ISSUES IN SPATIO-TEMPORAL DATABASE SYSTEMS: DATA MODELS, LANGUAGES AND MOVING OBJECTS

José Moreira
Universidade de Aveiro, Portugal

Cristina Ribeiro
Faculdade de Engenharia da Universidade do Porto, Portugal

Jean-Marc Saglio
Ecole Nationale Supérieure de Télécommunications de Paris, France

Michel Scholl
Cedric, Conservatoire National des Arts et Métiers de Paris, France

Keywords: geo-referenced systems, data management, spatio-temporal database systems, moving objects, data models and query languages

Contents

Summary

Spatio-temporal is a broad designation that may refer to a wide variety of systems dealing with the temporal evolution of spatial phenomena. Space and time are deeply ingrained in our commonsense reasoning patterns, and we have developed effective ways of taking their properties into account to manage their intrinsic complexity. It is therefore no surprise that representing and managing data where the temporal and spatial aspects are fundamental is a hard task. There are numerous application domains motivating the developments in this area. The problems posed by the integration of spatial and temporal dimensions of data require new solutions from database management systems. One application area which is receiving considerable attention concerns moving objects, represented as points in the spatial dimensions.

This article presents an overview of spatio-temporal database systems, describing the main features of spatio-temporal objects and proposing a classification of spatio-temporal systems based on the properties of the represented objects. This classification allows understanding the complexity of the dynamics that underlie the temporal evolution of spatial phenomena. A survey of some issues in the spatio-temporal domain is also presented. The focus is on data models and languages for spatio-temporal database systems with moving object applications in mind. Open issues for future research, in the light of the features described as useful to being handled by spatio-temporal systems, are also identified.

1. Introduction

With the increasing use of automated processes in fields like earth sciences, economical and socioeconomic studies, urban planning, traffic control, land information systems, environment protection or medical imaging, and the rapid increase in the availability of data from a wide variety of sources like satellite images, mapping agencies or GPS devices, we have witnessed in recent years an increase in the demand for systems that can model, store, manipulate, and interpret spatial data. In addition, spatial data often has an evolutionary behavior. As a consequence, not only must spatial information be permanently updated to represent the current status of real objects, but the history of their evolution also needs to be stored. Therefore, the new systems require functionalities for processing spatio-temporal data suitable for answering queries of the type "where and when," such as: "What were the boundaries of Porto City in January, 1990?" "Which properties were owned by John Smith in London from May 1992 to June 1997," or "Which ships were within a distance of 20 miles from Sines' harbor at 10 AM 11 June, 2001." This is the concern of what is commonly referred to as spatio-temporal database systems.

There is a general agreement in the database research community that efficient database management systems (DBMS) exist for the representation, manipulation, and querying of large volumes of attribute (also referred to descriptive) data, that is, numeric, alphabetic or alphanumeric data. Traditionally, these systems have been developed to capture a single snapshot of reality and they were not well suited for applications requiring the support of past, current, or even future data. The temporal databases community has studied this issue and there is now a consensus temporal extension (TSQL2), allowing the formulation of sophisticated queries over time, as well as modifications to previous states (if an error is detected or if more information becomes available) and to future states (for planning purposes). Spatial database research has been mainly concerned with the development of spatial data models, efficient access methods and algorithms for the implementation of complex spatial operations. These systems have become commercially available, first as proprietary GIS, such as Arc/Info from ESRI GIS & Mapping, and more recently as spatial extensions that can be integrated within traditional database systems, such as the Spatial Data Option from Oracle or the DB2 Spatial Extender from IBM.

A simple combination of temporal and spatial theories, as shown in this article and in other studies, is not sufficient to embrace all the new issues raised by the integration of spatio-temporal information in database systems. Presently, the challenge for developers of spatio-temporal systems is to find abstractions and architectures to implement generic systems, that is, to build systems with generic spatio-temporal data management capabilities that can then be tailored to the requirements of a particular application domain. Important issues in this context include the handling of spatio-temporal representations and data models, spatio-temporal query languages, and efficient access methods (see also *Spatio-Temporal Information Systems).* Such efforts require an accurate knowledge of the intrinsic features of spatio-temporal information.

