Langran and Chrisman 1988 p. 5, Langran 1989a p. 35).

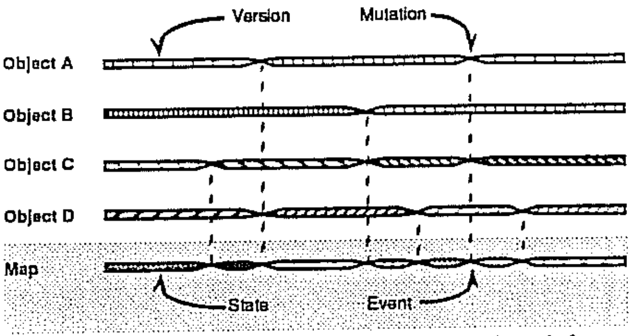

Figure 1: The topology of cartographic time (after Langran and Chrisman 1988)

This is not a two dimensional model, but it is the decomposition of a singly dimensioned model. With each object on a separate time line, temporal topology is implied by the adjacency of mutations and events on separate time lines. The decomposition and thus the implicit topology is necessary, because if all objects were to occupy the same time line, it would be impossible to resolve the various mutations and events.

Some models represent time as spatial metaphor (Hawking 1988, Rucker 1977, 1984, Szegö 1987, Thrift 1977). They reduce the complexity of spatial representation so that a temporal dimension may be shown in the familiar context of a third spatial dimension. Various paths through this space-time volume trace planar coordinate positions at various times (heights). These models have limitations for modeling the temporal topology of a database, because their single temporal dimension does not allow the portrayal of the many temporal relationships inherent in the construction and use of a database.

More importantly, while temporal and spatial dimensions are similar, in that they are metrics, their measures are not similar. Temporal measures are of changes in state, over time for a single entity, while spatial measures are of changes in state between two entities, over space and at a single time. This fundamental distinction between duration and distance is the most serious difficulty to be overcome in developing a model for a combined spatiotemporal topology. This discussion is concerned with the development of a model that fully describes the temporal topology of a GIS. Its combination with spatial topology is left for further research.

DIMENSIONS FOR TIME

Through multiple dimensions, this model seeks to portray explicitly and clearly two topological aspects of geographic data: 1) the temporal connectivity of data objects with each other, with their real world referents, and with the database or map; and 2) the full set of institutions which a GIS user may recognize as related to the mapped entities and to the GIS.

A dimension is a measurable property of some phenomenon. It is the property itself and not a measurement of the property nor the phenomenon, itself. While it is true that dimension often refers, in common parlance, to the spatial properties of length, breadth and depth, dimensions are not inherently spatial constructs. For example, some aspatial dimensional intersections are meals per month, dollars per year per person, and aircraft accidents per airline. Since dimensions are not inherently spatial, they may be used to relate various measures of time as well.

The user of a geographic information system (GIS) needs to deal simultaneously with many different contexts:

1) a set of mappable entities of the real world;
2) the data objects that represent the entities within a GIS;
3) all the institutions that, in the perspective of the user, have some direct interest in the mappable entities or the GIS.

Since a GIS represents the state and spatial arrangement of some collection of mappable phenomena, the events which alter their state or arrangement are of appreciable importance. This model is built upon the recognition that these events are an integral part of the mapping process, and that they may occur in any of the several contexts. For present purposes, the intuitive and obvious understanding of an event as a change of state (Langran and Chrisman 1988) is sufficient, but this is discussed more fully later, with mensuration.

Dimension as context. Langran and Chrisman expressed the view that an object may have more than one dimension, when they wrote that, "cartographic time is punctuated by 'events.' or changes, which are recorded along two axes" (1988 p. 4). This model postulates multiple temporal dimensions of a cartographic object and portrays them as being mutually perpendicular. The dimensions are the contexts within which a data object has meaning to a GIS user.

One especially appropriate instance of aspatial orthogonal dimensions is found in the analysis of training sets for the interpretation of remotely sensed data. Each of several spectral bands is considered a separate dimension for the measured reflectance of each pixel, and the data are statistically aggregated into multidimensional "clouds" in order to interpret land cover types (Lillesand and Kiefer p. 671-689).

Langran and Chrisman recognize "two concurrent clocks" and call them "world time" and "database time" (1988 p. 2, Langran 1989a p. 36, 1989b p. 218). In Langran's words, "world time traces real-world events, database time traces the history of database transactions" (1989b). This paper adopts this aptly descriptive terminology, but proposes alternate definitions for the two, since in a strict sense, time does not trace either events or their history. It is the measured data which are the traces, not the dimension in which those measurements are expressed. Thus, world time is the chronometric dimension related to some mappable entity, while database time is the chronometric dimension related to some data object. For a discussion of entities and data objects, see Morrison (1988 p. 24). The general events which are recognized in all dimensions are a beginning, a change and an end.

World Time (T_w). The events measured in world time are the beginning, discovery, change and extinction of some mappable entity. In practice, the primary event may be discovery rather than beginning, for the entity must be known before it may be measured, and in many cases this datum will be sufficient for the user's purpose. Subsequent events are perceptible and substantial changes in the entity. The entity ceases to exist at extinction. Some real world events are: the construction of a subdivision; the alteration of a river's course; the demolition of a road; a measurement of hospital admissions; an eruption of a volcano.

Database Time (T_d). In database time, the measured transactions of a data object are insertion, mutation and deletion (Langran 1989a, 1989b). It is seldom the case that a world event and the database transaction related to this event occur simultaneously. To account for the usual temporal disparity between the event and the transaction, two additional conditions are specified. A *prospective* database transaction occurs before its related event in world time, while a *retrospective* database transaction occurs afterward.

This discussion takes the view that insertion and deletion are the fundamental database transactions. A mutation is considered to be not a separate sort of a transaction, but a special case where the deletion and insertion of two related data objects have identical temporal measurements in each dimension (T_w and T_i).

Institutional Time (T_i). Other contexts or temporal dimensions related to each entity and data object are collectively called institutional time, though it is understood that there is a separate dimension for each institution related to the mappable entity. Institutional time is the chronometric dimension with respect to some institution that has a known, direct relation to the mappable entity. An institution is simply one or more persons. There may be zero or more recognized, relevant institutions. Institutional events are inception, commencement, modification, partial completion (way points), abandonment and completion.

Some institutional events are: the conception or preliminary planning of a road; building permit granted for some building; commencement of construction; intermediate inspection reports for some construction project; steel reinforcing installed; occupancy permit granted; expiration of building permit; occupation of the building; public opening of a road.

While not strictly related to a GIS, a bank deposit may provide a clear example of the interplay of world, database and institutional times. A customer might make a deposit of funds to a bank's after-hours depository at, say, 8:00 PM. The bank would consider the deposit to have been made at 10:00 AM, the opening of business, on the following day, while the deposit may be entered in to bank's database by an empolyee at yet another time, perhaps 9:15 AM. Thus the same event has been measured at three different times in each of the three contexts (dimenisons).

MENSURATION

An event is a simplification of a range of time. Since all things change continually, often imperceptibly, the notion of an event must be restricted to significant change: a change which is both perceptible (measurable) and substantial. Certainly, significance is dependent upon context. For example, a large river might alter its course by a few feet; in a rural context this might be insubstantial and imperceptible, while in an urban context it might be not only substantial and measurable, but catastrophic. Time may be the phenomenon that changes substantially, as with a periodic update of a database.

Temporal Resolution. Resolution is minimum detectable change, and it is related to both change of time and change of state in the relevant context. Because temporal measurements may be controlled, as a periodic update, or event driven, as recording ownership change upon sale of a parcel of land, the measurements may be made of entities which have not undergone significant change of state. Either the change in time or the change of state must be perceptible and substantial, in order for an event to be recognized and a new measurement to be made. In practice, both may be required to exhibit substantial change. At some scheduled periodic update, not all entities may be measured, nor may all objects be mutated or deleted; substantially unchanged features may simply carry over.

Some examples of differing resolutions in different contexts are: a database might be updated every ten years, annually, weekly, daily, hourly, etc.; a river's course might make a substantial change in five years or 100; weather conditions might change substantially by the hour but be measured every six hours; a building department might grant building permits with a minute-by-minute time stamp, but enter permits into their database on a daily basis, while the life of the permit might be two years.

While resolution may vary from one temporal dimension to another, this model requires that each temporal dimension be measured in common units. This discussion assumes these units to be of calendar time.

Temporal Error. Temporal error may be classified as mensural or contextual error. Mensural error is the difference between any particular time measurement of an event and the actual time that it occured. Contextual error is the result of the different views of the various contexts, where the times of different but related events are measured. When the events are not coincident in time, their temporal measures must differ. The inability of previous conceptual models to explain this fundamental difference between contextual and mensural temporal error has traditionally been dealt with by the conventional fiction of considering time to be constant while theme and location are then controlled or measured (Sinton 1978). It is exactly this temporal error, which is derived from contextual differences, that this model now seeks to incorporate and quantify.

It is important to note that temporal measurement is but one part of temporal resolution, which is also dependent upon change of state in the related context. Where an error of temporal measurement is less than the total discrimination of the resolution, the error is not significant. Conversely, where temporal errors of measurement are greater than the lowest level of resolution, they are significant errors.

Returning to the bank deposit example, the temporal differences all contain contextual error. There is no measurement error in the time of the opening of business, as this is given, not measured. However, the times of the customer's deposit and the insertion of the relevant object into the database may have been measured more or less accurately, thus, these may also contain mensural error. This model is concerned simply with all significant differences between time measurements on different axes (dimensions), but it does not further discriminate these temporal differences as mensural or contextual errors.

When measurements related to the same data object are identical in all dimensions, no error exists. That is to say, there is no temporal error when an entity, its data object and their related analogues in all other contexts exist at the same time. When there is some differing temporal measurement, the database is in error. Temporal error states have two determinants:

1) the type of database transaction (insertion or

deletion);

2) whether the transaction occurs before or after the associated world event (prospective or retrospective).

Table 1 summarizes all possible error states. The cells of the table show both the type of error involved as well as the orientation of the representative line segment in the 2-D version of the model. This orientation is discussed more fully, below, as the model is developed.

	Insertion	Deletion
Prospective	Commission ⊥ to T_w	Omission ⊥ to T_d
Retrospective	Omission ⊥ to T_d	Commission ⊥ to T_w

(⊥ = perpendicular)

Table 1: Temporal Error of Transactions

An object for which no entity exists is an error of commission. An error of omission occurs when some relevant entity is not represented in the database.

THE MODEL

From the GIS user's perspective, each data object is the focus of a set of related contexts; it is more than the representation of a single, real-world entity. For example, a linear feature might represent the existence of a section of roadway, but from the user's perspective, it may also bear a significant relation to a planning department, highway department, building department, contractor, census bureau, tax department, contiguous property owners, etc. The multiply dimensioned structure of the model is an expression of this relation of contexts, in which each axis represents some recognized, relevant context.

Thus, this temporal model is the orthogonal intersection of all relevant dimensions, with common units of measurement for all and a common origin. In practice, the point of intersection of the axes may be any convenient time, so long as it is common to all dimensions.

One Dimensional Representation. Figure 2 is a representation of the real-world temporal dimension for a section of road.

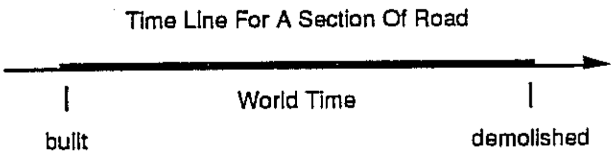

Figure 2: Time Modeled as a Single Dimension

This shows that a road section was completed at one date, that it existed for some period of time, and that it was demolished at some later date. A similar time line might be drawn for the data object which represents this road section. Its transactions would be the times when the data object was entered into the database and finally deleted from it. By combining the two, the model is similar to the familiar form of Figure 1, above.

Two Dimensional Representation. If these two time lines are set so they intersect at a right angle at some common time (Langran and Chrisman 1988 p. 2), then they form the orthogonal axes of a cartesian coordinate system. With this arrangement, one is able to represent explicitly the temporal relationship of related events measured on the two axes. A point in this system represents the times when the related events occur in each dimension. It is important to note, that a point may be plotted in the coordinate space only when a measured value exists on each axis. Returning to the above example, the creation of the road section is not mappable when it occurs, if its related data object has not yet been entered into the database. A value exists on one axis (T_w), but not on the other (T_d), so no point exists.

