Investigating U.S. Geological Survey Needs for the Management of Temporal GIS Data

Abstract
The U.S. Geological Survey is investigating its needs for management of temporal information in the National Digital Cartographic Data Base. Future production methods and customer requirements will probably require a capability to capture the history of individual features within the database. Updates to the digital database may be accomplished on a feature by feature basis, rather than through a revision of the entire map sheet, as has been the case in the past. Under this scenario, the ability to support queries involving the historical evolution of map features will require additions or modifications to the existing database structure.

Introduction
The U.S. Geological Survey (USGS), under the National Mapping Program, provides base cartographic and other geographic information covering the land area of the United States. The National Mapping Program also maintains thematic vector products such as land-use/land-cover maps. For this information to be useful, it must be updated on a periodic basis. However, there are many applications for which it is necessary to recover previous versions of the information to determine changes which have occurred over time. Traditionally, the need for historical information has been satisfied by retaining copies of previous versions of hard copy maps. The National Mapping Program is currently moving toward an all digital database, the National Digital Cartographic Data Base (NDCDB), which will be made available to the public in digital form, and support the production of traditional paper products. Although such a digital database has the potential to solve many of the problems of currency and availability of data, it also complicates the process of maintaining historical data. In an ideal world, the NDCDB would be updated continuously as changes occur and are reported, resulting in a perpetually current database. However, such a continual updating process may not be economically feasible. Even periodic updating of only selected features, however, complicates the process of recreating past states of the data unless a new archive version of the database is saved after each such update. Knowledge of these past data states is important to permit the reconstruction of historical data for analysis of many problems. The analysis of this temporal updating problem is the major emphasis of the research effort described in this paper.

The USGS also maintains large holdings of raster products. These include satellite data such as Landsat and Advanced Very High Resolution Radiometer (AVHRR) imagery, and derivative raster data such as Normalized Difference Vegetation Index (NDVI) data. The national Digital Orthophoto Quad (DOQ) program, currently underway, is being accelerated in cooperation with the Department of Agriculture. This program will provide a digital image database of major portions of the country which will be updated at regular intervals. These DOQs provide a source of digital image data of major portions of the country for use by the public and other Federal agencies. Also, the USGS has adopted the DOQ as the standard base media for updates to the base cartographic map series for the country. These image data are digital in nature, and are entered into the database as new images are acquired and processed. Unlike the vector data discussed above, these raster data sets have naturally occurring versions, corresponding to each acquisition of imagery. Analysis of change in the data is accomplished by a comparison of images taken at different times.Because raster data sets tend to have naturally occurring versions, and because the management of raster data has unique characteristics which are distinct from the management of vector data, the current research project concentrates primarily upon the analysis of the temporal nature of the USGS vector GIS data holdings.

The research project currently underway will define the needs for temporal cataloging of the National Mapping Program's vector data holdings, investigate the state of the art in temporal GIS systems, and outline a program through which the USGS can develop the systems necessary to maintain appropriate temporal information in the database of the future.

Current Programs in Digital Data Generation and Update at the USGS
The USGS is currently undertaking several initiatives which will lead to a largely digital mapping environment by 1995. These initiatives are required, not only to supply the users with the digital data which they require, but also to increase production efficiency. The five major thematic layers of the 1:100,000-scale series of digital data will be completed by 1995. These data will consist of the culture, hydrography, hypsography, transportation, and vegetation layers. Under a new, multi-agency initiative, the USGS has proposed a joint effort to complete the large-scale digital base cartographic data of the country by the year 2000. The USGS will be responsible for the base topographic data (1:24,000-scale USGS quadrangle data in digital form), the Bureau of Land Management will be responsible for the Public Land Survey Data,

Photogrammetric Engineering & Remote Sensing,
Vol. 59, No. 10, October 1993, pp. 1503–1507.

0099-1112/93/5910–1503$03.00/0

Clifford W. Greve
John A. Kelmelis
Robin Fegeas
Stephen C. Guptill
Nancy Mouat
National Mapping Division,
U. S. Geological Survey, Reston, VA 22092.

and the National Ocean Survey will be responsible for digitizing the Nautical Charts.

In addition to these programs, a joint program is underway with the Census Bureau to integrate the Topologically Integrated Geographic Encoding and Referencing System (TIGER) data with the 1:100,000-scale database, by making TIGER data conform to the metric accuracy of the map data while preserving the information content of the TIGER data. The USGS is pursuing the development of techniques to permit the merging of these data sets.