This article starts with a characterization of spatio-temporal data, identifying the main features to be taken into account and their impact on a database representation. The various approaches to capturing the spatial and temporal features of data in a database are analyzed, together with the proposed data models and languages. The focus is then put on the somewhat simpler problem of representing moving objects and the specific cases of monitoring systems (dealing with the history of objects' movement) and real-time systems (where present and near-future events are central). The open and ongoing research is highlighted in the conclusions.

2. Characterization of Spatio-Temporal Systems

A search for existing and potential applications of spatio-temporal database systems shows that the application domains are varied and that the objects to be represented have greatly differing properties. A classification would therefore be useful, not only to bring together applications with identical requirements, but also to identify the elementary components that must be handled by spatio-temporal systems, facilitating a modular development and a better understanding of the extent of the issues that these systems must deal with.

For the purpose of classification of spatio-temporal systems, the view here is centered on the spatio-temporal aspects of objects. Physical objects may be characterized by various attributes other than the spatio-temporal ones, storing values for object features of different kinds: descriptions, physical quantities, and in general domain-specific attributes. These may also change over time, but the analysis in the classification presented here is focused on the temporal evolution of spatial properties.

2.1. Properties of Spatio-Temporal Objects

The classification of spatio-temporal systems can be based on the properties of the objects they represent. The two main aspects taken into account are the spatial properties of physical objects that may change over time and the way in which these changes occur. Spatial properties of objects that typically change over time are their location, direction, size, and shape. Location and direction are used to capture the movement of objects; their evolution provides information on change that can be modeled as geometric transformations such as translation and rotation. The mutation of objects over time is apparent in their size and shape properties. The information about size can be used to capture variations in scale, that is, the expansion and contraction of the objects, whereas the information about shape is required to capture deformations. The most complex case for mutation arises when it involves the aggregation or the fragmentation of complex objects, such as objects that are composed of other elementary objects. In these cases, it is often necessary to handle functional relationships between the objects, in order to make it possible, for instance, to answer questions about the exchange of components between complex objects.

The values of spatial properties may change continuously or discontinuously over time. When changes are continuous, the values of a spatial property may be modeled as a function of time represented as a continuous line in a graph. (Figure 1).