An important feature of this model is the line of synchronicity: the simple linear function $T_d = T_w$. It represents the condition of being "up to date". On the most fundamental level, once the database and world times are synchronized, one can be treated as the other, the map is seen as an accurate reflection of the world at large. This synchronicity is, at once, both the ideal and, in the absence of contrary evidence, the assumed state.

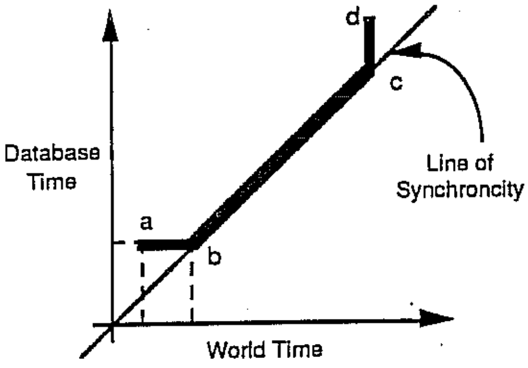

Figure 3: A Model of Two Dimensional Time

Figure 3 is a representation of two temporal dimensions for a road section. It shows a simple history: (a) the section is built; (b) it is entered into the database; (c) the road is demolished; and (d) it is deleted from the database.

Point (a) represents the different times of construction and insertion. Point (b) represents the first instance of synchronicity, when the entity and its data object both exist. All the points from (a) to (b) are plotted simultaneously. The segment (a,b) represents, for a retrospective database transaction, the time period that the database was in error, because an extant entity was not represented. The segment (b,c) represents the period of temporal accuracy. The segment (c,d) represents the time when the defunct entity was still erroneously represented in the database. It cannot be plotted until the date of extinction is known, but it then continues with invariant world time, until it terminates upon the object's deletion. Both of these asynchronous segments are typical of retrospective database transactions. For brevity and clarity, asynchronous segments are called bristles.

The separation of a bristle from the line of synchronicity is a graphic illustration of the lead or lag between the occurrence of a real world event and its representation in the database. But while the bristle is a clear indication of both the type and duration of temporal error, it can show that error only when it has become apparent in both dimensions, i.e. when it is measured on each axis. In the interim, synchronicity and, thus, the accuracy of the map is assumed.

Assumptions about accuracy differ according to the user's perspective. From the perspective of world time, synchronicity is assumed between measurements on T_W, while from the perspective of database time, it is assumed between measurements on T_d. In short, the world observer assumes that what is seen is in the database, while the GIS user assumes that the entities that are represented internally actually exist. The latter assumption is undoubtedly the more frequent of the two.

In his discussion of Land Information Systems (LIS) as communication devices, Bédard enumerates the uncertainties caused by "fuzziness associated with the definition or identification of a spatial entity or ... in the description and location (in space and *time*) of this spatial entity" (1986 p. 148). [emphasis added] This uncertainty about time is the result of temporal error, whether or not it is adequately quantified. Bédard recognizes that uncertainty (error) is affected by the methods of identifying land parcels as well as the precision of measurements, both internally and externally, with respect to the LIS (1986 p. 151 - 156).

By appropriate choice of methods of identification and measurement, the uncertainty is reduced, but it cannot be eliminated. What remains must be "absorbed" by the makers or users of a GIS (Bédard 1986 p. 157). This absorption of error and the assumption of accuracy are one and the same. This model provides an explicit representation of this aspect of Bédard's communication model of LIS.

In the example, when the feature was demolished at (c), the database user, being unaware of the fact, assumes that the time plot for the object continues along the line of synchronicity. When the fact is known and the deletion is made, the synchronous plot must be truncated at (c) and the segment (c,d) inserted. It is the multidimensional nature of this model that allows this accurate portrayal of the behavior and assumptions of the database user. This portrayal is not possible in models with only one temporal dimension.

In the 2-D form of this model, line segments have only three possible orientations:

1) bristles perpendicular to the database time axis represent errors of omission (extant but unrepresented entities);
2) bristles perpendicular to the world time axis represent errors of commission (nonexistant but represented entities);
3) segments upon the line of synchronicity represent temporal accuracy.

Either error state may occur at either insertion or deletion. For instance a subdivision might be represented in the database before it was actually built, or some object might be mistakenly removed from the database while its relevant entity still existed. These last two are both prospective database events, where the database transaction preceded the world event.

At this point, the first tracks of the polygon of time are just visible. In Figure 3, the area below the bristle (a,b) and above the line of synchronicity is a right isosceles triangle. In this model, this is always the case, because all bristles are perpendicular to some axis, therefore the distance (a,b) equals that from (a) to the line of synchronicity measured perpendicular to T_W. The conditions are similar for all bristles.

The important characteristics of the model are:

1) The line of synchronicity is the collection of all points where $T_W = T_d$.
2) Measurements of temporally accurate database objects are always plotted onto the line of synchronicity.
3) Measurements of temporally inaccurate database objects cannot be plotted until a measurement is made in each dimension. When they are plotted, the resultant bristles are always perpendicular to one of the axes, and their length is proportional to the length of time of the error.
4) The orientation of a bristle and the identification of the related database event as an insertion or deletion present an unmistakable signature of the type of error which the bristle represents. In general, insertion bristles are perpendicular to the axis with the latest measurement, while deletion bristles are perpendicular to the axis with the earliest measurement.

These general characteristics are independent of the number of dimensions being modeled.

Multiple Dimensions. The two-dimensional model may be expanded into three dimensions, with the addition of an institutional dimension. In the example, this third dimension might add the perspective of a State Highway Department. Figure 4 is a representation of a 3-D temporal structure with three data objects. Again the line of synchronicity is all of the points which have the same value in each dimension.

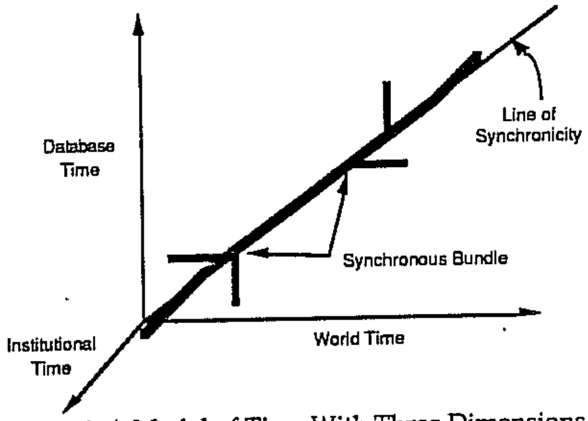

Figure 4: A Model of Time With Three Dimensions

As in the 2-D model, the bristles' orientations are significant, and they represent the time period during which a particular data object is in temporal error. The bundled (coincidently synchronous) data objects represent the period of temporal accuracy with respect to both the world referents (entities) and some other relevant context. This simple example is not meant to suggest that a database might ever achieve complete synchronicity among all extant data objects and all contexts.

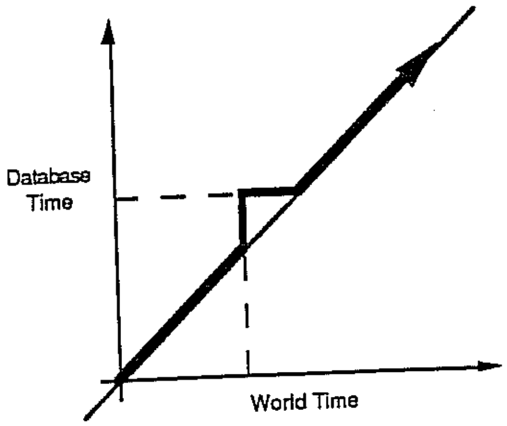

Figure 5: The Mutation of a Data Object

Finally, this model embodies the idea that bristles may also represent transitional events or mutations between two "versions" (Langran and Chrisman 1988, p. 5) of an object. For instance, a section of graded dirt road might be paved, and being a significant change in terms of some GIS, this change is later recorded in the database. This temporal error is represented by a pair of subsequent synchronous segments connected sequentially by a deletion and an insertion bristle. Figure 5 represents this situation. Note that both bristles are retrospective bristles and that the bristle of deletion is perpendicular to the world time axis at the time when the event (repaving) occurred, while the bristle of insertion is perpendicular to the database time axis at the time when the mutation of the data object occurred.

In this model of multidimensional time, the mutation of a data object is more complicated than in Langran and Chrisman's model of Figure 1. In the current model, the mutations are related to the states and events of entities as well as the states and events of the database (the collective map). Thus, while both the real world and database events have instantaneous measurements, the model clearly shows both the time difference between the two events and the order of the events. Parenthetically, the mutation is the clearest example of the polygon of time, which is, as ever, a right isosceles triangle.

While there may well be a practical limit, there is no theoretic limit to the number of dimensions of similar temporal models. Each additional dimension would represent an additional, relevant context. The inability to graphically represent them may make them less useful for some explanatory purposes, but their use for analysis would be similar to that already shown.

CONCLUSIONS

Synchronicity in the multidimensional model is an explicit expression of the temporal accuracy of the various contextual relations of data objects with each other and, by extension, the map. This range of function is greater than that provided by the "parallel lines" or the 3-D spatiotemporal models.

The multidimensional model also points the way to an expanded notion of topology in a GIS. When designing geographic information systems, it may be useful to implement the topology of data objects not only over space and time but also among a number of related contexts. For instance a municipal database might find use for topological information relating data objects to various departments (Building and Land, Assessor, etc.) or to other interested parties (property owners with active building permits, state highways, Federal property, etc.). It is beyond the scope of this discussion to speculate as to how such additional topology might actually be implemented, but it is useful to consider the full range of potential needs as data models are designed. It might be cheaper in the long run to design a comprehensive model with a greater range of function than currently anticipated, rather that to design a simpler, more restricted model which will have to be redesigned all the sooner.

ACKNOWLEDGEMENTS

The genesis of this model can be found in Langran and Chrisman's description of the "essentially orthogonal" nature of world and database time (1988 p. 4). This author's construction has been greatly affected by the insights and criticisms of Dr. Chrisman, to whom I am abundantly grateful.

US National Science Foundation Grant SES 87-22084 provided partial support for this paper.

REASONING ABOUT GIS USING TEMPORAL AND DYNAMIC LOGICS

M. F. Worboys
Midlands Regional Research Laboratory and Department of Computing Studies
University of Leicester, Leicester LE1 7RH
United Kingdom

October 10th, 1990

ABSTRACT. This paper is a tutorial introduction to temporal logic in the context of spatio-temporal information systems. It begins with a brief description of modal logic. The main part of the paper discusses temporal and dynamic logics. It concludes by briefly indicating possible applications of these logics to GIS.

1. Introduction

The purpose of this paper is to introduce readers to temporal logic and to apply the methods developed to some simple examples in temporal geographic information systems. In order to develop the material, the paper begins with more general modal logics. Temporal logics are constructed as special cases of modal logics. Particular temporal logics which have special properties are then discussed. A distinction is drawn between temporal logics, used for modelling and reasoning with the information in temporal GIS, and dynamic logics, used for modelling the action of the information system itself as it evolves over time. In order to make the exposition more concrete, simple examples of the application of these logics to temporal GIS are presented. The paper concludes with some general observations about the use of temporal logics in the representation of spatio-temporal information. The technical presentation is necessarily brief. For a fuller treatment, the reader could consult [Hughes and Creswell, 1972, 1984] for general modal logics, and [Rescher and Urquehart, 1971] for temporal logics.

2. Brief description of modal logic

Logic has a most important role to play in the representation of knowledge. However, not all types of knowledge can be represented using traditional propositional and predicate logics. Such logics are useful for formalizing the reasoning within static domains, as would occur for example within an unchanging database. But for more dynamic situations, such as occur in spatio-temporal systems, classical logic is not sufficient. One approach to the representation of spatio-temporal knowledge is to use modal logic. Specific modal logics concerned with temporal situations are tense or temporal logics. Modal logic applied to general dynamic domains is referred to as dynamic or action logic. We begin by describing the most general of these logics, namely the modal system K.

When studying modal logic we shall confine ourselves to modal propositional logic. The vocabulary of modal propositional logic is formed by adding the modal operators $\Box$ (for "necessarily") and $\Diamond$ (for "possibly") to the vocabulary of propositional logic. In all the systems which we shall study $\Diamond\varphi$ is taken to be an abbreviation for $\neg \Box \neg \varphi$. The intuition we are relying on here is that φ is a necessary truth if, and only if, φ is true in all possible scenarios (or at least all those which may be envisaged from the present scenario). φ is a possible truth if, and only if, $\neg \varphi$ is not a necessary truth.

