

# CHOROCHRONOS: Research on Spatiotemporal Database Systems

Timos Sellis

*Inst. of Communication and Computer Systems and Department of Electrical and Comp. Engin.  
National Technical University of Athens, Zografou 15773, Athens, Greece  
timos@dbnet.ece.ntua.gr*

## Abstract

*Spatiotemporal database management systems can become an enabling technology for important applications such as Geographic Information Systems (GIS), environmental information systems, and multimedia. In this paper we address research issues in spatiotemporal databases, by providing an analysis of the challenges set, the problems encountered, as well as the proposed solutions and the envisioned research areas open to investigation in the EU funded project CHOROCHRONOS.*

## 1. Introduction

Temporal databases and spatial databases have long been separate, important areas of database research, and researchers in both areas have felt that there are important connections in the problems addressed by each area, and the techniques and tools utilized for their solution. Many papers in temporal databases conclude with phrases such as “the ideas in this paper may be extended to spatial databases.” Similarly, many papers in spatial databases suggest that techniques developed for spatial databases apply to temporal databases, by restricting attention to one dimension only. But so far relatively little systematic interaction and synergy among these two areas has occurred. CHOROCHRONOS aims to achieve exactly this kind of interaction and synergy, and aims also to address the many real-life problems that require *spatiotemporal concepts* that go beyond traditional research in spatial and temporal databases (support moving objects is a good example of the latter [8, 21, 22]).

*Spatiotemporal database management systems* (STDBMSs) can become an enabling technology for important applications such as Geographic Information Systems (GIS), environmental information systems, and multimedia. As a step towards the realization of this technology, CHOROCHRONOS was established as a *Training and Mobility Research Network* with the objective of studying the design, implementation, and application of STDBMSs. The participants of the network are the Institute of Computer and Communication Systems of the National Technical University of Athens, Aalborg University, FernUni-

versität Hagen, Università Degli Studi di L'Aquila, UMIST, Politecnico di Milano, INRIA, Aristotle University of Thessaloniki, Agricultural University of Athens, Technical University of Vienna, and ETH. All these are established research groups in spatial and temporal database systems, most of which have so far been working exclusively on spatial or temporal databases. CHOROCHRONOS enables them to collaborate closely and to integrate their findings in their respective areas. The network has the following main objectives.

- To stimulate research in the areas of spatial and temporal databases.
- To allow researchers working on spatial and temporal databases to improve their understanding of each other's work, to integrate their results, and to avoid duplication of work. The design and partial implementation of an STDBMS architecture will result.
- To allow researchers working on temporal and spatial databases to cooperate with researchers from *other disciplines* that are faced with spatial and temporal information and that would benefit from spatiotemporal database technology.

To achieve these objectives, CHOROCHRONOS pursues an extensive research program, covering issues related to the ontology, structure, and representation of space and time; data models and query languages for STDBMS; graphical user interfaces for spatiotemporal information; query processing algorithms, storage structures and indexing techniques STDBMSs; and architectures for STDBMSs. The participants also organize workshops with the participation of researchers from other disciplines faced with temporal and spatial information.

The network is coordinated by Timos Sellis (National Technical University of Athens); more information about the network and its activities can be found at <http://www.dbnet.ece.ntua.gr/~choros>.

## 2. Overview of Research in CHOROCHRONOS

Put briefly, a spatiotemporal database is a database that embodies spatial, temporal, and spatiotemporal database concepts, and captures simultaneously spatial and temporal aspects of data. All the individual spatial and temporal

concepts (e.g., rectangle or time interval) must be considered. However, attention focuses on the area of intersection of the two areas, which is challenging, as it represents inherently *spatiotemporal concepts* (e.g., velocity and acceleration). In spatiotemporal data management, the simple aggregation of space and time is inadequate. Simply connecting a spatial data model to a temporal data model will result in a temporal data model that may capture spatial data, or in a spatial data model that may capture time-referenced sequences of spatial data.

The main technical goal of CHOROCHRONOS is to study the issues involved in the design and implementation of an STDBMS. To achieve this goal, we carry out research covered by the following six tasks.

- Task 1: *Ontology, Structure, and Representation of Space and Time*. This task involves the study of temporal and spatial ontologies, including their interrelations and their utility in STDBMSs. In addition, structural and representational issues as they have been articulated in spatial and temporal database research are considered in order to obtain a common framework for spatiotemporal analysis.
- Task 2: *Models and Languages for STDBMSs*. The focus is here on three topics: (i) the study of languages for spatiotemporal relations, (ii) the development of models and query languages for spatiotemporal databases, and (iii) the provision of design techniques for spatiotemporal databases. This work builds on previous proposals and covers relational and object-oriented databases.
- Task 3: *Graphical User Interfaces for Spatiotemporal Information*. Research in this task has two goals: (i) the integration and extension of the participants' expertise in graphical interfaces for temporal and spatial databases, and (ii) the development of better visual interfaces for specific applications (e.g. VRML for time-evolving spaces).
- Task 4: *Query Processing in Spatiotemporal Databases*. Techniques for the efficient evaluation of queries are the focus of this task. These studies cover a variety of optimization techniques, ranging from algebraic transformations to efficient page/object management.
- Task 5: *Storage Structures and Indexing Techniques for Spatiotemporal Databases*. This task involves the comparison of existing storage and access structures for spatiotemporal data and is expected to lead to novel structures that enhance the most promising existing structures.
- Task 6: *The Architecture of an STDBMS*. The research activities outlined above are being carried out with a broader goal in mind: the development of an architecture for an STDBMS.