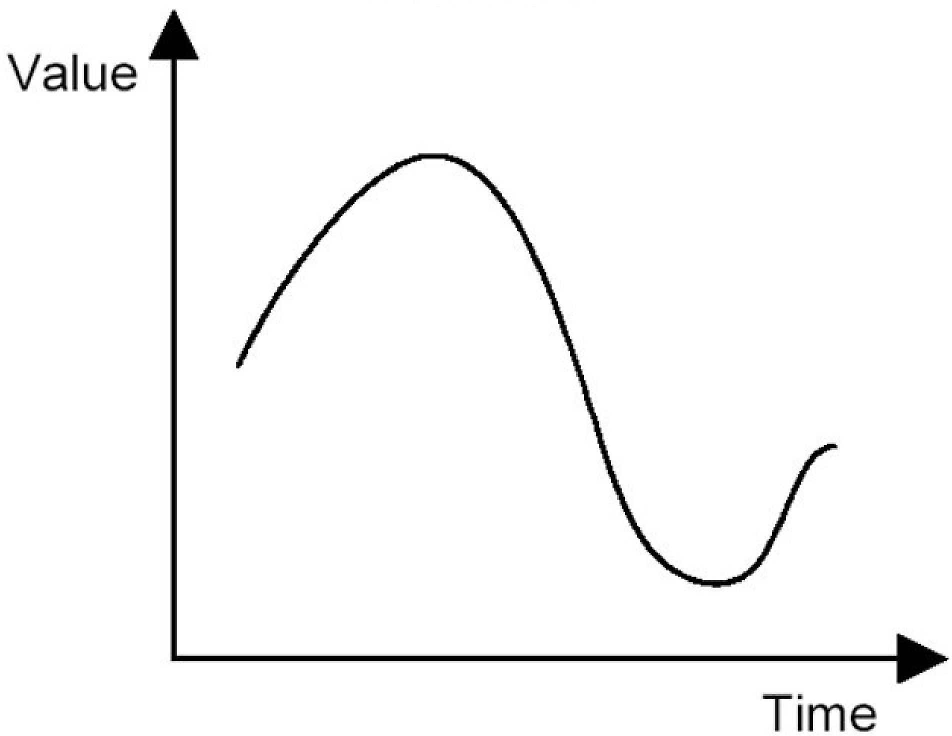

Figure 1. Continuous change

When the values of spatial properties change from one value to another at a precise moment in time we are in the presence of discontinuous change. In this case, the management of change can resort to the kind of solutions proposed for managing versions in temporal databases. A version is a snapshot of the state of the data at a specified time. The temporal dimension may represent time instants or time intervals. In the former case, a version records a value at a particular time instant. Between two time instants, the values of the objects' attributes are either considered undefined or assigned some conventional value (Figure 2).

Versions based on time intervals represent successive states of an object. In this case change can be viewed as caused by events and, between two events defining a time interval, the values for the objects attributes are considered constant (Figure 3).