The formation rules for modal propositional logic are as for propositional logic, with the addition of:

If φ is a well-formed formula then $\Box\varphi$ is also well-formed.

2.1 PROOF THEORY

The proof system we consider first is called the system K (K for Saul Kripke, the inventor of possible world semantics for modal logic). The tradition in modal logic is to treat the proof theory as an extension of a Hilbert style classical logic. That is to say, the axioms are all theorems, and the rules of derivation allow one to derive theorems from theorems.

When a wff φ is provable as a theorem in K, we shall represent this fact by $K \vdash \varphi$, or sometimes, when there is no risk of ambiguity, by $\vdash \varphi$.

Axioms:

$K \vdash \varphi$, where φ is a classical tautology.

$K \vdash \Box(\varphi \rightarrow \psi) \rightarrow (\Box\varphi \rightarrow \Box\psi)$ (distribution axiom)

Rules of Derivation

$$\frac{K \vdash \varphi \rightarrow \psi \qquad K \rightarrow \phi}{K \vdash \psi}$$ (modus ponens)

$$\frac{K \vdash \varphi}{K \vdash \Box\varphi}$$ (necessitation rule)

2.2 SEMANTICS: KRIPKE FRAMES AND MODELS

These semantics are usually attributed to Saul Kripke. In fact, the idea seemed to occur to Kripke, Hintikka, and Kanger independently. The seminal paper that is usually quoted is [Kripke 1963].

A *Kripke frame* $\langle W, \rho \rangle$ is a set W together with a relation ρ on W. The relation is often called an *accessibility relation* on W. The idea is that W is a set of possible scenarios, or worlds, and ρ is a relationship of accessibility between these worlds. That is $w_1 \ \rho \ w_2$ if, and only if, from the scenario of w_1, w_2 is considered a possibility.

A *Kripke model* $\mathcal{M} = \langle W, \rho, \mathcal{V} \rangle$ of a set of wffs $\mathcal{F}$, is a Kripke frame together with a *valuation function* $\mathcal{V}$, such that:

$$\mathcal{V}: \mathcal{F} \rightarrow \mathcal{P}(W)$$

and for all $\varphi, \psi \in \mathcal{F}$:

$$\mathcal{V}(\varphi \wedge \psi) = \mathcal{V}(\varphi) \cap \mathcal{V}(\psi)$$
$$\mathcal{V}(\neg\varphi) = W - \mathcal{V}(\varphi)$$
$$\mathcal{V}(\Box\varphi) = \{w: \forall x \in W \text{ (If } w\rho x\text{, then } x \in \mathcal{V}(\varphi))\}$$

The intuition here is that $\mathcal{V}(\varphi)$ consists of all worlds in which φ is true. The last statement may be interpreted as stating that $\Box\varphi$ is to be true in world w if, and only if, φ is true in every world accessible to world w. We may represent a Kripke model by a labelled directed graph as in the diagram below. The worlds are represented by nodes, the accessibility relationship by edges, and each node is labelled by propositions which are true in the world which that node represents.

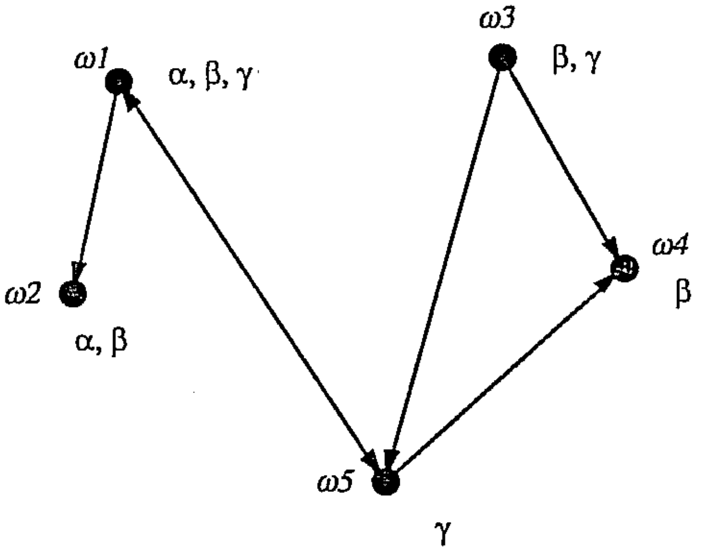

The diagram shows that world w_2 is accessible from world w_1; world w_1 is not accessible from world w_2; worlds w_1 and w_5 are co-accessible; propositions α, β and γ are true in world w_1 but only α and β are true in world w_2; and so on.

Let $\mathcal{M}$ be a Kripke model $\langle W, \rho, \mathcal{V} \rangle$ of a set of wffs $\mathcal{F}$. Let $w \in W$ and $\varphi \in \mathcal{F}$. Then:

$\langle \mathcal{M}, w \rangle \models \varphi$ if, and only if, $w \in \mathcal{V}(\varphi)$.

φ is *valid in the Kripke model* $\mathcal{M}$ if, and only if, $\langle \mathcal{M}, w \rangle \models \varphi$, for all $w \in W$. (In this case we write $\mathcal{M} \models \varphi$, or often just $\models \varphi$ if no ambiguity arises.)

φ is *valid in a Kripke frame* $<W,\rho>$ if, and only if, φ is valid in $\mathfrak{M}$, for all Kripke models $\mathfrak{M}$ which are based upon $<W,\rho>$. (We write $<W,\rho> \models \varphi$, or often just $\models \varphi$ if no ambiguity arises.)

Soundness and completeness theorem

φ is provable in K if, and only if, φ is valid in all Kripke frames

3. Temporal systems

Temporal or tense logic allows the possibility of formulating tensed propositions. Temporal logic is a modal logic formed by the introduction of some operators which express tenses. The basic vocabulary is as follows:

$G\varphi$: φ will always be the case.
$H\varphi$: φ has always been the case.
$P\varphi$: φ has been the case (at some time in the past).
$F\varphi$: φ will be the case (at some time in the future).

G and H correspond, roughly, to two different versions of $\Box$, and P and F to two different versions of $\Diamond$. P is $\neg H \neg$, F is $\neg G \neg$.

3.1 THE MINIMAL SYSTEM: TEMPORAL K

The system below formalizes relationships between the temporal operators G, H, P, and F. This system, which is a development of Kripke's work, described in Section 1, is attributed to E.J. Lemmon and D. Scott (notes not fully published).

Axioms:

All propositional tautologies.
$G(\varphi \rightarrow \psi) \rightarrow (G\varphi \rightarrow G\psi)$
$H(\varphi \rightarrow \psi) \rightarrow (H\varphi \rightarrow H\psi)$
$\varphi \rightarrow GP\varphi$
$\varphi \rightarrow HF\varphi$

Rules of Derivation:

$$\frac{\vdash \varphi}{\vdash G\varphi} \qquad \frac{\vdash \varphi}{\vdash H\varphi}$$

and modus ponens.

3.1.1 Semantics for Temporal K

Frames and models are defined as before, except that we have to extend our account of the valuation function to take account of G, H, P, and F.

$$\mathcal{V}(G\varphi) = \{w \mid \forall w'(w\rho w' \to w' \in \mathcal{V}(\varphi)\}$$
$$\mathcal{V}(F\varphi) = \{w \mid \exists w'(w\rho w' \wedge w' \in \mathcal{V}(\varphi))\}$$
$$\mathcal{V}(H\varphi) = \{w \mid \forall w'(w'\rho w \to w' \in \mathcal{V}(\varphi))\}$$
$$\mathcal{V}(P\varphi) = \{w \mid \exists w'(w'\rho w \wedge w' \in \mathcal{V}(\varphi))\}$$

Our intuition here is that the possible worlds in a temporal frame are instants of time, and that $w\rho w'$ if, and only if, w is earlier than w'.

3.1.2 Soundness and Completeness Theorem

φ is provable in Temporal K if, and only if, φ is valid in all Kripke temporal frames.

3.2 LINEAR TIME

In 3.1, there is no restriction on the structure of the temporal frames. It is possible to quite easily impose restrictions which may more closely mirror our view of the structure of time for particular applications. For example, we may view time as a linear stream. Absolute time in Newtonian physics has this structure. The instants of time may be imagined to be arranged along a line, as shown below.

This restriction on the temporal frame structure is mirrored by adding a further axiom to Temporal K:

$$PF\varphi \to (P\varphi \vee \varphi \vee F\varphi)$$

There are appropriate soundness and completeness results.

3.3 BRANCHING TIME

Branching time is a weaker system, in that every linear time frame is branching. The intuition that we intend to capture here is that at any present moment there is the possibility of a number of different futures, but only one past. The diagram below shows the kinds of frames possible.

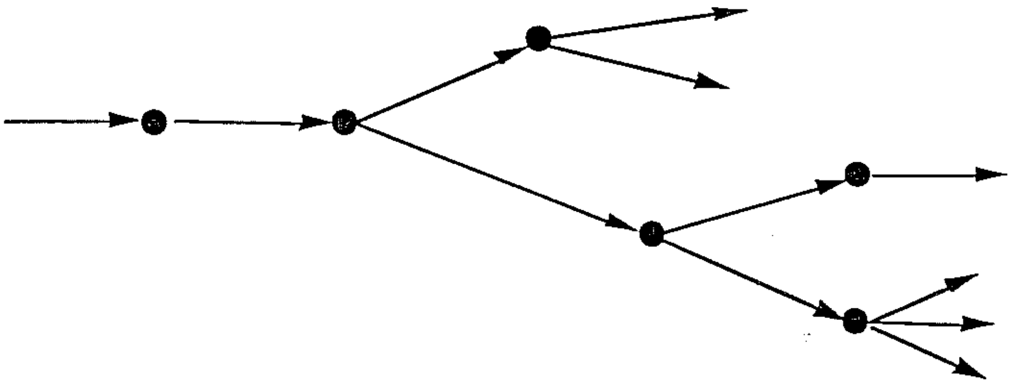

The axioms to be added to Temporal K are the following:

$G\varphi \rightarrow GG\varphi$

$H\varphi \rightarrow HH\varphi$

$(H(\varphi \vee \psi) \wedge H(\varphi \vee H\psi) \wedge H(H\varphi \vee \psi)) \rightarrow (H\varphi \vee H\psi)$

Again, there are appropriate soundness and completeness results.

3.4 OTHER LINEAR TEMPORAL SYSTEMS

It is possible to refine our systems to capture more and more the structural features of the particular temporal dimension in which we wish to work. Starting from linear time, it is possible to express the conditions that time is to have a beginning or an ending. The axiom scheme that may be added to express the proposition that time to have an ending is:

$$G\varphi \vee FG\varphi$$

The intuition here is that when time has an ending, then at any given moment, for any proposition φ, either we are at the end of time, in which case $G\varphi$ is true, or we shall be at the end of time at some point in the future, in which case $FG\varphi$ is the case. A symmetrical axiom can be given to specify that time has a beginning.

An interesting condition that can be imposed upon the temporal dimension, and one very much bound up with measurement processes, is density. We may wish to stipulate that the temporal dimension is continuous, that is can be subdivided arbitrarily finely.

Consider the pair of axiom schemes:

$$GG\varphi \rightarrow G\varphi$$
$$HH\varphi \rightarrow H\varphi$$

or equivalently:

$$F\varphi \rightarrow FF\varphi$$
$$P\varphi \rightarrow PP\varphi$$

In order to gain some intuitive understanding of these axioms, suppose that the time line below is dense.

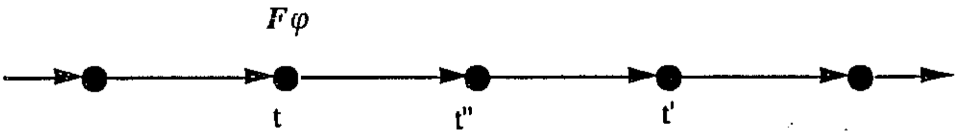

Consider an arbitrary time t on the line. Assume at time t that $F\varphi$ is true. Then there exists a future time, say t', when φ is true. Assuming the densieness of the time line, there is a time t'' strictly between t and t'. At t'', it is the case that φ is true at a future time, namely t', and so $F\varphi$ holds at t''. At time t, it is the case that $F\varphi$ is true at the future time t'', and so $FF\varphi$ holds at time t. To sum up; at time t, $F\varphi$ implies $FF\varphi$. Therefore $F\varphi \rightarrow FF\varphi$ holds at time t.

