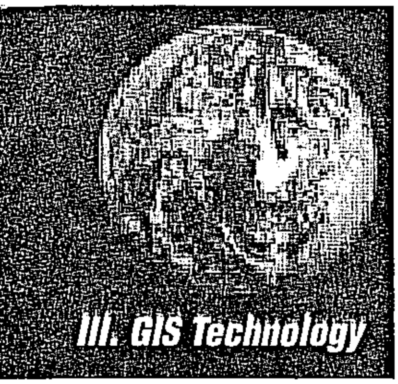

Temporal GIS Keeps Data Current

By Khaled Al-Taha and Andrew Frank

"To draw conclusions from temporal data, a logical framework with inference rules is used."

A GIS is not just a collection of maps stored in a computer, but a database that should correctly model a particular region of the world. Since the world is changing continuously, GIS applications require that data collection.be updated. Every GIS is, by definition, a database that manages spatial data and includes facilities to update its data so it continues to reflect the current situation. Only a GIS that includes time and space in its model of reality can be called a "temporal GIS." It must contain the current information *and* keep the data that describe the previous states and when they were applicable. A temporal GIS answers "*when*-and-*where*" questions, such as: What were the boundaries of New York City in June 1908? Where were New England forest areas in 1880? Which properties were owned by John Doe in Baltimore from August 1981 to June 1988? When a GIS includes time, it can demonstrate the enormous advantage of a database that models reality compared to a conventional collection of printed maps.

GIS data will need updates regularly. Is every GIS, therefore, a temporal GIS? For a temporal GIS, data would have to be related to time intervals. A temporal GIS contains data that not only describe the present situation, but also past or future situations, and allow queries that can refer to past, present or future data (Figure 1). This is an extension of a spatial database qualification, where data must refer to space and where spatial query methods are present.

What Applications Need a Temporal GIS?

Like all temporal databases, the major goals of a temporal GIS are data storage and retrieval with respect to spatial location and temporal validity, and the extraction of the world's situation at a given time in response to temporal queries. Such facilities are needed to analyze and understand phenomena by tracking their change over time, to predict future systems' behavior by looking into their past, and to plan and take action that will lead to a desired future situation (Shoham and Goyal, 1988).

Managing time-variant information, which describes dynamic, ever-changing situations, is necessary for many GIS users. Most scientific models of the world link static and dynamic aspects, i.e., describe future states in terms of present and past states and the rates of change between them. In many applications, such as urban dynamics (Forrester, 1969) and global-change studies, time-related data are crucial and must include the notion of "change" (e.g., the popular computer game SimCity, which simulates the development of a city in time and space). If we look only at a present state, no change rates can be computed. However, temporal data use in GIS applications varies and has different requirements. Applications range from those that require a small database for recording the history of a town, to extremely large temporal and spatial data collection for studying global change.

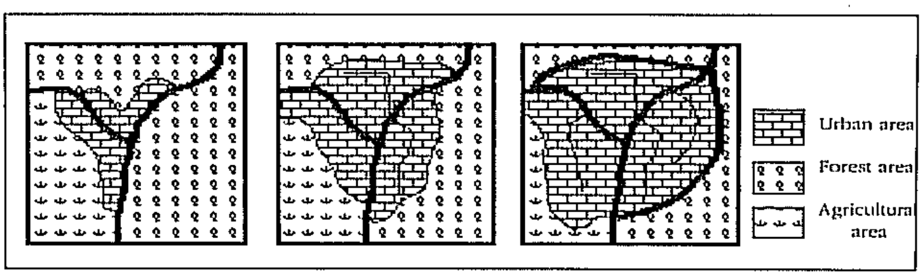

Figure 1. A temporal GIS allows queries that can refer to past, present or future data. For example, urban area expansion and land-use history easily can be determined.

We separate two groups of temporal GIS users: those who deal with small-scale mapping and collect global data for large areas of the Earth (such as politicians, meteorologists and geographers) and those who deal with large-scale mapping and local data collections for smaller, specific areas (often planners or engineers). The data processing of the first group differs significantly from the applications used by the second group, and the temporal data management requirements are different. The differences are size and type of data, precision and abstraction level, type of queries posed, temporal aspects of the data and the temporal-reasoning model. This article describes these differences and presents an example for each case. The two temporal GISs will be referred to as *globally oriented*, such as research and science applications, or *detail-oriented*, such as engineering and administrative applications. The grouping is somewhat arbitrary, and many applications will show properties of both. The ultimate temporal GIS will be able to deal with both cases, but that is currently not a realistic goal.