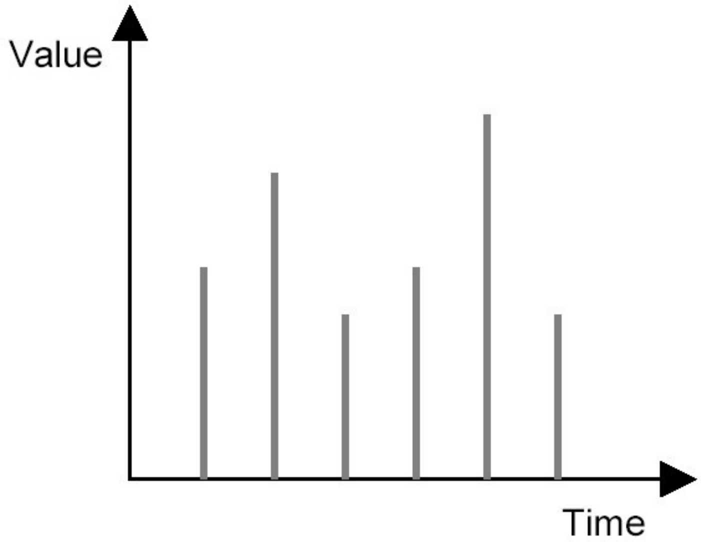

Figure 2. Event recording

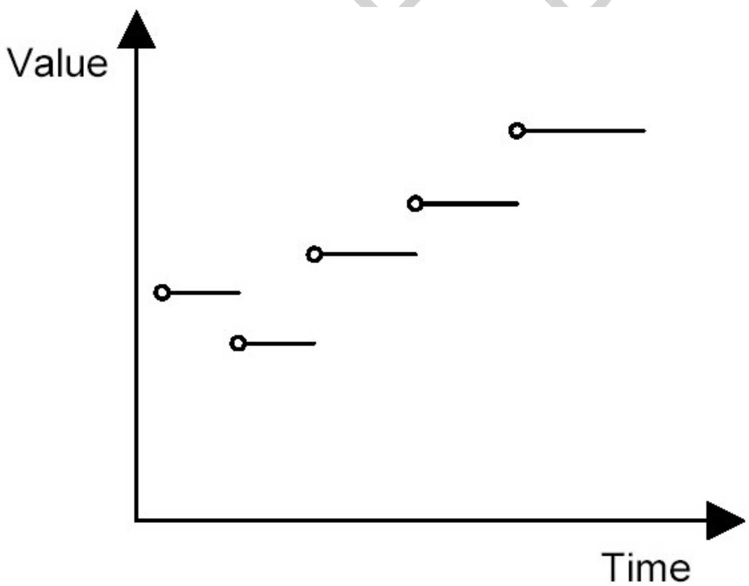

Figure 3. Change of state

2.2. Examples of Spatio-Temporal Systems

When classifying the allocation of spatio-temporal systems to categories, it is important to note that the complexity of their applications may differ significantly (see Table 1). For example, some systems only deal with changes in one kind of properties. This is the case with systems that handle rigid moving objects, where the only properties relevant are those related with the movement, that is, the location and the direction of the objects. The same is true for a system designed to monitor the deformation—changes in shape and size—of an island due to tides, where location and direction are not subject to change. Other systems must deal with changes in several kinds of properties. For example, an application designed to monitor the evolution of pollution spots in the sea should cope not only with the movement of the spots, but also with the changes on their extent—size and shape—and possibly with the fragmentation of one spot into several different spots.

A similar analysis can be carried out for systems where discontinuous change is assumed. For example, a shoreline will change over time because of coastal erosion. To track this (possibly continuous) change, the deformations of the curves that define the shoreline might be only periodically recorded. For an historical system studying the migration of populations it may be necessary to record not only the movement of populations but also the changes in the territory occupied at a given historical period, and even the possibility of aggregation and fragmentation of populations and territories.

	Continuous change			Discontinuous change	
	Moving Objects	Island	Pollution Spots	Shorelines	Migration of Populations
Movement	√		√		√
Mutation					
Simple deformation		√	√	√	√
Aggregation and fragmentation			√		√

Table 1. Classification of spatio-temporal applications

Different levels of detail may involve different recording requirements. For example, a fire monitoring system for a fire brigade operational center should be able to represent the number, location, and extents of burned areas, and to update this information, say, every three hours. In this case, the fire might be represented by polygons with discontinuous evolution over time.

However, fire brigades need more detailed information to control the fire. They need the location of the fire, as well as the local weather conditions, to isolate the focal points of the fire and prevent it from spreading into other areas. This information should also be updated more frequently than the higher level one.

It should be noted that the choice between continuous and discontinuous change might be determined either by the nature of the system and the data to be modeled, or by the practical possibility of capturing or acquiring data with sufficient rigor. These kinds of decisions belong to the modeling phase of the system.

3. Data Models and Languages for Spatio-Temporal Systems

A spatio-temporal database is a database designed for the representation, manipulation and querying of spatial data and spatial data changing over time. All individual spatial and temporal database concepts must be considered. However, as various studies have already shown, a simple combination of the spatial and temporal aspects of data is not appropriate to capture the whole diversity of the spatio-temporal phenomena that the users need to represent in a database system. Rather, the temporal evolution of spatial objects must be investigated in order to develop inherent spatio-temporal database concepts.

Early work on spatio-temporal systems started with the study of systems for concrete applications with their own specific problems, and was usually developed by domain experts. The main focus was on applications for the representation of discontinuously changing spatial data. Several data models and structures have been proposed, a few principles have been defined, and many problems, difficulties, or obstacles to the design and implementation have been encountered (for static modeling, see also *Conceptual Modeling of Geographic Applications)*.

Strongly encouraged by the *Chorochronos* network funded by the European Commission, the European spatial and temporal database communities started with systematic research on the integration of spatial and temporal concepts in database systems. The main focus is on systems for the representation of continuously changing spatial information. The representation of moving objects in database systems is becoming a major topic of research in this context.

It is important to note that temporal databases distinguish between valid time and transaction time. The *valid time* of a fact is the time when a fact is true in the modeled reality and is therefore part of the user data while the *transaction time* is the time when a fact has been stored in the database and may be retrieved. In what follows, all references are to the valid time concept.