4 Dynamic Logic

The temporal logics discussed in Section 3 may be used to model time in an historical database. The other main type of temporality existing in an information system is the temporal dimension in which the system itself resides. We require a new formalism to model the changes that the information itself undergoes. The formal tool which has been developed to model 'system time' is *dynamic*, or *modal action* logic,.

We begin with the semantics of dynamic logic, and take the set W of possible scenarios to be the set of possible states of a knowledge base, and allow one world to be accessible from another if there is a transition from the former to the latter by means of a sequence of transactions (e.g. updates) on the knowledge base. There will be a set of accessibility relations, each relation corresponding to a transaction sequence (or *action*). More formally:

$w \; \rho_a \; w'$ if, and only if, we can get from state w to state w' via action a.

Suppose that a is some action. Then we can understand $[a]\varphi$ to mean that after action a has been performed upon the knowledge base, then φ must be the case.

[a] can be treated as a modal operator. A little reflection will quickly show that the necessitation rule and the distributive axiom are entirely reasonable when [a] is understood in this way. We also have an extra set of axioms concerning negation which can be asserted if each state of the knowledge base is a consistent state. As before, <a> is an abbreviation for $\neg\,[a]\,\neg$.

4.1 PROOF SYSTEM

Axioms:

All propositional tautologies
$[a](\varphi \rightarrow \psi) \rightarrow ([a]\varphi \rightarrow [a]\psi)$ for all actions a.
$[a]\neg\varphi \rightarrow \neg[a]\varphi$ for all actions a.

Rules of Derivation:

$$\frac{\vdash \varphi}{\vdash [a]\varphi}$$

and modus ponens.

4.2 SEMANTICS

An *extended Kripke frame* $\langle W, \{\rho_a : a \text{ an action}\}\rangle$ is a set W together with a set of relations ρ_a on W (one for each action a). Models are defined as before, except that the result of the valuation function $\mathcal{V}$ acting on $[a]\varphi$ is given by:

$$\mathcal{V}([a]\varphi) = \{w \mid \forall w'(w\,\rho_a\,w' \rightarrow w' \in \mathcal{V}(\varphi))\}.$$

There are soundness and completeness results, as for the earlier system. The topic of dynamic logic is developed fully in [Harel, 1979], and its application to knowledge base specification is discussed in [Khosla et al, 1986].

5. Current work on temporal issues in databases and GIS

For a database management system to be labelled temporal, it should be capable of supporting a model of a time-varying world. There have been three main areas of research here: the development of temporal query languages, the development of an underlying logical model and theory (cf relational theory for relational databases), and of course the question of physical implementation.

Characteristic problem areas for temporal databases are storage, object identity, and representation of temporal knowledge. Regarding storage, it is clearly inefficient to store the contents of the entire database of information recorded at different times. Much of the data may be time-invariant, and this method would lead to redundant duplication of information. However, it is not clear at this time what is the best approach to storage for temporal databases. [Langran, 1989] lists various approaches taken in the literature.

Object identity is required if the system is to recognise differing versions of the same time-varying object. The principle here is that an object has an identity independent of its attributes. Thus if the attributes change, the object is still intact. Object identity is not supported by the pure relational model, although it can be argued that it is supported by Codd's extended model RM/T [Codd, 1979]. However, RM/T is not implemented in any currently marketed system known to the author. Object-oriented database systems have addressed this problem of object identity.

5.1. TEMPORAL KNOWLEDGE REPRESENTATION

Temporal knowledge may be represented using the modal formalism outlined in section 3. There is a distinction between *registration* time (also called transactional, recording history, or system time), when a transaction takes place and the *logical* time (also called valid, historical, extrinsic, or object level time) at which the events recorded actually happened. At one extreme, a database may be *dynamic* in that the full transaction history is recorded, and yet contain information with no temporal reference. At the other, a database may be *static* with no continuous stream of transactions, (maybe with no update facility), but contained temporally referenced information.

For a dynamic database holding temporally referenced data, transactions may be *retroactive* (i.e. coming into effect after the time to which the data was referenced), or *proactive* (i.e. coming into effect before the time to which the data was referenced).

There may be more than one type of logical time to be recorded in the same database. For example, if written forms are used, the time of the document which contains the observation may be important. Thus, in the harbour information system described below, notification of the placement of a buoy in the harbour at coordinates (45,46) may have the following times associated with it::

office historic time	notifying document dated	28 August
registration time	data entered into GIS	1 September
harbour historic time	buoy placed in harbour	8 September.

5.2. DATA WITH A SPATIO-TEMPORAL REFERENCE

Future GIS should be capable of managing data referenced to both space and time. Langran (1989) proposes that a temporal GIS should be able to provide answers to such questions as:

Where was this feature located two years ago?
How has this feature changed over the last five years?
Has this feature moved?
Have these two features ever been juxtaposed?

5.2.1 Extended HYDRO

We use the example HYDRO given in [Langran, 1989]. HYDRO contains marine information. Each piece of information relates to a feature, gives an (x,y)-position, and gives the date of the source document which refers to it.

On 30th August, HYDRO contains information:

Feature	X	Y	Effective
Beacon	38	38	8 March
Shoal	23	32	2 May
Hazard	56	34	3 March
Bell	44	28	8 August.

On 1 September, the following transactions occurred:

Buoy	added proactively; coordinates (45,46); taking effect from 8 September. (Plans are made to place it in the harbour the following week.)
Shoal	coordinates altered to (15,21); new survey information; document dated 18 August.
Light	added; coordinates (12,18); document dated 18 August.
Hazard	deleted.
Bell	deleted.

On 5 September, the following transactions occurred:

Hazard	restored.
Light	coordinates corrected postactively to (18,18); document dated 18 August.

Example constraints:

Below is given a subset of the possible integrity constraints which the harbour system must satisfy. A distinction is made between static constraints, which relate to individual states of the system, and dynamic constraints, which operate on temporal sequences (in both historical and registration times) of states of the system.

Static:	Each feature has a unique coordinate set.
Dynamic (historical):	Lighthouses do not move.
Dynamic (registration):	The date of a document specifying an observation must be earlier than the date of insertion of the corresponding data into the database.

Example queries:

Next is given a pair of typical retrieval requests to the system. Both queries include time as an integral component.

Retrieve all features currently active within a spatio-temporal region.
Retrieve all features that may never vary temporally.

Formalization

We illustrate the use of a typed tense first order predicate logic in this context by formalizing two of the above constraints.

That each feature has a unique coordinate reference may be expressed as:

$\forall x, x', y, y' : \mathbb{R}$ (the class of real numbers)
$\forall f : \text{Feature}$
$(\text{coordinates}\ (f, x, y) \wedge \text{coordinates}\ (f, x', y')) \rightarrow (x = x' \wedge y = y')$

That lighthouses do not move may be expressed as:

$\forall x, y : \mathbb{R}$
$\forall l : \text{Lighthouse}$
$\text{coordinates}\ (l, x, y) \rightarrow G(\text{coordinates}\ (l, x, y))$

Notice that only the second constraint requires a temporal operator.

5.2.2 Boundary change example

The next example shows up some of the special problems of a database of spatio-temporally referenced data. The difficulties of transforming data from one spatial reference unit to another are well known. These difficulties become compounded if the spatial units are time-variable. A brief discussion of these matters is given in [Hearnshaw et al, 1989]. We illustrate with a local example. The city of Leicester, UK, is partitioned into a number of district electoral wards for purposes of the census and local administration. The partition is regularly modified as boundaries are changed and wards are inserted or deleted. Attribute data is collected, which is referenced to the partition at a particular time (or set of times).

Imagine that at time t_1, Leicester was divided into two wards, A and B. At time t_2, in order to adjust for population movements, the boundary between A and B was altered. At time t_3, ward B was divided into wards C and D. Throughout this time period, feature f is always in ward A. We illustrate this in the diagram below.

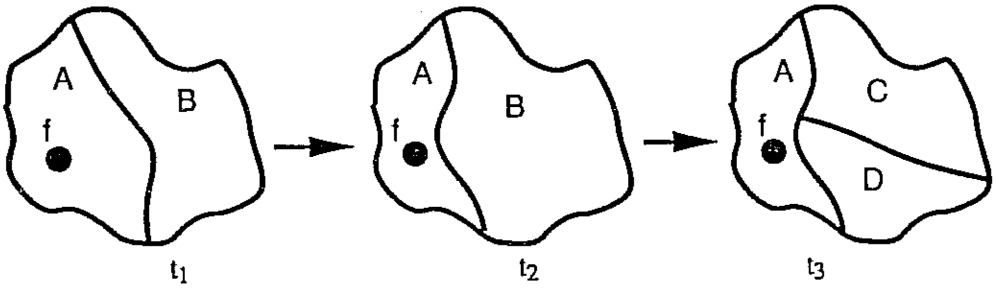

The knowledge representation problem is how to formalize this information. To illustrate the approach using tense logic, the knowledge that at time t_1 feature f is and will always be inside region A may be expressed as:

$G\ (\text{inside}\ (f, A))$

We can also imagine a query language which is a modal extension of the relational calculus which could be used to retrieve information stored in a database about such a system. In [Worboys et al, 1990], we develop an object retrieval language (ORC) which interacts with an object oriented database. ORC is an object version of the domain relational calculus, considered in [Lacroix and Pirotte, 1977]. It is a straightforward matter to furnish ORC with a modal component to handle spatio-temporal object systems.

6. Conclusions

Little is yet known about the introduction of the temporal dimension into spatial knowledge bases. Indeed, the study of general temporal databases is still undeveloped. But it is clear that spatio-temporal systems have many practical applications and will certainly be constructed. It is therefore important to establish their theoretical foundations.

Using classical logic, we may be able to reason quite well within a spatial domain which contains information about the state of affairs at particular point in time. We need extra machinery to be able to reason about a spatial domain which is changing over time. A further complication occurs since the knowledge base which holds this information may itself be dynamic, knowledge being inserted, updated and deleted over time. Temporal and dynamic logics are presented in this paper as formalisms for handling these complex spatio-temporal systems. We do not wish, however, to make the claim that the formalism of modal logic is sufficient to capture all of our temporal concepts, or all temporal constructions within natural language. However, such formalisms are a first step. This work is in its early stages. The details of the representation of spatio-temporal knowledge, and in particular the use of modal logic to represent topological/geometrical knowledge in a temporal setting, are still to be worked out.

Regarding implementation, the basic problem is how to store the information efficiently and effectively. At one extreme, we could store all the geometry/topology of the partition at each change time, even if only a single boundary is altered. This uses a great deal of storage, but the model is simple and retrieval will be fast. At the other extreme, store a start state of the partition and then only changes (*deltas*) to the partition. This *temporal spaghetti* will be low on storage space but the model is messy and retrievals will be inefficient. A compromise would be to store the whole partition, store the deltas, and after a specified number of deltas store the whole partition again, more deltas, and so on. The question of temporal database storage is still unanswered in the basic relational case, and so *a fortiori* for spatial data. A promising approach is to store the historical spatial information as a single spatio-temporal structure, and use dynamic logic to model the dynamic nature of the holding information system. A further paper is being written on this approach.

We certainly do not claim that modal logic holds all the answers to temporal knowledge representation. We hope to have acquainted the reader with some of the issues in this fascinating area, and indicated some of the directions for future research.

Acknowledgements

The author thanks Nick Measor, University of Leicester, for introducing him to modal logic and for interesting discussions.

References

CLIFFORD, J., WARREN, D.S., (1983), Formal semantics for time in databases, ACM TODS, 8, pp. 214-254.

CODD, E. F., (1979), Extending the database relational model to capture more meaning, ACM TODS, 4 (4), pp. 397-434.

HAREL, D., (1979), *First-Order Dynamic Logic*, Springer-Verlag, Berlin.

HEARNSHAW, H. M., MAGUIRE, D. J., WORBOYS, M. F., (1989), An Introduction to Area-Based Spatial Units: A Case Study of Leicestershire, Midlands Regional Research Laboratory, Research Report No. 1.

HUGHES, G.E., CRESSWELL, M.J., (1968), *An Introduction to Modal Logic*, Methuen, London (reprinted with corrections 1972).

HUGHES, G.E., CRESSWELL, M.J., (1984), *A Companion to Modal Logic*, Methuen, London

KHOSLA, S., MAIBAUM, T.S.E., SADLER, M., (1986), Database Specification, in *Database Semantics* (Ed. STEEL, T.B., MEERSMAN, R.), Elsevier, Netherlands.