Globally Oriented GIS

Scientists who study global phenomena, such as climate changes, need a database that represents large areas such as a country, a continent or the Earth, and will handle a variety of data such as topography, meteorology, vegetation, animal habitat and their changes over time. Scientists compare data in a specific area to detect changes, e.g., monitor the change of forest boundaries (Figure 1), trace a tree disease, follow the course of migrating birds or study the advances made in the distribution of "killer bees."

The data in such applications primarily describe measurable quantities and are often — especially for global-change applications (Mounsey and Tomlinson, 1988) — reported as an average value for a regular area of the world (raster cell). Scientists start with data sets for certain points in time and space and compute new data sets with methods based on mathematical analysis and calculus. One can think of each data set as describing a value of a variable for a large set of discrete points in time and space, but otherwise similar to a simple variable in an equation from "Physics 101." Dana Tomlin (1990) has formalized such operations using conventional map overlay methods as inspiration. It is necessary to extend this map algebra to include time in the data and in the operations, i.e., include rate-of-change operations.

Scientists use a temporal GIS in an explorative mode, building models and verifying them with available data, to understand relations between different phenomena and to identity the effect of one variable on another, e.g., test the effect of acid rain on the destruction of plant species, the relation between the ozone hole and rising temperature or other questions related to environmental challenges that affect our existence.

Detail-oriented GIS

GIS for engineering and design applications deal with smaller areas and handle a smaller variety of data. Because of the detail level involved, these GISs also may be large (e.g., the size of a GIS for a county may range between one and a few hundred gigabytes). An example of an engineering GIS is a *multipurpose cadastre* (National Research Council, 1980) that combines layers of information involving cadastral data (ownership, mortgages and spatial property representations), topographic data (houses, roads, trees, lakes and rivers), land use (agricultural, residential, commercial and industrial use), utility lines (telephone and power lines, water and sewer installations) and future plans (planned zones for new roads, dams, residential and industrial areas, and natural reserves).

The major differences between a detail-oriented GIS and a globally oriented GIS are the following.

Specific requirements. Engineering or administrative databases aim to handle and solve specific problems with clearly defined entities. Their operational goals often are defined in laws prescribing the management of some resources. Therefore, such systems have clearly defined requirements.

Incremental changes. Most updates on engineering GISs are represented and expressed in discrete time units, such as years, days or hours, and the time units are semiuniform. For a data set, it is possible to apply uniform time granularities, e.g., "minutes" can represent deed-recording time, "days" can represent power and telephone lines, and "months" and "years" can represent the building-date of a house.

Data use is mostly for administrative decisions requiring facts such as: When did John buy parcel 61, or who were the owners of the houses at the north side of Main St. on April 22, 1970?

Time Dimension Complexity

Any database has to deal with a large amount of data. The major issues in every database are managing data (e.g., data structures and the physical design and storage mechanism), optimizing response time and efficiency, guaranteeing integrity (methods to ensure there are no internal contradictions), safeguarding against system failure and developing concurrency control (a mechanism to avoid problems when multiple users attempt to update the same data element at the same time). It is advisable to use a database management system (DBMS) for a GIS to:

- assure that data storage and data treatment in application programs is separated (Codd, 1982);
- provide a general-purpose query language and optimization methods for execution; and
- include a transaction manager which, despite hardware or software failures, guarantees the integrity of the data stored even if accessed or updated by different users at the same time.

In addition to these services provided by the DBMS for administrative uses, nonadministrative applications like GIS require additional services, primarily for handling spatial data (Frank, 1988). The temporal aspect of databases can be formalized using a knowledge-based or information-based approach. The first approach was developed by the artificial intelligence (AI) community, and the latter comes from the DBMS community (Barrera and Al-Taha, 1990).