In order to accomplish these enormous tasks, the USGS will rely increasingly on the private sector. Much of this work will be accomplished through normal contracting procedures, wherein a private firm will be placed under contract to perform a service or to deliver a complete digital product. In addition to these normal contracting procedures, however, the USGS is instituting a practice called "Innovative Partnerships" in order to take advantage of the large amounts of digital data being generated by numerous private and public entities within the United States for various purposes. Data obtained under the Innovative Partnership arrangement can consist of entire database layers, but will more frequently consist of additions to existing layers. Each submitted addition may cover many conventional map sheets. For instance, a major power line, highway, or pipeline would add data to single layers of many map sheets, but would not constitute the totality of any layer of any map sheet. Such additions will be inserted into the NDCDB by the National Mapping Division (NMD) to generate the updated database.

The Nature of the Problem

In the initial stages of populating the NDCDB, updates are being made on a quadrangle basis. In those cases where digital data already exist in the database, the archival record of the existing version is preserved by making a copy of that version of the digital information on some media. Thus, there is still a capability for retrieval of previous versions of the database. As time progresses updates will be made to the data through various means, creating still additional versions. So long as these versions are preserved, the capability to recover past states of the data is preserved. The Innovative Partnership arrangement is a good illustration of the need to add a temporal capability to the current data model. In many cases, data from Innovative Partnerships will provide partial updates for one or more layers covering several quad sheets. It would be desirable to add these data to the database as soon as they are available, to ensure that the database is as current as possible for all potential users. However, to preserve an archival copy of all of the affected quad sheet files each time such an update is made will produce a much larger volume of archival snapshots than would be generated by the current practice of performing complete updates to quad sheets on a periodic basis. In fact, one can envision a situation in which geographic data are updated on a transactional basis, in much the same way that textual information for such applications as inventory control is updated in current practice. Under these circumstances, maintaining any viable chronology of the data would become very unwieldy if a copy of the entire quad sheet file had to be made each time that a transactional update occurred.

For the NMD to be able to maintain a current base of cartographic vector GIS data, a data structure must be developed which will incorporate the capability of reconstructing the data as they existed at any chosen time in the past. Even though the total database will not be populated until the year 2000 under optimistic assumptions, the incorporation of temporal information into the data entered prior to that date is necessary to provide a basis for future updates (Langran, 1992). Thus, it is important that NMD determine how the temporal data will be handled, so that the database can be built with the necessary temporal information included.

Preliminary Investigation of Temporal GIS Technology

Although a far more thorough review of the literature will be undertaken during the execution of the first phase of this project, a discussion of the background knowledge accumulated to date will serve to place the problem in perspective, and to provide a basis for the requirements discussion to follow.

In addition to the spatial aspects of time in GIS, there are several questions as to the nature of what is necessary to document the history of geographic data. First is the debate of what time tags should be recorded for each feature. There are at least three times which are relevant. The first has been called logical time (Snodgrass and Ahn, 1985), or event time (Copeland and Maier, 1984), and many other terms. This is the time at which the change actually occurred. Note that the disappearance of a feature is also an event, so that duration of existence of a feature is implicitly included in event time. Second is observation time, or the time at which the event was observed (Langran, 1992). Third is physical time or database time, or the time at which the event was added to the database (Snodgrass and Ahn, 1985; Langran, 1992). These three times need not be, and frequently are not, identical. In fact, upon further consideration, they need not even be ordered consistently between events (Newell *et al.*, 1992). It is entirely possible that a scientist could first enter the current status of the information into a database, and at a later date add historical data depicting the state of the information at a previous time. In this case, database time for the historical data would have a more recent date than the database time assigned to the original state of the data, which would be the inverse of the relationship between the event times. Likewise, the time of observation of two changes to the physical reality represented by the data need not reflect the order of the event times at which these changes actually occurred. Thus, in the general sense, three time tags may need to be represented within the data structure.

Recent literature has emphasized the need for records of evidence and process information in the temporal domain (Kelmelis, 1991; Langran, 1992). In this context, not only is the temporal duration of the three-dimensional domain important, but also what evidence led to the change in the database, and what caused the change. Looking at the context of a two-dimensional database, with time as the third dimension, the change in land cover from forest to barren ground is an event in time. The evidence might be aerial photography, but a record of this evidence would not normally be included in current attribution schemes in typical GISs available today. Further, the fact that the process (for example, a fire) which caused the change may need to be stored in the database increases the complexity of storing all relevant information.