KRIPKE, S, (1963), Semantical Analysis of Modal Logic I: Normal Propositional Calculi, Zeit. Math. Logik Grund. 9, pp. 67-96.

LACROIX, M., PIROTTE, A., (1977), Domain-Oriented Relational Languages, *Proc. 3rd Int.Conf.VLDB*.

LANGRAN, G., (1989), A review of temporal database research and its use in GIS applications, Int. J. GIS, 3 (3), pp. 215-232.

RESCHER, N., URQUEHART, A., (1971), *Temporal Logic* , Springer-Verlag.

WORBOYS, M. F., HEARNSHAW, H. M., and MAGUIRE, D. J., (1990), Object-oriented data and query modelling for geographical information systems, 4th International Symposium on Spatial Data Handling, Zurich.

Tracking the Temporal Polygon: A Conceptual Model of Multidimensional Time For Geographic Information Systems

Marcus Lester
Department of Geography, DP 10
University of Washington

Presented at:
The Temporal Workshop
University of Maine
October 13, 1990

Abstract. With improvements in technical ability to acquire data and to analyze and present information, temporal measurement is an increasingly significant component of geographic data quality. A model of multidimensional time may provide a useful conceptual structure for expressing temporal relationships of geographic information. It portrays temporal information as being composed of many dimensions: the contexts within which each datum has meaning to the user of a GIS. Patterns that develop as temporal data are plotted in multiple dimensions can both qualitatively and quantitatively describe the temporal accuracy of a GIS database.

Abstract

A workshop on temporal relations in Geographic Information Systems (GIS) was held on October 12-13, 1990, at the University of Maine. The meeting was sponsored by the National Center for Geographic Information and Analysis (NCGIA). Seventeen specialists gathered from the fields of Geography, GIS, and Computer Science to discuss users' requirements of temporal GIS and to identify the research issues posed by these requirements.

We found

- Two paradigms of time as users understand it: one that sees time as a continuum, in which evolution is gradual and often described by differential equations, and another in which time is a sequence of intervals and changes are caused by events.
- The need for including exploratory mechanisms appropriate for searches in large datasets.
- The convenience of including views that reflect a user's need e.g., temporal/spatial aggregation, and rules to construct these views and to enforce consistency.
- The need for the manipulation of inexact or approximate information, including provisions for approximate queries, for assessing the quality of data (i.e., its fitness for use), and for evaluating the propagation of errors in derived data.
- A need for a closer communication between GIS specialists and computer scientists. There is also a need for a division of labor, in which GIS specialists establish the requirements in formal terms, and the computer scientists provide optimal technical solutions. These needs for collaboration are in agreement with the findings of the NCGIA specialist meeting on very large spatial databases [SMIT90].

It became clear at the meeting that handling time related data in a GIS — a spatio/temporal information system — poses very challenging problems which may be quite difficult to solve.

Contents

1 Introduction

The National Center for Geographic Information and Analysis (NCGIA), created by the National Science foundation in August 1988, is organized around research initiatives that focus on specific topics relevant to GIS. The initiatives are planned well in advance to concentrate the efforts of researchers from several institutions, inside and outside the NCGIA, for a period of a twelve to eighteen months.

One of these initiatives, planned since the inception of the NCGIA, relates to the Temporal Relations existing in GIS. This initiative, known internally as "Initiative 10", has been designed to

1) Understand the modeling of time (continuous time, discrete time, and events) as it relates to GIS;

2) Assess inference methods in temporal logic and deduction strategies in non-monotonic systems;

3) Compare modeling of states to methods of modeling incremental changes to different GIS applications;

4) Study the architecture of a temporal GIS;

5) Extend the methods for dealing with multiple and alternative representations to include temporal aspects.

The start of Research Initiative 10 is programmed for the fall of 1991; meanwhile, significant independent work on the temporal aspects of GIS has been in progress in several universities, notably the University of Washington at Seattle, the University of Syracuse, Pennsylvania State University, and the University of Maine. It seemed therefore convenient to hold a workshop prior to the initiative's commencement.

First, the students presently working on time-related topics will benefit from a preliminary meeting and a frank exchange of opinions and experiences. The meeting will also help them in the planning of their work and in its integration within the broader context of the research initiative. It is hoped that they will contribute advanced research at the initiative's specialist meeting.

Second, many of the activities necessary for undertaking the research initiative (e.g., literature search, expert's identification) need to be done in advance. The GIS work-in-progress meeting was a good opportunity to bring together a few noted specialists in logic and databases and form a critical mass for the discussion of user's requirements and the preliminary identification of research issues in temporal GIS.

The workshop was planned for October 12^{th} and 13^{th} in Orono, Maine. Seventeen experts attended the workshop[1]; a majority of the participants were geographers, with engineers and computer scientists composing the rest of the group.

The workshop was organized so that sessions of presentations, either of tutorial or of expository character, were interweaved with two-hour discussions in which topics selected among the participants were discussed.

There were eight presentations:

- three short tutorials: one on temporal database query languages, one on temporal data models, and one on temporal logic.
- three presentations on issues in temporal representations; and
- two presentations on applications of temporal GIS: one for the study of global change, and one for cadastral systems.

In addition to the presentation, sessions were held for the exposition of several temporal GIS' applications, and for discussions on temporal representations, on temporal reasoning, and on the need for object identity.

Finally, during the last session the participants presented their individual conclusions on the meeting.

This document reports the results of the discussions. Since the issues considered in a session frequently overlapped or were complemented by those of other sessions, the results have been rearranged into four topics:

1. the requirements of a temporal database,
2. the representations of time,
3. software requirements of a temporal GIS,

---

[1] U. Dayal (DEC); M. Worboys (U. Leicester); G. Wuu (Bellcore); N. Chrisman, M. Lester (U. Washington); M. Monmonier, I. Vasiliev (U. Syracuse); J. Kelmelis, D. Peuquet (Pennsylvania State U.); M. Armstrong (U. Iowa); K. Al-Taha, R. Barrera, K. Beard, M. Egenhofer, A. Frank, S. Hornsby, and M. McGranahan (U. Maine).

4. existing solutions.

This rearrangement, compounded with the occurrence of intense interactions among the participants, made it difficult to cite the originators of each concept and comment; thus, such citations have been completely omitted.

To accelerate the process of publication, this report has not been circulated among the participants. We expressly acknowledge the important contributions by each one. The results of the workshop are the intellectual property of all of the participants, and this should be considered as a partial attempt to report them from the editors' point of view.

The remainder of this document is organized as follows: the next section presents two examples of temporal GIS posing rather diverse requirements; afterwards, each of the four previously cite topics is treated in a section, and the workshop's conclusions are gathered at the end. Finally, an appendix includes some relevant papers contributed by the participants.

2 Examples of temporal GIS

Two applications presented during the workshop will now be described. The first one, presented by K. Al-Taha, deals with the requirements of cadastral systems; the second one, based on an exposition by J. Kelmelis, depicts the role of a GIS in modeling Global Change.

Both applications need the support of a temporal GIS. Their requirements, however, are radically different, since they correspond to extremes of the gamut of applications. These examples will be used in section 3.1 to illustrate the qualitative and quantitative paradigms of GIS.

2.1 GIS and cadastral systems

A cadastral system records and documents legal information related to real-estate, such as ownership, mortgages, sales agreements, value and taxation, land use, rights of way, rights to use and to cultivate the land, etc. In any cadastral system, a flow of temporal data is caused by people's transitions or business transactions, since events such as sales agreements, adverse possession, loan's obtention, bankruptcy, owner's decease, etc. will affect a property's ownership.

Depending on how the ownership of real-estate is transferred, cadastral systems are divided into two classes:

Government intervention systems. In this type of cadastre, a government act, usually the registration of the new owner, is required for the ownership's transfer of real-estate. In such systems, land is defined as a union of irregular spatial objects (parcels). These systems deal with the owner relationship between people and parcels, as well as additional relationships referring to other rights in the land; these relationships change over time.

Interesting examples of government supported systems are the Torrens and the Swiss cadastral systems [USDA74] [WYLE76].

Deed registry systems. In this type of cadastre, deed registration is voluntary and unrecorded deeds may be valid. In these systems a property may not correspond to a well-defined spatial object; e.g., a person can sell one half a plot without specifying the exact location of the sold portion. In deed registry systems it is difficult to decide on the current owner of a parcel or on the location of a parcel's boundaries; only a thorough search of all the evidence collected in the registry can lead to an "opinion". This type of deed registration is typical of English-speaking countries [USDA74] [HINT90] [ONSR89].

Both types of cadastral systems require

- a mechanism to record changes in legal status caused by sales agreement, inheritance, confiscation, hostile or adverse possession, etc.
- a method for including spatial and topological updates such as modifications to the size or shape of a parcel, mergers or subdivisions of parcels, etc.
- a model of time based on events, to support the updates mentioned above
- at least two perspectives of time, to model both the lifespan of an event and the registration time of the corresponding transaction.
- a capability for temporal reasoning, to answer questions such as "who owns this parcel?", or "which one of two mortgages has precedence?".

Hunter and Williamson proposed such a system, named the "Historical digital cadastral database"[HUNT90] . This system can process temporal graphical data by adding time-coding attributes, and it can also display and report the parcellation of a region at any time in the past. It cannot support, however, topological changes or temporal reasoning.

An example of temporal aspects in cadastral systems. To show some of the events affecting a cadastral system and their consequences, let us consider the following story:

John, an old man, lived alone after his only daughter Ruth married and moved out of state. John's house accidentally burned; he was fortunate to escape the fire but had to be moved to a hospital and spent his last days in a nursery home. The insurance money, after paying the mortgage on the property (parcel number 9 in Figure 1), was partially used to support John's stay at the nursery home. Jeff, John's neighbor, needed a car garage next to his house (parcel 8) and wanted to get part of John's abundant property for that. As he could not contact Ruth, he decided to build the garage and maybe buy the property later. After a few years, the Community built a pedestrian path as a shortcut to the school behind John's property; this path split John's land into two parts. Jeff felt that the part attached to his property was essential to him so he decided to enlarge the garage, add a one-bedroom apartment on top for his son Joseph, and enlarge his backyard as well. After Jeff died, his wife sold the property to the Jackson family and moved to a down-town apartment.

Ten years later, Barbara, John's granddaughter, wanted to build a new house in place of her grandfather's house. To finance the new building, she needed a title opinion required by the bank. Only then did she find that her grandfather John had allowed a ten-foot right-of-way to his neighbor Jeff, and had given him a twenty-foot strip of his back yard to enlarge the school yard (parcel 37). One day, when she went to school to take her son to kindergarten, she accidently learned from the principal that part of John's property was taken by the community to build a walking path to the school, and that the rest was attached to the neighbor's property. When she referred to the surveyor to verify the boundary, it was also found that John's property consisted of two parcels (parcels 9 and 18), his own and an annexed parcel to it that was owned by his wife. All these changes are illustrated in Figure 1.

As presented in the example above, we can subdivide changes into different categories, such as

Legal changes:

Following are examples of legal actions that can lead to an ownership transfer of a property:

- Sales agreement– in 1908, John purchased parcel 9.
- Inheritance– in 1962, John died and his property transferred to Ruth, his daughter.
- Adverse possession– the Jackson family obtained an adverse possession on part of the land after publicly using it more than a certain number of years.