3.1. Discontinuous Change

The representation of spatial objects is usually a combination of geometric and descriptive data. The geometric part, commonly referred to as *feature*, has often a complex structure, that is, it is composed of several primitive geometric components also referred to as elements. For example, shorelines are represented by sets of line segments. For such objects, two types of change are possible: attribute change and geometric change. In attribute change, the whole feature (for example, a shoreline) undergoes a change. With geometric change, the whole feature can be moved or deleted, a new element can be added to a feature, an existing element can be moved or deleted in a feature, and part of a feature can be reshaped or deleted. However, successive updates of geometry can cause edge effects and introduce artifacts that cause data to appear piecemeal or incorrect. If only a local area is altered, the borders of that area may no longer match the surrounding unaltered area. This problem has been studied at the US National Ocean Service and two approaches have been considered. One consists in leaving the dangling line segment to dangle and let the viewer, who sees the gap, to use imagination to close the gap. The other consists in adding a closing line and tag it to make it evident that it is an extra line.

Another issue concerns the techniques for timestamping spatial objects with complex geometry. The most intuitive and inexpensive approach is to timestamp the entire object associating a temporal value at its birth (its creation in the database) and another temporal value at its death (its deletion in the database). However, it is obvious that this solution has only limited expressiveness in what concerns the temporal properties of the object throughout its life. To obtain more control over the time granularity, it is necessary to timestamp the referenced primitive geometric components. Two solutions have been proposed. One consists of merging time and space at the primitive geometric component level (point, line segment, or polygon). The other consists of merging time and space at point level only, by forming a composite class of a spatial point (*point*) and a temporal point (*tpoint*). This class is used to construct other primitive spatio-temporal object classes, of higher dimensions, such as *tpolygon*.

Object identity becomes an issue in applications that must control the aggregation and fragmentation of spatial objects. The distinction between constructive change and non-constructive change has been proposed to deal with this problem. A non-constructive change takes place when a certain status of the object changes in terms of spatial or descriptive attributes. A constructive change is the change that creates a new object out of an existing one with a new concept different from the previous one. Non-constructive changes are simpler: it is sufficient to give a unique identifier to each object along the time line. The history of constructive changes is kept using an object-oriented data model based on a class named identity. An identity number is assigned to each version of an object, and a doubly linked list is maintained with pointers to the next and previous identity numbers.

Another approach, intended to support the development of applications for the management of discretely changing spatial data, is an extension of an object-oriented database system. Information is represented as spatio-temporal objects encapsulating descriptive, temporal, and spatial data. A set of operators for the evaluation of temporal, spatial, and spatio-temporal queries is also proposed. The framework provides a kernel of spatial and temporal class hierarchies to be used to build applications supported on the object-oriented concepts of inheritance, composition, and polymorphism.

-
-
-

Bibliography

Grumbach S., Rigaux P., Scholl M. and Segoufin L. (2000). The *Dedale* prototype. *The Constraint Databases*, pp. 365–382. Heidelberg, Germany: Springer Verlag. [This chapter presents a linear constraint database system for spatial and spatio-temporal applications.]

Güting R., Böhlen M., Erwig M., Jensen C., Lorentzos N., Schneider M. and Vazirgiannis M. (2000). A foundation for representing and querying moving objects. *ACM Transactions on Database Systems*, **25**(1), pp. 1–42: ACM Press. [This paper presents an abstract data model for representation of continuously changing spatial information.]

Güting R. and M. Schneider, Moving Objects Databases (2005). Morgan Kaufmann Publishers. [This book focuses on modelling, storage and querying of moving objects data.]

Koenen R. (2001). *Overview of the MPEG-4 Standard*. ISO/IEC JTC1/SC29/WG11. URL: http://www.cselt.it/mpeg/standards/mpeg-4/mpeg-4.htm. [This document gives an overview of the MPEG-4 standard, explaining which pieces of technology it includes and what sort of applications are supported by this technology.]

Koubarakis M., Pernici B., Schek, H.-J., Scholl M., Theodoulidis B., Tryfona N., Sellis T., Frank A., Grumbach, S., Güting, R., Jensen C., Lorentzos N., Manolopoulos Y. and Nardelli E., (Eds.) (2003), *Spatio-Temporal Databases: The ChoroChronos Approach.* Springer-Verlag, Lecture Notes in Computer Science 2520. [This book presents *Chorochronos*, a European research project on spatio-temporal databases, and presents the major achievements in this area.]