A Knowledge-based Approach

To draw conclusions from temporal data, e.g., in processing a query, a logical framework with inference rules is used. Logicians have developed various models in temporal reasoning and methods to handle temporal data. Historically, the first attempts to handle temporal information in AI studies adopted a change-based approach, which incorporates time as an attribute only instead of a separate dimension; only the changes in the databases were recorded. Until the next change, the last state of the world remains unchanged. Examples of change-based approaches are *situation calculus* (McCarthy and Hayes, 1981), which views the world as a set of states or situations, each of which is a snapshot of the world; and *dynamic logic* (Pratt, 1976), a framework for reasoning about programs based on modal logic. New states are generated by change indicators, such as *actions* in situation calculus and *programs* in dynamic logic. The change-based approach supports only instant actions with instantaneous results and does not support concurrent or overlapping actions.

Later, the time-based approach was introduced, which considers the passage of time as its only major change indicator. At any given time, a user can infer a conclusion based on the knowledge stored in a database represented by true and false assertions or propositions. To implement a time-based approach, one must choose a temporal logic, e.g., classical first order logic, reified sentences or modal temporal logic (Shoham and Goyal, 1988), and time primitives (time points, time intervals or both) for time representation and interpretation. Examples of the time-based models are the interval-based model after J.F. Allen (1984), the point-based model by D. McDermott (1981) and the temporal model by Y. Shoham and N. Goyal (1988).

An Information-based Approach

The information systems referred to in this section are the systems that use a classical DBMS that supports the retrieval of stored information in response to a user's query. In a DBMS, the answer to a query is based only on the data that already was stored in the database.

Temporal DBMSs developed in the past decade are subdivided into four main designs: *static*, *rollback*, *historic* and *temporal* (Snodgrass, 1987). The major difference stems from the time perspectives that exist in these databases: *valid* time or real-world time, which describes when an action happened in the real world; and *transaction* or database time, which describes when the information was entered into the database (Figure 2). A static database does not deal explicitly with time, a rollback database deals only with the time a fact is acquired, a historic database is only concerned with when a fact became valid, and a temporal database contains both valid and transaction time.

		transaction time
	Static Database	Rollback Database
valid time	Historic Database	Temporal Database

Figure 2. Time perspectives and the various temporal databases.

Static database. A static database, or snapshot database, represents a snapshot of the world at a particular time. In such a database, one can ask, "What is the population of Orono now," but not "What was the population of Orono last year?" Time is used in a static database as an attribute, and old information is overwritten; therefore, it is not possible to see what the state of the database was at an earlier time.

Rollback database. A rollback database stores all past states, indexed by time, of the snapshot database as it evolves (Figure 3a). Deletion of old information is not permitted; instead, a new state is generated and the old state is preserved. Such an approach needs a representation of transaction time. In a rollback database, one can get a snapshot of the database at a given time *t*, i.e., what was recorded at that time. Changes may be made only to the most recent snapshot state. An example of a query that can be answered in a rollback database is, "What was the population of Orono in 1951 according to the 1980 version of the database?"

Historic database. In many applications, such as historical geography, the time the information was valid is more relevant than the time it was recorded. Another way to store temporal information is in a historic database organized by "valid time" (Figure 3b). Errors can be corrected as soon as they become evident, which is not possible in a rollback database; however, no record is kept of the errors that have been corrected. Unlike rollback databases, historical databases cannot be interpreted as they were in the past. So in a historical database, one still can ask, "What was the population of Orono in 1951?" However, he or she may get different answers depending when the query was asked and when errors were corrected.

Temporal database. A temporal database combines the properties from a rollback and a historical database. It handles both transaction and valid time (Figure 3c). So in a temporal database, one can view information to be valid as of a past date (a property of rollback databases), and one also can go back and correct errors as soon as he or she learns about them (a property of historical databases). Temporal databases can support queries such as, "Was the population count performed and published in the same year?"

To demonstrate the principal differences between a historic and a rollback database, consider the population of Orono in 1980 and in 1990. Assume Orono's population was 10,000 in 1980. The 1990 count was first published in 1991, and the population was 11,000. Assume that an adjustment to this count was announced in 1992, and the population was corrected to 10,500 (Figure 4). For a historic database, because any wrong information is corrected as soon as the right information becomes available, Orono's population was 10,000 in 1990 (according to the count of 1980 and before the new count took place), 11,000 in 1991 (after the new count was published in 1991) and 10,500 in 1992 (after the new adjustment was announced in 1992). Therefore, if one asks a historic database what the population of Orono was in 1990, the answer will vary depending on when the question was asked. If the question is asked in 1990, the answer will be 10,000. Similarly, if the question is asked in 1991 and 1992, the answer will be 11,000 and 10,500, respectively.