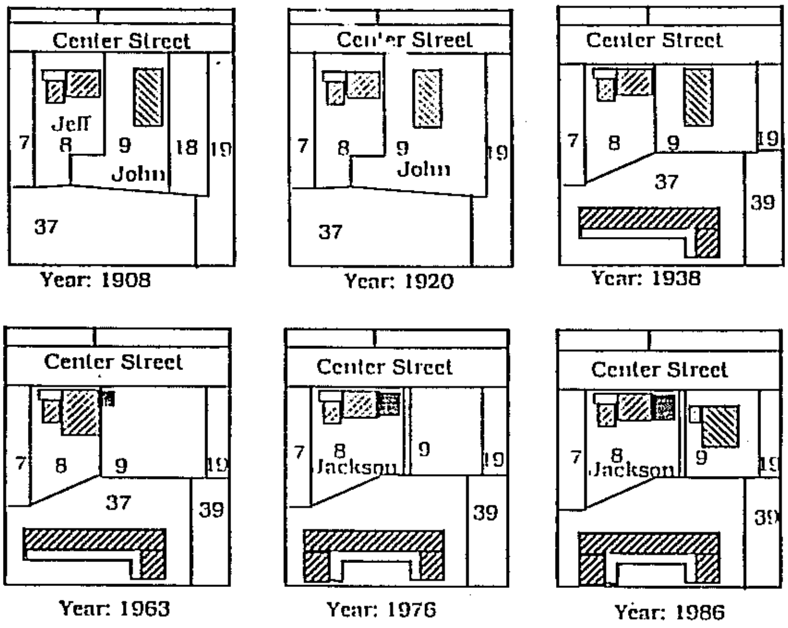

Year: 1908 Year: 1920 Year: 1938

Year: 1963 Year: 1976 Year: 1986

Figure 1: Spatial representation for parcel 9 and its neighborhood

Geometric changes: We can also observe geometric changes to the properties:

- Size – e.g., the areas of parcels 8 and 9 were changed by the enlargement of Center Street.
- Shape– e.g., parcel 9 became a quadrilateral in 1938, although it was originally L-shaped.
- Join– after the unification of two (or more) parcels, some of them may disappear and new ones may be created. Thus, parcel 18 ceased to exist after joined with parcel 9.
- Subdivision– a parcel can be partitioned, yielding one or more new parcels. e.g., parcel 19 was split into 19 and 39.

Topological changes: Some Topological changes common in cadastral systems are the following:

- Neighborhood– e.g., before 1964, parcels 8 and 9 were neighbors; in that year, however, they were separated by the pedestrian path. Parcel 19 became a neighbor to parcel 9 in the year 1920, and stopped sharing common boundary with the school in 1938.
- Boundary elements (points, lines)– e.g., in 1964, parcel 8 gained two more boundary points.

Land-use changes: The way a land is used (such as agriculture, forest, roads, public parks, housing, industry, and the like), or is planned to be used in the future, can change as well. Changes in the use of land imply direct changes to its value, and, consequently, to its tax.

Typical queries to be supported in cadastral systems are:

Who owned parcel 8 when the house burned?

Did property x have a mortgage in 1987?

Who were John's neighbors in 1933?

What percentage of owners in the flood area are insured?

What is the title abstract (history) for property x?

Does John have a good title for property y?

2.2 Global Change and GIS

The discovery of the ozone hole in Antarctica, the results of Atmospheric General Circulation Model runs which attempt to predict the future changes due to increased carbon dioxide, and a growing concern for the effects of rapid population increases have highlighted the need to consider our planet as a system. This fast succession of events illustrates the intensity of our new awareness: the ozone hole was discovered over Antarctica between August and October 1984. After the findings were published in early 1985, the National Science Foundation launched an expedition to Antarctica to study the phenomenon before the austral spring the same year. The public became concerned about this and other global change issues and, during 1988 and 1989, more than a hundred bills either directly concerned with global change or concerned with some global change issue were introduced to Congress. Very few scientific events have raised the public's ecological consciousness in such short time.

Earth, to be considered as a system, needs to be modeled and analyzed as a system. To help in this endeavor, the Working Group in Global Change was formed as part of the Committee on Earth and Environmental Sciences [2]; this group has identified a set priorities for research in global change.

A sample of the priority research issues is [CES89]:

Climate and Hydrologic Systems: the role of clouds in the radiation budget of the atmosphere; oceanic circulation patterns and the redistribution of energy within the oceans; the fluxes of water and energy between the atmosphere, biosphere, and land and ocean surfaces; and others.

Biochemical Dynamics: the fluxes of radiatively and chemically active species between the atmosphere, biosphere, and land and ocean surfaces; the atmospheric cycling and transformations of radiatively and chemically important trace species, and others.

---

[2] The government agencies currently conducting scientific research into global change under the Committee on Earth and Environmental Sciences are the Departments of Agriculture, Energy, Interior, the Environmental Protection Agency, the National Aeronautics and Space Administration, the National Oceanic and Atmospheric Administration, and the National Science Foundation

Ecological Systems and Dynamics: the structure and function of biological systems on various time scales; the response of species, ecological communities, and natural and managed ecosystems to carbon dioxide, climate, and physical/chemical stresses, and others.

Earth System History: natural variability of climate on all time scales; responses of ecosystems to climate change; changes in the composition of the Earth's atmosphere; changes in sea level; and others.

Human Interactions: establish long-term comparable, cross-national data bases that encompass human activities such as land-use practices, energy transformations, legal and regulatory requirements, and others.

Solid Earth Processes: coastal erosion and wetland loss caused by sea level change; the role of subaerial and submarine volcanism in contributing radiatively important gases, aerosols, heat, and fluids to the atmosphere and the ocean; permafrost change and atmospheric methane concentrations; and others.

Solar Influences: the effect that changes in the influx of solar radiation has had on environmental changes.

Data management and modelling are of paramount concern within the global change research community. Of particular concern is the weakness in the results from the Atmospheric General Circulation models. Figure 2 illustrates the linkages between the Atmospheric General Circulation Models, the Oceanic General Circulation Models, and the models of the land surfaces. GIS will play a critical part in adequately modeling the surface of the Earth. Of course, since it is change that is being modeled, the GIS must include a temporal dimension.

As shown in Figure 3, the study of global change is a continuing process that advances in cycles where

- data is gathered from the research sites; based upon this information, relevant processes are first hypothesized and later tested with further data.
- processes are integrated into models; this integration requires an appraisal of a process' relevance to the model, an evaluation of the convenience of aggregating it with other processes, and the determination of the proper level of spatial and temporal aggregation both for the basic and the derived information.

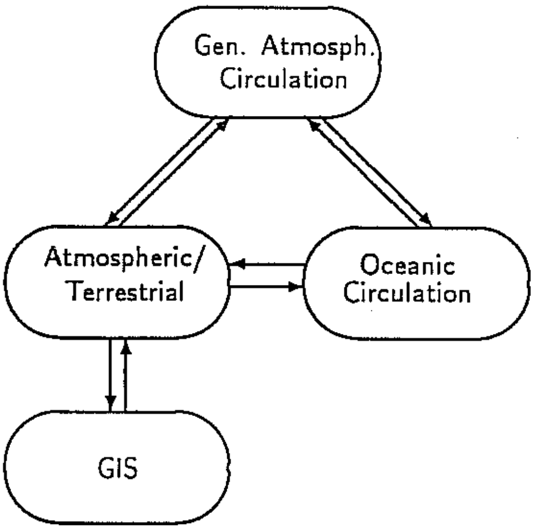

Figure 2: Atmospheric Models

- the model is put into operation and its results analyzed.

- the results are assessed, and as a consequence, new studies may be designed.

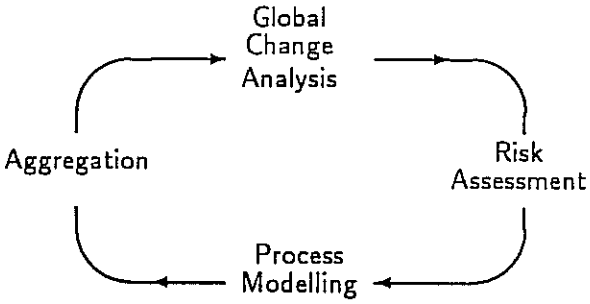

Figure 3: Model cycle for Global Change

Modeling global change has four characteristics:

i) it relies on the integration of quantitative models based on the physical and biological sciences;

ii) it is inherently open: no limitations are set on the type of phenomena to be explored, on the kind of models to be built, or on the type of data to be used;

iii) it needs simplification and aggregation mechanisms. The efficiency of a model may benefit from dropping some considerations and from integrating several processes. For the same reason both time and space need to be aggregated or refined to an adequate resolution; and

iv) it is a multidisciplinary effort that will last for several decades.

3 Requirements of temporal GIS

In addition to the presentations described in section 2, a whole session was devoted to the description of several typical applications. Rather than discussing these applications in detail, this section will present the conclusions reached after their analyses. These requirements relate to the type of systems to be modelled, to the richness or scarcity of data, to the models of time, and to the logical formalism needed in temporal GIS; therefore, a subsection is assigned to each one of these items.

3.1 Quantitative and Qualitative GIS

The spectrum of user requirements from temporal GIS has two extremes:

On one end lie GIS that imply a discrete representation of time and of space. In them, time is considered as a sequence of intervals, with events that punctuate their beginning and their end. The characteristics of a modelled entity remain unchanged within a temporal interval; any modification to these characteristics constitutes an event. The properties of this time are a consequence of a set of axioms. Depending on the chosen axioms, time can branch to the future or to the past, or a 'next' interval may be defined, etc. Sometimes this type of GIS establishes only a partial order among events. It is then not always possible, based on this partial order, to decide the relative position of any two pairs of intervals, or to conclude whether they intersect or not.

In these systems, the model of space prescribes a controlled subdivision of spatial objects. Spatial objects have topological relations defined among them, e.g. connectedness, disjointness, etc., and a metric is not strictly necessary. These GIS resort to mathematical logic to describe temporal behavior, either by prescribing the possibility or impossibility of future events, or by utilizing cause-effect relationships. The models used in these GIS are inherently *qualitative*; the cadastral GIS (Section 2.1) as well as some aspect of geological GIS exemplify these systems [TURN90].

On the opposite end of the spectrum are GIS that consider a continuous time and a continuous space. Their model of time is similar to that of Physics; it admits only one past and one future and has a measure associated to it. This metric provides both a total order on sets of disjoint intervals and a manner of deciding whether or not two intervals are disjoint. Space, in turn, can be indefinitely divided and the topological relations between its subdivisions expressed also by a metric. These GIS resort to differential equations to describe the evolution of the modelled phenomenon through time; their data is used either as inputs to these equations or as means for verifying the adequacy of the model. In short, the models used in this case are *quantitative*; they can be exemplified by those used in the Global Change project (Section 2.2).

The two types of systems described above will be referred to respectively as *qualitative* and *quantitative* GIS. We warn the reader that these two types are the extremes of a spectrum, that they have been discussed only as means of illustration, and that most applications lie somewhere between them.

3.2 Data-rich vs. Data-poor

Remote sensing, particularly that from satellite borne sensors, has changed the manner in which the geosciences use information: instead of being data-starved they have become data-rich.

In a data-poor environment, the scarcity of data is compensated for by a deep knowledge of the underlying phenomena. The little data available in it is used to deduce rather complex results, by visualizing sequences of potential processes and selecting one of them; this selection requires a moderate amount of evidence backed by a large set of rules. An example of this environment is found in Geology in general [FRAN90] [TURN90]; geologists must draw maps of regions of the earth's crust based on a few isolated observations, e.g. the outcrops. In data-poor environments the scientists are aware of both the phenomena to be studied and the processes to be observed; in this paradigm, the paucity of data is compensated by a large body of *a priori* hypotheses based on common sense or on previous experiences.

In the data-rich environment the scientist is deluged with a flow of highly unstructured information that arrives periodically, usually from several sources or sensors. She does not know beforehand what to look for in the stream(s) of data; the only selectivity criteria is that of being *interesting*, i.e., not following the expected trend. Once interesting patterns are found, the processes that govern them have to be explained; these processes are frequently unknown and can only be inferred by a repeated observation of the occurrence of these patterns. Planetary exploration gives an example of this environment: the astronomers who accessed Voyager-1's data were suddenly flooded by a plethora of results that had to be carefully examined to locate interesting phenomena. The causes of these phenomena first had to be hypothesized to later test its validity with further data analysis. In a data-rich environment, found mostly in quantitative temporal GIS, there is no *a priori* knowledge of the kind of the phenomenon of the studies, much less of the rules that govern it.

To be suited for a data-rich environment, a temporal GIS must

i) be inherently open; no limitations should be set on the phenomena to be explored, on the models to be built, or on the type of data to be used.

ii) handle simplification and aggregation for both the spatial and temporal domains. The reasons are multiple: a need for reducing the volume of data that may be collected at the rate of terabytes per day; the elimination of unnecessary considerations from a model or the need for fusing several submodels into one; and the adequacy of the level of aggregation to that of the involved discipline, since data-rich environments need to cater to several scientific branches and since each discipline has its own level of temporal and spatial aggregation. (Section 4.3).

iii) provide exploratory capabilities that facilitate the search for interesting features.

iv) assess data quality [CHRI83] i.e., its fitness of use for a purpose.

This is not to say that all temporal GIS are data rich. Indeed, many important applications of temporal GIS are known to exist in "data-poor" environments. The applications in geology discussed in the previous paragraphs are a case in a point. Another example is that of cadastral systems (Section 2.1); their information is very limited and facts are only collected if they are relevant under the current legal rules: there is no possibility of exploratory data analysis or of hypothesizing. Somewhat similar are archeological applications [ZUBR90], where facts are located in time and space, but where many intervening events may be missing.