Kuper G., Libkin L. and Paredaens J. (2000). *Constraint Databases*, pp. 175–314. Springer Verlag. [This book presents a comprehensive survey of the field of constraint databases.]

Moreira J., Ribeiro C. and Saglio J.-M. (1999). Representation and manipulation of moving points: an extended data model for location estimation. Cartography and Geographic Information Science (CaGIS), Special Issue on Dealing with time, *Journal of the American Congress on Surveying and Mapping (ACSM)*, 26(2), pp. 109–123. [This paper presents a data model for physical representation of information about objects' movement in relational database systems.]

Rigaux P., Scholl M. and Voisard A. (2001). *Spatial Database Systems with Application to GIS*, 440 pp. San Francisco, CA: Morgan Kaufman. [This book is an in-depth survey of special techniques for dealing with spatial data, particularly in the field of Geographic Information Systems (GIS).]

Wolfson O., Sistla P., Xu B. and Zhou J. (1999). Tracking moving objects using database technology in Domino. *Lecture Notes in Computer Science*, **1649**, pp. 112–119. [This paper presents an overview of the Databases fOr MovINg Objects (DOMINO) project for representation of information about the location of moving objects in real time systems.]

Zaniolo C., Ceri S., Faloutsos C., Snodgrass R., Subrahmanian V. and Zicari R. (1997). *Advanced Database Systems*, pp. 97–160, 267–314. San Francisco, CA: Morgan Kaufman. [This book presents a survey of some of the most important research topics on database systems including, among others, spatial and temporal databases.]

Biographical Sketches

José Moreira received his Ph.D. in computer science from Ecole Nationale Supérieure des Télécommunications de Paris (France) and Faculdade de Engenharia da Universidade do Porto (Portugal) in 2001. He is currently Invited Assistant Professor at the Department of Electronics and Telecommunications of Universidade de Aveiro and researcher at IEETA, a non-profit institute associated with the same university. His background includes data structures and databases. His research interests cover spatio-temporal database systems and digital libraries.

Cristina Ribeiro (born in Espinho, 1958) received her Ph.D. in Informatics, Artificial Intelligence branch, from Universidade Nova de Lisboa, 1993. She is currently Auxiliary Professor at the Department of Electrical and Computer Engineering, Faculdade de Engenharia da Universidade do Porto (FEUP), lecturing on Data Structures and Information Retrieval. She has been a Researcher at INESC, a non-profit

R&D Institute associated with Porto University, since 1985. Her research interests include temporal reasoning, constraint programming, and information retrieval.

Jean-Marc Saglio received his first M.S. degree in engineering from ENST, Paris, in 1966, another in mathematical statistics from Alger University, in 1973, and a third in computer science in Paris 6 University in 1985. He has been senior research engineer and teacher in the Computer Science and Networking Department of ENST until 2006. His interests covered spatial and temporal databases in an object-relational perspective. For five years, between 1997 and 2002, he was involved in research projects dealing with GIS, spatial indexing and metadata management, which became his main concern during following years. He had a wide experience of many kinds of database benchmarking ranging from industrial to research implications. He was an expert in groups working for TPC (in Bull), X/Open (Bull) and ISO (AFNOR).

Michel Scholl received his Ph.D. in computer science from the University of California at Los Angeles in 1977, and his French Thèse d'Etat from the University of Grenoble in 1985. He was with INRIA, Rocquencourt, France for twelve years, where he headed the Verso database group, prior to joining the faculty of CNAM, Paris. Since 1989, he has been a full professor of Computer Science at CNAM. He manages the database group in the research laboratory CEDRIC of CNAM. His background includes computer-assisted instruction, packet-switched radio networks, data structures and databases. His current research interests include spatial databases and image databases. He has served as program committee member of many database conferences, including ACM SIGMOD, VLDB, EDBT and ICDE