There are many possible views of the problem of time in GIS. The intended application of the GIS affects the view that might be chosen. Potential temporal GIS capabilities include inventory, analysis, updates, quality control, scheduling, animation, and static mapping (Langran, 1988). Inventory involves storing a complete description of a study area, and accounting for changes in the state of the area. Analysis requires the ability to explain and forecast a region's processes. Updates require that outdated information be replaced by

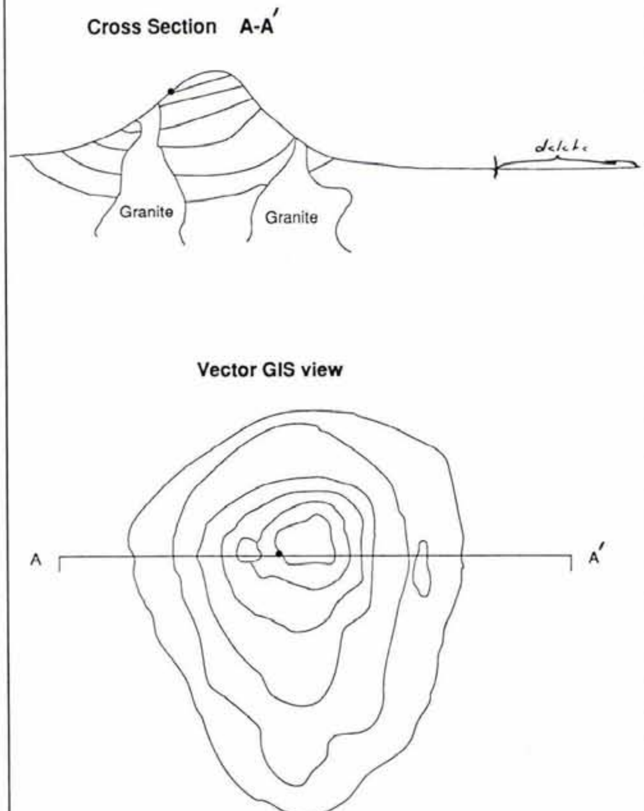

Figure 1. A cross section of geological formations, shown with the normal two-dimensional vector GIS representation.

current information and also be retained and not discarded. Quality control requires that current data be consistent with previous versions, and often requires information, such as lineage, in addition to the geometry and attributes commonly contained in GISs. Scheduling involves detecting threshold database states which trigger some action. Animation involves a display of a dynamic summary of processes. Static mapping is the generation of current versions of classic cartographic products. Each of these capabilities place demands upon the structure of the temporal GIS.

In general, the problem of time in geographic information systems can be viewed as adding a temporal dimension to the general two- or three-dimensional GIS problem (Kelmelis, 1991). The usual representation of the *Z* dimension in vector GIS data is as an attribute attached to point coordinates or to features, but is not considered in the topological structure except in some cases to ensure locational consistency at nodes. These current systems might be referred to as 2.5-dimensional GIS. The general three-dimensional case is defined herein as the problem of dealing with geographic data in a context where the bounding polygons of familiar two-dimensional vector GISs are replaced by general three-dimensional solids. As far as has been determined in the research conducted to date, the concept of topology in this three dimensional space, while it has been discussed, has never been implemented in a general solution to the problem. Adding the fourth dimension of time merely serves to add further complexity to the already complex general three-dimensional problem.

In the most general case of temporal GIS, there are four continuous dimensions (Kelmelis, 1991). Such a system would be required, for instance, if it were required to view the changes in a three-dimensional data structure continuously through time. Real-time aircraft simulators, where objects move continuously through three dimensions over time and it is desired to view the process continuously (or at least at very small increments in time), would at first appear to be an implementation of this case of four continuous variables. These simulation problems are soluble today through the use of massive computer power, and the simplifying fact that many of the elements, such as the terrain, are static in the simulation environment. The moving elements, such as aircraft, are modeled as unchanging solid objects which move only in position and attitude during time. These simplifying assumptions, coupled with the fact that the number and type of features is restricted, makes this special case of the general four continuous dimensional problem soluble with today's technology and GIS models.