3.3 Models of time

Cadastral systems require a time modelled by a sequence of intervals in which an entity's data remain invariant; the changes that delimit an interval do not occur periodically. The temporal intervals considered in the global change example, on the other hand, are punctuated by regular events (days, years, etc.); in this example, a new temporal interval simply announces another tick of the clock and does not necessarily imply a change in the data.

Many other models of time, beside the two cited above, may be needed by other applications: a time that is cyclic, to describe recurrent processes, such as the climate per day or year; a time that branches into the future, to explore alternative scenarios; a time that is infinitely divisible, to model instantaneous change; a time that terminates, to model limited horizons, etc.

An axiomatic treatment of the models of time can be found in [RESC71]; the inclusion of extra axioms allows the definition of the models of time as lattices of types.

3.4 Querying requirements

Depending again on the application, the processing of queries in a system may be based on

First order predicate calculus, for queries that can be solved by considering time as an extra attribute, e.g., "Did John register any deeds in Orono during 1990?".

Modal temporal logic [WORB90], for queries as "What was the first deed that John registered?".

Dynamic logic [WORB90], if the application contemplates actions that modify the contents of the database, such as, "if I modify 'A': will 'B' be affected by the integrity constraints?".

Non-monotonic logic [SHOH88], if reasoning is to be done in data-poor environments and the system is forced to "jump into conclusions" by providing an extra body of hypotheses based upon common sense or past experiences.

4 Representations of time

The discussions on this topic are divided into three categories: the perception and models of time, the representation of time by metaphors, and the need for manipulating several kinds of time in a GIS. The three categories are discussed in the following subsection.

Visualization was found to play an important role in the discussions on the representations of time. It is clear that GIS need to incorporate a rich and varied set of options for viewing maps, for displaying statistical graphics, and for performing animation; it is also clear that all these methods will have to be extended to include time and temporal data. The discussions on the subject were restricted to those touching temporal matters, and in general, observations on visualization in GIS were excluded.

4.1 Temporal abstractions and temporal perception

To describe time to a computer and to perform temporal reasoning, the concept of time must be abstracted and formalized, i.e., a subset of the characteristics and operations of time needs to be described rigorously. The most widely used abstraction of time is that of Physics, which consists of a set of symbols with its own operations and logic. Physical time is thought of as a linear continuum with a direction, often represented by the real numbers. Although physical time has been found to be an appropriate model for a multitude of phenomena besides those of physics, it is not suited for expressing all situations.

Physical time is based on exact measurements and the topology of its intervals is a consequence of this metric; there are disciplines, e.g., Geology, Archeology, or Anthropology, in which the measurements of time are inexact, incomplete, or inexistent. These disciplines use a partial order, not given by a metric, to describe relationships among events.

The perception of time by humans, on the other hand, is not founded on abstractions, but rather on primary experiences encountered during the initial years of life [LAKO87]. Temporal abstractions like those mentioned above are constructed in later ages; however, to be comprehended or communicated, these abstractions must be metaphorically related to early-life experiences.

One of the findings of the NCGIA initiative on Spatial Languages was the existence of more than one perception of space; it seems likely that there are also several ways of perceiving time [MARK89].

4.2 Graphical representation of time

The most widely used metaphors represent time by space. These metaphors are the ones favored by cartographers. Cartographers represent time in maps in three forms [VASI90]:

> **Time as the base map.** In this alternative, a map does not depict change; time, however, is the frame of reference within which information is placed. For example, even if a road map does does not contain temporal information, the time in which it is valid can be deduced from an inspection of the list of revisions supplied by the editor.
>
> **Time is a mapped variable.** For example, the dates of independence of a continent's countries can be represented by different colors.
>
> **The combined time map.** Here, time is both a mapped variable and part of the frame of reference.

The representation of time by space is not restricted to maps, e.g., PERT diagrams represent time as bars and show time as flowing from left to right. A slightly similar representation, Petri nets, is used to represent causality in addition to temporal precedence.

Cartographers use a panoply of symbols and conventions (arrows, etc.) for the representation of time by space. It would be extremely convenient to know what sort of symbols ought to be associated with (or dissociated from) a particular kind of temporal information. An absolute categorization seems infeasible. First, the representations of time in terms of space may be culturally dependent: e.g., there are peoples for whom the future lays behind them while they face the past [LAKO80]. Second, metaphorical views such as seeing time as moving past the observer ("time flies"), or perceiving it as a static entity in which the observer moves ("we move ahead in time"), are not contained in all languages.
Some consideration was given to the general possibilities for representation of time on maps:

Space by time. In this metaphor, the inverse of time-by-space, space is explored using a "travel" allegory.

Time by time. In these metaphors, used in "virtual reality", the temporal dimension is a controlled variable which the user can traverse in both directions and at variable rates.

Time by other graphical variable. Some other metaphors were cited, as "time by color", "time by intensity", "time by degree of deterioration", etc.

4.3 The need for several kinds of time

The range of applications in temporal GIS requires the use of different abstractions of time, such as the continuous and discretized times of Section 3.3. These different models correspond to data types that are elements of a lattice; each of these types has a set of operations (e.g. next, previous, etc.), that can be inherited by its successors in the lattice. Two principles of the object-oriented paradigm advocated in the section 5 can be of a great help in representing time: the principle of inheritance will prevent the redundant programming of operations, while the principle of abstraction will isolate the user from unnecessary details of the internal implementation of temporal data types.

In the remainder of this work we will make a distinction between a *model* of time and a *kind* of time. A "model" of time is a mathematical abstraction expressable by axioms; a "kind" of time is a human category based on the uses given by the application.

Temporal GIS are embedded in institutions and carried out by people; the roles assigned by them will dictate the kinds of time that are necessary. Even if all users and applications of a GIS employ the same abstraction of time (e.g. a dense, linear time such as that of physics), there would still a definite need for having several kinds of time. An example in which time is modelled as linear and dense in all applications, and yet three kinds of time is discussed in [LEST90]:

World time, needed for describing the real occurrence of an event,

Institutional time, necessary for describing the legal validity of an operation,

Database time, used for time-stamping a transaction.

The use of these three times is be illustrated by a night deposit into a bank: A client delivers the money at 8PM (world time); the bank is responsible for it after it opens, at 8:30 AM next day (institutional time); finally, the deposit is recorded into the database at 9:15 AM (database time). In this case it is not the model of time that is different. The data type that models the time of the institutional and world time is exactly the same; the consequences associated to the different kinds of "time of deposit" is what makes them different.

Other examples of the necessity of several kinds of time can be found in [LANG88] and [ALTA90a]. Two kinds of time, named as *valid* and *transaction*, have been identified in temporal databases [SNOD85]; the latter is equivalent to Lester's database time, while Lester's world and institutional time are two instances of valid time. As will be seen in the next section, reasoning with these two kinds of time will require two different logical formalisms.

Finally, besides a connection between time and roles, there is interplay between the use given to information and its required level of temporal (and spatial) aggregation. It was noted that scientific disciplines are intrinsically associated with a level of aggregation and a scale of events in time or space. A seismologist is interested in milliseconds while a geologist studying plate tectonics deals in thousands of years, yet both study the same phenomenon.

5 DBMS requirements posed by temporal GIS

The discussion on requirements led to the characterization of a number of features needed by the software system that manages a temporal GIS. Some of this features are

A need for identity. A GIS, to support temporal change, must have mechanisms to decide whether an object has merely altered its appearance or has become an entirely new item (Heracleitus' observation that the river we step into is the same one as yesterday and yet it is not). The answer to this question depends on the application: two applications may render different answers when considering the same real-world entity.

Existing systems have attacked this problem following either a value-based approach, similar to that of the relational model [DATE90], or an identity-based approach, like the one of OO DMBS [KIM90]. In the value-based approach, entities are totally characterized by the values of their attributes; the user is responsible for providing an *ad hoc* value (a key) to tag an entity so it can be recognized after changes. In the identity-based approach, the system perceives each object with a personality of its own, regardless of external appearances; the system bears the responsibility of tagging individual objects.

The participants concurred that, of these approaches, the one provided by object-oriented systems was the better suited for the implementation of temporal GIS; they also agreed that many problems lurked in the identification of a system's objects, i.e., in answering questions such as: Is this an object or an object's feature?, or, Are these two objects the same?, etc.

These problems are consequences of how identity is supported in object-oriented systems [WEGN90]: *every* object, when created, is provided with an immutable, system-given identifier that serves as its handle; this non-reusable identifier is generally kept hidden from the programmer. The participants described, during the discussions, practical situations in which providing objects with immutable identities ranged from irrelevant to annoying since an object's identity may change with the perceptions or the needs of a user.

And yet, a system-provided identity appears as a lesser evil; its alternative, providing entities with user-given keys, is a source of potentially greater problems.

There have been proposals [STEI89] to mix both approaches: having some objects with an immutable identity, and characterizing others by *references*, i.e., user-given keys similar to those found in relational systems.

Diverse times. Different GIS applications need different models of time: some, like those related to oceanography, will need a linear, infinitely divisible time; others, such as cadastral systems, will need a linear time in which events occur at irregular times; some others, like those used to forecast the consequences of actions, will necessitate a time that branches into the future, etc.

To be useful in the implementation in temporal GIS, a DBMS should also be able to manage several temporal variables within a DBMS; an example of this requirement is the application of Section 4.3 that needs to combine institutional, world, and database time.

Diverse Logical Formalisms. Existing temporal DBMSs are based on first order predicate calculus [SNOD87]. These systems are useful in a wide range of queries; if temporal reasoning is to be supported, however, other logical tools are needed.

The way temporal events are related in time determines the logic needed for temporal reasoning, e.g., modal logic is adequate if the relationship is that of *possibility*; if *causality* is involved, then dynamic logic is necessary.

Managing approximate information. This requires two capabilities:

i) The support of approximate views and approximate queries. Both are consequences of the exploratory requirements. Views, to be discussed in Section 6, are mechanisms to reduce and combine the data so it conforms to a user's needs; the definition of approximate views is convenient since the user is not absolutely certain of the conditions that produce an interesting feature. Approximate queries are needed since the user may not want (or does not know how) to formulate a query rigorously and instead formulates demands such as "Give me the items that look like X" or "Give me those that are close to Y".

ii) The management of quality-associated information. This, in turn, needs mechanisms for mixing data with diverse qualities that should be provided when data of several sources is combined, and for procedures for propagating data imprecision when quality-associated views are used. This requirement is similar to that of *tagged* information found in statistical databases [SHOS85].

Aggregation operations. These are needed to reduce the amount of information to be processed and to set the spatial and temporal resolution of the data to the levels relevant to the user's scientific discipline. In particular, support is needed for

i) Constructing aggregations. This, in turn, requires facilities for defining complex objects, e.g., templates that indicate how the elements of a lattice or the components of a grid are recursively aggregated into larger objects, and how to derive the new attributes at the different levels of aggregation. This also requires facilities for checking the consistency of user-defined complex objects, e.g. verifying the compatability of the composing objects, the correct semantics of the aggregated attributes, etc.

ii) Evaluating aggregate attributes. This involves supplying efficient algorithms for computing these attributes, consistency constraints and procedures for fixing constraint violations, and algorithms for propagating changes along aggregation hierarchies.

iii) Operations on complex objects that include modification of their structure, aggregation in other complex objects, etc. Of particular interest are queries that provide approximate results based on a superset of the aggregations [VRBS90] or that provide quick estimates on the values of aggregate functions [ROWE85].

The identification of the aggregate-processing capabilities needed in a temporal GIS is a research issue; a starting point would be to consider

i) The extensions of the relational model for statistical databases [OZSO85], [OZSO89], that, among others, provide constructs for specifying the format of statistical tables, a library of built-in aggregate operators, and specially tailored query-optimization strategies.

ii) The facilities provided in OO DBMS for the creation, modification, and version-maintenance of complex objects [KIM88].

iii) Existing categorizations of temporal propositions, depending on how an assertion's validity propagates from a given time-interval to others containing it or contained within it [SHOH88]. For example, a proposition is *downward-hereditary* if its validity over an interval implies its validity on each of the subintervals; a proposition is *clay-like* if whenever it holds on two consecutive intervals, it holds in their union.

iv) Existing mechanisms for propagating changes across derived objects [ELLI90] [HUDS90] [ROSE89].