On the other hand, for a rollback database, new information will generate a new database state. So there will be a 1990 state where the population is 10,000, a 1991 state where the population is 11,000 and a 1992 state where the population is 10,500. Therefore, the answer will be the same (10,000) if one asks the same question to a rollback database, despite the date the question was asked since there is only one 1990 state, which has not changed.

In a temporal database, if one asks what Orono's population was in 1990, the answer will be 10,000; if the same query is asked as of 1991 and 1992, the answer is 11,000 and 10,500, respectively. This facility is important for legal and administrative decisions. Decisions are made with the best data available at the time. If a decision is appealed, it must be considered in the frame of what was known then — not based on hindsight.

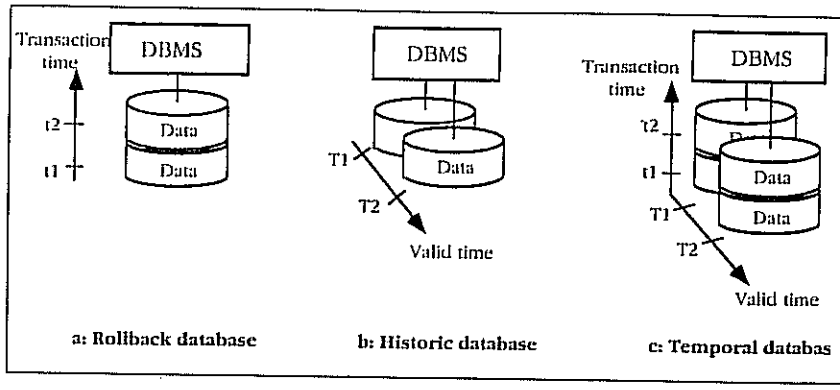

Figure 3: An overview of temporal databases.

Temporal GIS History

The first work of note on temporal GIS technology is that of U. Basoglu and J. Morrison (1978). They built a U.S. Historical County Boundary Data File that contains changes to county boundaries during the past 200 years, using a "historic database" design. Another early database that implements time encoding is the Norwegian Socioeconomic Database, which handles administrative boundaries for the period 1770-1980.

G. Langran (1989) has explored a wide range of applications for temporal GIS, such as forest resource management, urban and regional management, research and development, electronic navigation charts and transportation. M.F. Worboys (1990) has shown the role of modal logics in a GIS and concluded that, at least for some GIS applications, a modal logic is necessary to capture the meaning of the data. G.J. Hunter and I.P. Williamson (1990) have presented a method of storing and processing temporal geographic data by the addition of time-encoding attributes to data elements. They also have created a historical digital cadastral database to show this method. The National Center for Geographic Information and Analysis has selected the problem of temporal data as the focus of one of its research initiatives. Recently, R. Barrera and K. Al-Taha (1990) have presented a literature survey on temporal issues in the fields of AI, DBMS and GIS. Al-Taha also is working on applying various methods used in AI and DBMSs to develop a temporal reasoning framework for cadastral systems.

Next Steps

The combination of the tools for AI and DBMS and their respective models provide the power needed for future GIS. There is an urgent need for temporal GIS, both for engineering administrative and global-science applications, but there also are user groups from the social sciences, primarily historic geography, anthropology and archeology, that need access to data collection with powerful temporal and spatial analytical methods. Engineering and administra-

What was the population of Orono in 1990?

date we ask the question	Historic database	Rollback database
1990	10000	10000
1991	11000	10000
1992	10500	10000

What was the population of Orono in 1991?

date we ask the question	Historic database	Rollback database
—	—	—
1991	11000	11000
1992	10500	11000

What was the population of Orono in 1992?

date we ask the question	Historic database	Rollback database
—	—	—
—	—	—
1992	10500	10500

Figure 4. Different answers due to various time perspectives and error correction methods in a historic and a rollback database.

tive applications appear to be easier because their specifications can be derived from their function within the administrative or legal context. On the conceptual level, it is not easy to determine which of the different temporal logic systems should be used; for the implementation, the interaction between spatial and temporal operations must be explored and formalized. On the other hand, scientific global databases seem feasible because they are based primarily on a concept of time similar to the one used in physics; therefore, calculus methods become applicable.