As mentioned above, in the typical 2.5-dimensional vector GIS problem encountered in practice today, the only continuous variables are *X* and *Y* (or latitude and longitude or whatever horizontal reference is desired). *Z* is usually treated as an attribute. In those rare instances where there are multiple features defined at a point in horizontal space, they are treated as an exception. Topology is defined with the highest dimensional figure being the two-dimensional polygon. The concept of multiple values of a single attribute for a single horizontal area (such as rock type as a function of elevation) is in general undefined. The concept of a continuous view of the data in the *Z* dimension is not contained in the database structure. As illustrated in Figure 1, the subsurface profile plot of rock strata could not be made from the two-dimensional topology, even though each point on the bounding curves has the elevation of the surface recorded. Only the information relative to the surface exposure of the rock formations is available spatially (although strike and dip might be recorded as attributes). However, for most GIS problems, where the third dimension is essentially an attribute which does not enter into queries in a continuous sense, the current structure suffices as a good compromise between an efficient solution and total mathematical rigor.

Just as current systems do not treat *Z* as a continuous variable, current systems do not treat time as continuous. The current practice in the temporal domain is to undertake GIS efforts as projects. Photography is acquired, and the data are analyzed and placed into a GIS, possibly by updating an existing database for the project area. Change analysis from prior time periods is accomplished by loading the previous database from its archival media, and using it as a basis for comparison with the new data. Based upon this comparison, the change of the data with time is determined. In this current handling of time, not only is time excluded from the topological structure, but it is also quantized under the condition that all features are tagged with the time at which the project is archived.

The next level of temporal quantization amounts to time tagging features, in much the same way as elevations are tagged in current two-dimensional GISs. If all features, attributes, and topological relationships were tagged, it would be possible to reconstruct the GIS database as it existed at any point in time. It would be somewhat similar in the *X,Y,Z* case to asking for all features which lay between 1000 and 1500 metres of elevation, but not exactly, because placing elevation tags on points does not allow for recovering the topology at, for instance, the 1250-metre elevation plane to be reconstructed. One might refer to this formulation, imposed on the standard 2.5-dimensional GIS, as a 2.5 + 0.5-dimensional GIS.

At this 2.5 + 0.5-dimensional level, there have been some implementations. Langran (1990) describes such a system that was designed by Intergraph Corporation for the National Ocean and Atmospheric Administration. This system stores temporal information in feature attribute relational tables. The system is somewhat restricted in that dead features do not reference graphics other than by a bounding rectangle; however, they do reference the current version of that feature. Thus, although there is a record that a feature existed in the past, the topology cannot be reconstructed as it would have been at that past time. If, on the other hand, a feature moves, but still remains a feature, the previous location is preserved through linkages to previous geometry tables.

The current NMD Digital Line Graph (Enhanced version) concept includes a Feature History Flag providing information about the origination of a given feature in a data set (i.e., new, revised, modified, deleted, or unchanged). This information relates to the latest preceding version only, i.e., it goes back only one generation. Thus, unless all versions are preserved, there is no traceability from the current database to the entire history of the feature, without consulting each precedent version in turn. Spatial changes are not retained, only the historical record referencing a permanent feature identifier.

Some GIS systems, such as SMALLWORLD, have at least some capability of handling version numbers for features in the data, which translates to a rather coarse quantization in the time dimension. In the SMALLWORLD concept (Newell *et al.*, 1992), each user can modify the basic database, and that user's modifications will be stored in a particular version of the database for that user. If there are subsequent versions of the database, it is imperative that no modifications be made to the antecedent version, or it may invalidate any subsequent version. The combining of version trees from different users for an eventual update of the primary database is discussed briefly. If, for instance, one user has modified a feature which another user has deleted, the resolution of that conflict may well require human intervention. Thus, although versions allow temporal update without locking other users out of the data base, it is the opinion of the authors that they scarcely, as implemented by SMALLWORLD, solve the universal problem of time management in GIS. It is, however, a beginning.

Preliminary Assessment of USGS Needs for Temporal GIS

As mentioned above, the USGS is currently involved in several initiatives, most notably the Innovative Partnerships initiative, which will likely lead to a requirement for dynamic updating of its geographic database. To prepare for this requirement, adequate capability for processing the data must be developed, without allowing the resulting system to become so overly complex that it will be difficult (if not impossible) to develop, or inefficient to use after it has been developed.

USGS is not attempting to solve the general temporal GIS problem for users of spatial data. The USGS National Mapping Program role is to gather, consolidate, and archive data so that they are available for users, both public and private. Any analysis which USGS performs on the data under the National Mapping Program is for purposes of validation and research, and thus the internal data structure need not support the full gamut of possible applications. It is important, however, that the data structure be sufficiently robust to portray the data and its temporal attributes, so that the user can reformat the data in response to a particular application while retaining the necessary information content of the original data.