6 Existing Techniques

Sections 3 and 5 posed the requirements of temporal GIS; this section will review some available software technology relevant to these requirements. As could be expected, some of these features come from other non-conventional applications of DBMSs, mainly from those used for CAD CAM.

Views. Views are DBMS-provided mechanisms for deriving information from raw data and/or other views. Views are needed for implementing a user's concept of *interesting*, for disposing of unnecessary detail, for including the context that makes visual information intelligible, and for rendering spatial or temporal information at the level of aggregation needed by the application.

Consider, for example, a regional authority interested in studying the propagation of an insect pest. The authority has access to a database containing sightings of insects; each sighting consists of a date, a geographical location, an insect's species, and a rough estimate of the insect's population per unit area. An appropriate view would select the records related to the pest, group them by county and month, and count the records at each group. Finally, the view will display both the infested areas and the new appearances, with a regional map as a background, using a time-as-time or a time-as-space metaphor (Section 4).

Rules. A rule has three components: an *event* that specifies the circumstances under which a rule should be considered, a *condition* on the state of the GIS that determines whether or not the rule is applicable, and an *action* to be taken if the condition is fulfilled [DAYA88]. An *active* database is one that, among other types of information, stores triggers and actions of rules. Rules fulfill a double purpose in a temporal GIS: to enforce the consistency of information and to specify how a user's view is to be implemented, whether by computing it anew or by storing past results. In the latter case, a rule also specifies the moment in which changes are reflected into the stored view. Rules may also be used to propagate changes along aggregate hierarchies, specifying how changes in the lower-level objects propagate upwards.

If a GIS must have exploratory capabilities, rules should be dynamically changeable and easy to browse; if a user's concept of interesting is to be implemented by a view, users should be allowed to have their own sets of rules.

Long transactions. An exploratory process may have an extended lifetime and require transactions that span across several sessions, referred to as *Long transactions* [BANC85]; one method to achieve this property is to provide information with *check-in*, *check-out* mechanisms. These mechanisms are implemented by means of data-structures called *locks*; long-transaction locks, in opposition to those found in conventional DBMS's, must be kept into stable storage. To work in cooperative environments, each session of a long transaction may have several *short* transactions [ROTZ90]. The changes applied during a short transaction are hidden from all other transactions by means of *short locks*; these short locks are released after a pseudo-commit, the results of the short transaction being then visible to the other members of the cooperating group. Pseudo-committed information is kept invisible from users outside the cooperating group prior to a long transaction's commitment.

Temporal DBMS and reasoning systems. Temporal GIS can be used in a wide variety of environments, with requirements ranging from information retrieval to temporal reasoning [WORB90]. Temporal DBMS's [ELMA90a] are ideal for information-retrieval environments since they are designed for query processing and since they are implemented with a primary consideration on efficiency. Reasoning systems are designed for generality; although they can be used for the same purposes of a DBMS, their performance in these tasks, compared to that of a DBMS, is one or two orders of magnitude lower. A temporal GIS needs both querying and reasoning capabilities; its design should make both approaches available, and the GIS should apply each one of them where they are most efficient.

Object Orientation. Object Orientation is the most useful mechanism available for modeling space and time. The reasons are manifold: first, it clearly separates the properties of an operation from its implementation details, allowing the design of simpler and cleaner information systems. Second, it allows the inheritance of operations across type hierarchies; therefore, the different models of time (continuous, discrete, branching, etc.) can be modelled uniformly. Finally, it can implement views by means of aggregate objects.

It was noted that object identity, as implemented in existing OO systems, should be used with discretion , and that OO systems, to be applicable in exploratory environments, should add types just as easily as objects since there is no *a priori* knowledge on the type of objects needed.

7 Conclusions

Conclusions

Here are some of the conclusions reached by the participants on the characteristics and requirements of temporal GIS:

- The detection of two different sets of users' requirements, named as the *quantitative* and *qualitative* cases.
- The need for exploratory capabilities.
- The need for supporting views, rules, spatial-temporal aggregation operators, and long transactions.
- The need to manipulate inexact or approximate information, including provisions for approximate queries, for assessing the quality of data (i.e., its fitness for use), and for evaluating the error propagation to derived data.

The discussions in the workshop centered on the GIS requirements for effectively representing temporal-spatial information; many important issues were left untouched, e.g. the implementation of storage structures, of temporal indices, of efficient methods for solving queries. Finally, the participants agreed on the necessity of

i) a detailed analysis of the requirements of temporal GIS prior to the formulation of a set of design specifications. This analysis should detect the kinds of processes likely to be captured in such a GIS, and, as a result, articulate a circumscribed set of constraints for the GIS design.

ii) the establishment of a continuing interaction between the members of the GIS and computer science communities. Two activities were suggested:

- the documentation of the requirements of temporal GIS by a set of examples. Two examples of requirements' specification, related to statistical and scientific databases respectively, can by found in [SHOS82] [SHOS84].
- that GIS and database specialists attend each others professional meetings.

Acknowledgments

The authors thank John Kelmelis for his comments on an earlier version of the draft and Ms. Sue Comins for proofreading this report.

Bibliography

[ALTA90a] K. Al-Taha, R. Barrera, "Temporal Databases: their applicability to the ISDS", Internal memorandum, Department of Surveying Engineering, University of Maine, June 1990.

[BANC85] F. Bancilhon, W. Kim, H. Korth, "A model for CAD transactions" Proceedings of the 11^{th} international conference on very large databases, pp. 25-33, August 1985.

[CES89] The Committee on Earth Sciences, "Our Changing Planet: A U.S. strategy for global change research", January 1989.

[CHRI83] N. Chrisman, "The role of quality information in the long-term functioning of a geographic information system", Cartographica 21(2,3), pp. 79-87, 1983.

[DATE90] C. J. Date, An introduction to database systems, 5^{th} edition, Vol. 1, Addison Wesley, 1990.

[DAYA88] U. Dayal *et al.*, "The HiPAC project: combining active databases and timing constraints", SIGMOD Record, 17(1), pp. 51-70, March 1988.

[ELLI90] H. J. C. Ellis *et al.*, "Extending the behavorial capabilities of the object-oriented paradigm with an active model of propagation", Proceedings 18^{th} annual computer science conference, February 1990.

[ELMA90a] R. Elmasri, G.T. Wuu, "A temporal model and query language for ER databases", Proceedings 6^{th} international conference on data engineering, pp. 76-85, January 1990.

[ELMA90b] R. Elmasri, G.T. Wuu, Y. J. Kim, "The time index: an access structure for temporal data", Proceedings of the 16^{th} international conference on very large databases, pp. 1-12, August 1990.

[FRAN90] A. Frank, T. Buyong, "Geometry for three dimensional GIS in geoscientific applications" in Proceedings, NATO advanced research workshop on three dimensional modeling with geoscientific information systems, Keith Turner, Ed., Luwer, Dordrecht, Netherlands.

[HINT90] R.J. Hintz and H.J.Onsrud, "A methodology for upgrading real property boundary information in a GIS using a temporally efficient automated survey measurement management system", Department of Surveying Engineering, University of Maine, 1990.

[HUDS90] D. L. Hudson, "Autonomous view-stages: materialized support for view update propagation", Ph. D. Thesis, Department of Surveying Engineering, University of Maine, May 1990.

[HUNT90] G. J. Hunter, I.P. Williamson, "The development of a historical digital cadastral database", International journal of geographic information systems, 4(2), 1990.

[KIM88] W. Kim, H.T. Chou, J. Banerjee, "Operations and implementation of complex objects", IEEE transactions on software engineering, 14(7),, pp 985-996, July 1988.

[KIM90] W. Kim, Introduction to object-oriented databases, MIT Press, 1990.

[LAKO80] G. Lakoff, Metaphors we live by, U. of Chicago Press, 1980.

[LAKO87] G. Lakoff, Women, fire, and dangerous things: what categories reveal about the mind, U. of Chicago Press, 1987.

[LANG88] G. Langran, N. R. Chrisman, "A framework for temporal geographic information", Cartographica, 25(3), pp. 1-14, 1988.

[LEST90] M. Lester, "Tracking the temporal polygon: a conceptual model of multidimensional time for geographic information systems", presented at the Temporal Workshop, NCGIA, University of Maine, October 1990 (see appendix).

[MARK89] D. M. Mark, A. U. Frank, M. J. Egenhofer, S. M. Freudschu, M. Mc Granaghan, R. M. White "Languages of spatial relations: Initiative 2 specialist meeting report". National Center for Geographic Information and Analysis. Report 89-2, (1989).

[ONSR89] H.J. Onsrud, "The Land Tenure System of The United States", Zeitschrift FORUM des Bundes der offentlich Bestellten Vermessungsingeneure, 1/1989, pp. 23-28, 1989.

[OZSO85] G. Ozsoyoglu, A. M. Ozsoyoglu, "Statistical Database Query Languages", IEEE Transactions on software engineering, SE-11(10), pp. 1071-1081, October 1985.

[OZSO89] G. Ozsoyoglu, V. Matos, A. M. Ozsoyoglu, "Query processing techniques in the summary-table-by-example database query language", ACM Transactions on database systems, 14(4), pp 526-573, December 1989.

[RESC71] N. Rescher, A. Urquhart, Temporal Logic, Library of Exact Philosophy, Springer-Verlag, 1971.

[ROSE89] A. Rosenthal, U.S. Chakravarthy, B. Blaustein, J.A. Blakeley, "Situation monitoring for active databases", Proceedings 15th international conference on very large databases, pp. 455-464, August 1989.

[ROTZ90] K. Rotzell, "Transactions and versioning in an OODBMS", Proceedings OODBTG Workshop SPARK/DBSSG/OODBTG, pp. 55-61, October 1990.

[ROWE85] N. C. Rowe, "Antisampling for estimation: an overview", IEEE transactions on software engineering, 11(10), pp. 1081-1091, October 1985.

[SHOH88] Y. Shoham, N. Goyal "Temporal reasoning in artificial intelligence" in Exploring artificial intelligence, H. Shrobe, Ed., Morgan-Kaufmann, pp. 419-438, 1988.

[SHOS82] A. Shoshani, "Statistical databases: characteristics, problems and some solutions", Proceedings 8th international conference on very large databases, pp. 208-222, September 1982.

[SHOS84] A. Shoshani, F. Olken, H.K.T. Wong, "Characteristics of scientific databases", Proceedings 10th international conference on very large databases, pp. 147-160, August 1984.

[SHOS85] A. Shoshani, H. K. T. Wong, "Statistical and scientific database issues", IEEE transactions on software engineering, SE-11(10), pp. 1040-1047, October 1985.

[SMIT90] T. Smith, A. Frank, "Very large spatial databases: report from the specialist meeting", Journal of visual languages 1(2), pp. 291-309, 1990.

[SNOD85] R. Snodgrass, I. Ahn, "A taxonomy of time in databases", Proceedings SIGMOD conference, pp. 236-246, May 1985.

[SNOD87] R. Snodgrass, "The temporal language TQUEL", ACM transactions on database systems, (12)2, pp. 247-299, June 1987.

[STEI89] J. Stein, T. L. Anderson, D. Maier, "Mistaking identity", 2nd international workshop on database programming languages, pp. 161-164, June 1989.

[TURN90] K. Turner, Ed., Proceedings of the NATO advanced research workshop on three dimensional modeling with geoscientific information systems, Luwer, Dordrecht, Netherlands.

[USDA74] U. S. Department of Agriculture, U. S. Department of Commerce, "Land Title Recording in the United States", Special Studies No. 67, 1974.

[VASI90] I. Vasiliev, "Examples of the treatment of time as a variable on maps", presented at the 1990 AAG in Toronto (see appendix).

[VRBS90] S.B. Vrbsky, J.W.S. Liu, K. P. Smith, "An object-oriented query processor that returns monotonically improving approximate answers", IFIP OODB Conference, July 1990.

[WEGN90] P. Wegner, "Concepts and paradigms of Object-Oriented programming", OOPS Messenger, 1(1), pp. 7-87,August 1990.

[WORB90] M. F. Worboys, "Reasoning about GIS using temporal and dynamic logics", presented at the Temporal Workshop, NCGIA, University of Maine, October 1990 (see appendix).

[WYLE76] S. Wyler, B. Wyler, The Swiss Civil Code, english version. Remak Verlag Zurich, 1976.

[ZUBR90] E. Zubrow (editor), Interpreting space, Taylor & Francis, 1990.

APPENDIX

Presentations