Finally, database implementation strategies and the data structure and performance issues that are specific for GIS have not been addressed. One must assume, however, that there are interactions between spatial and temporal data structures and access methods. At first glance, temporal GIS for administrative use — at least for all systems where the information is used to determine legally relevant actions — must make valid and transaction time available, and queries of the type, "Who was owner of parcel 723 on May 5, 1988, as recorded on May 15, 1988?" answerable. A GIS to store scientific data to study global change or other large-area phenomena likely will not need this facility. It appears that the structure of such databases is determined more by the data available than the present needs of the researchers — not to exclude new uses. In this environment, methods to inform scientists about available data and visualization methods become crucial.

Khaled Al-Taha is a graduate assistant and a Ph.D. candidate in the Department of Surveying Engineering at the University of Maine. Andrew Frank is an associate professor in the Department of Surveying Engineering at the University of Maine. They can be contacted at the University of Maine, 107 Boardman Hall, Orono, ME 04469.

Shoham, Y. and Goyal, N. 1988. **Temporal Reasoning in AI.** *Exploring AI.* Shrobe ed. Morgan Kaufmann: pp. 419-438.

Forrester, J. W. 1969. *Urban Dynamics.* Cambridge, Mass., M.I.T. Press.

Mounsey, H. and Tomlinson, R. eds. 1988. **Building Databases for Global Science.** *Proceedings of the IGU Global Database Planning Project.* Tylney Hall, Hampshire, UK. London: Taylor & Francis.

Tomlin, C. D. 1990. *Geographic Information Systems and Cartographic Modeling.* Englewood Cliffs, N.J.: Prentice Hall.

National Research Council. 1980. *Need for a Multipurpose Cadastre.* National Academy Press, Washington, D.C.

Codd, E.F. 1982. **Relational Database: A Practical Foundation for Productivity.** *Communications of the ACSM.* 25 (2): pp. 109-117.

Frank, A. 1988. **Requirements for a Database Management System for a GIS.** *Photogrammetric Engineering & Remote Sensing.* 54 (11): pp. 1557-1564.

Barrera, R. and Al-Taha, K. 1990. *Models in Temporal Knowledge Representation and Temporal DBMS.* NCGIA Technical Report 90-8, University of California, Santa Barbara.

McCarthy and Hayes. 1981. **Some Philosophical Problems From the Standpoint of Artificial Intelligence.** *Readings in Artificial Intelligence.* pp. 431-450.

Pratt, V. R. 1976. **Semantical Considerations on Floyd-Hoare Logic.** *Proceedings of the 17th Foundation of Computer Science.* IEEE: pp. 109-121. New York.

Allen, J. F. 1984. **Towards a General Theory of Action and Time.** *Artificial Intelligence.* 23 (2): pp. 123-154.

McDermott, D. 1981. *A Temporal Logic for Reasoning about Processes and Plans.* Computer Science Department, Yale University, RR 196.

Snodgrass, R. 1987. **The Temporal Query Language TQuel.** *ACM Transactions on Database Systems.* 12 (2): pp. 247-298.

Basoglu, U. and Morrison, J. 1978. **The Efficient Hierarchical Data Structure for the U.S. Historical Boundary File.** *Harvard Paper for GIS.* Dutton ed. Addison-Wesley. Massachusetts.

Langran, G. 1989. *Time in Geographic Information Systems.* University of Washington, Ph.D. Thesis.

Worboys, M. F. 1990. **Reasoning about GIS Using Temporal and Dynamic Logics.** *Proceedings of the Temporal GIS Workshop.* University of Maine, Orono, Maine, Oct. 12-13.

Hunter, G. J. and Williamson, I.P. 1990. **The Development of a Historical Digital Cadastral Database.** *Int. J. Geographical Information Systems.* 4 (2): pp. 169-179.

1991-92 International GIS SOURCEBOOK™