Thus, while some application may require time to be treated as a continuous variable, to do so within the National Mapping Program internal model appears at this time to be unnecessary, because the USGS has no information which would permit the changes to features contained in the data to be observed in a continuous manner. The USGS is unlikely to have available information to trace the continuous nature of building a road, such as each truckload of earth moved, but only information which indicates the status of the road at the times at which the status was observed. Thus, it would appear that time tags on both attributes and objects in a 2.5-dimensional database would be adequate to represent the state of the features at the times at which observations were made. As discussed above, such a database might be called a 2.5 + 0.5-dimensional database.

There is still the problem, however, of what time to represent in the system. Traditionally, map versions have corresponded to what was introduced above as observation time. The map legend states that the map was made with photography of such and such a date, with field completion on another date. That field completion date, for all practical purposes, constitutes the time tag for the features checked, with date of photography being the time tag for other features. The time at which the map is actually completed and entered into the database becomes the database time. Event time is not currently a function of the USGS map making process, as it is not typically acquired unless by accident in the process of reporting of sources for map information. However, event time is recorded for the dates on which primary control data are established. Event time may also be available for some classes of features derived under the Innovative Partnership program described above.

It would be simplest, in a GIS for USGS's National Mapping Program purposes, if only one time were required. There are applications for which such a system might not be adequate, however. If only database time were recorded, for instance, potential users of the data would have to verify the actual event time of any features appearing in the data if the exact time of the creation of the feature were critical for the application. If the time given were observation time, the above problem might be alleviated, because in many instances observation time might be the closest approximation to event time which was available. However, modifications to the database are made in database time, and thus the modification records to the topology will be in database time, no matter what the time tag might be. Thus, if only observation time were stored, it might be difficult to determine the appropriate time periods in database time to use to detect change for a specific feature. The ramifications of these issues on the design of the GIS structure are yet to be determined, but there seems to be a higher complexity to tagging by observation time than to tagging by database time. It would appear, however, that having observation time available to the user is a requirement, so dual tags may be required. Preliminary conclusions are that observation time may be included as an attribute to the features, while the time tags to the topological relationships would be in database time.

As mentioned above, event time is not generally known to the USGS and, therefore, would not be a likely candidate for inclusion in the database. However, Innovative Partnerships might provide event time for some features. Whether USGS should store that information for those features while compiling or updating other features with only observed time, which would result in a data base containing incomplete or even inconsistent information, is critical. The question is whether the best information available should be stored and, if so, would such a mixed set of temporal information be more confusing than illuminating or more expensive than cost effective.

Conclusion

The discussion in this paper describes the first efforts in a continuing research program within the USGS to define the characteristics of a temporal GIS structure for the NDCDB. Clearly, the definition of a temporal data model to meet the needs of the NDCDB is not simple and will require a careful balance between cost and capability. As technological capabilities change, additional opportunities for data acquisition and information storage may arise. In addition, new GIS and process modeling capabilities along with increased user sophistication will place new and greater demands upon the data. Standardization efforts will also have an impact, in that the temporal requirements will have to be reflected in revisions to current standards and in the definition of new standards. The conclusions of the current research effort, then, must include assessment of current needs and capabilities, costs and benefits, and projected future technologies and applications. The eventual system must meet current requirements, be flexible enough to permit modification to meet future requirements, and be implementable within the limited resources available. There will be many difficult decisions to be made along the way to the solution.

References

Copeland, G., and D. Maier, 1984. Making Smalltalk a Database System, *Proceedings of the Sigmod 84 Conference*, June, pp. 316–325.

Kelmelis, John A., 1991. *Time and Space in Geographic Information: Toward a Four-Dimensional Spatio-Temporal Data Model*, Ph.D. Thesis, Pennsylvania State University.

Langran, Gail, 1988. Temporal GIS Design Tradeoffs, *Proceedings GIS/LIS 1988*, pp. 890–899.

———, 1990. Tracing Temporal Information in an Automated Nautical Charting System, *Cartography and Geographic Information Systems*, 17(4):291–299.

———, 1992. States, Events, and Evidence: The Principle Entities of a Temporal GIS, *Proceedings GIS/LIS 1992*, pp. 416–425.

Newell, R.G., David Theriault, and Mark Esterfield, 1992. Temporal GIS-Modelling the evolution of Spatial Data In Time, *Computers and Geosciences*, 18(4):427–433.

Snodgrass, Richard, and Ilsoo Ahn, 1985. A Taxonomy of Time in Databases, *Proceedings of ACM-SIGMOD International Conference on Management of Data*, 14(4):236–246.