After this brief outline of the research areas pursued in CHOROCHRONOS, we proceed to give more detailed de-

scriptions of some of these areas. The emphasis of the material to follow is on ontological and modelling issues.

### 3. Spatiotemporal Data Modeling

In this section we address some issues involved in the ontology of spatial entities as well as the ontology of space itself, and issues corresponding to the development of conceptual and logical models, along with respective languages, for spatiotemporal data.

#### 3.1. Ontological Issues

Regarding the *ontology of spatial entities*, in order to model change in geographic space, a distinction is made between *life* (the appearance and disappearance, and merging and splitting of objects) and *motion* (the change of location over time) of objects. At its outset, this research identified and investigated prototypical situations in the life and movement of objects in geographic space. Particular focus has been on the life and movement of socio-economic spatial units.

A formal model of object lifestyles has been developed, in which the notion of *lifestyle* is applied to cover a variety of changes that an object or a group of objects can undergo through time [15]. The possible temporal constructs are formalized using the Haskell programming language. These constructs, when combined, have enough expressive power to represent very complicated worlds of change.

Regarding the *ontology of space*, a region-based ontology of geographic space was proposed. Spatial objects are located at regions in space. The concept of *exact location* is a relation between an object and the region of space it occupies. Spatial objects and spatial regions have a composite structure, i.e., are made up of parts. The ways in which parts of objects are located at parts of regions of space are captured by the notion of *part location*. Since there are multiple ways for parts of spatial objects to be located at parts of regions of space, multiple part location relations are identified, and a classification of part location relations is provided.

The work on ontology is coordinated by the group at the Technical University of Vienna. Additional work on spatiotemporal ontology is also carried out in Task 2 of CHOROCHRONOS where specifications of spatiotemporal objects (e.g., moving points and regions) and operations on these are defined.

#### 3.2. Models and Languages

Models and languages for spatiotemporal database systems is a central activity in CHOROCHRONOS, as it serves as the basis for several other tasks (for example

query processing and indexing). In this area, several aspects of spatiotemporal data modeling are covered: logical data models (Section 3.2.1), conceptual data models (Section 3.2.2), and ECA rules (Section 3.2.3). Special emphasis has also been put into issues relating to the representation of uncertain information (Section 3.2.4).

**3.2.1 New Data Models for Spatiotemporal Information.** This research may be divided into two categories: a) research initiated in the project that focuses on tightly integrated spatiotemporal support, and b) previously initiated efforts that have dealt mainly with temporal aspects, but are being extended to also cover spatial aspects. We consider in turn research in each category.

An important effort led mainly by the group at FernUniversität Hagen, focused on identifying the main requirements for a spatiotemporal data model and a spatiotemporal DBMS. Based on a *data type approach* to data modeling, the concept of *moving objects* has been studied [5, 6]. This has led to a research plan consisting of several steps leading from data model specifications (at two different, abstract levels) to implementation aspects. We have also investigated the basic properties of spatial, temporal, and spatiotemporal data types. Having specified mappings between different data type models and their relational DBMS embeddings, a precise landscape of the models' relative expressive powers has been drawn in [7]. Finally, a concrete spatiotemporal data model for moving objects has been provided. Here the focus is on a systematic classification of operations on relevant data types that facilitates a highly generic specification and explanation of all types and, in particular, of the operations on them. A detailed description of this model, including formal specifications of the types and operations is presented in [8] along with examples demonstrating the data model in action.

A significant effort in the area of models and languages deals with the use of constraint database models. These models constitute a separate direction that CHOROCHRONOS researchers are exploring in modeling spatiotemporal information. The Verso database group at INRIA has been working on DEDALE, a prototype of a constraint database system for spatiotemporal information [9, 12]. DEDALE is implemented on top of the O2 DBMS and features graphical querying. DEDALE offers a linear-constraint abstraction of geometric data, allowing the development of high-level, extensible query languages with a potential for optimization, while allowing the use of optimal computation techniques for spatial queries. It was also shown in [11] that the complexity of manipulating data of higher dimension depends not so much upon the global dimension, which as it is well known constitutes an exponential factor, but upon the orthographic dimension. In particular, in the context of spatiotemporal data (dimension 3) [10] developed techniques to represent it with

an orthographic dimension 2, thus leading to an efficient evaluation of queries.

Other research teams pursue research in the second category of models and languages, i.e. extending the relational ones. The team at Aalborg has continued previous research on temporal query languages. First, a framework consisting of temporal relations and algebraic operations on these has been provided, within which query language properties, e.g., the notion of snapshot equivalence, the reducibility of algebraic operators, and the notions of point-based and interval-based semantics have been studied formally [1]. This framework may be generalized to cover also spatial aspects. Second, the core of an SQL-based language, STSQL, has been proposed [18]. This language generalizes previous proposals by permitting relations to include multiple temporal as well as spatial attributes, and it generalizes temporal query language constructs, most notably statement modifiers, to apply to both the spatial and temporal attributes of relations. Because space and time are captured by separate attributes, the STSQL is intended for applications that do not involve storing the movement of continuously moving objects [3].

Another approach is pursued by UMIST. This group has developed the TAU system, which offers a formal, integrated framework for developing temporal object databases [13]. TAU is being extended to support also spatiotemporal data. TAU is fully upward compatible with the ODMG de-facto standard. It provides a rich set of temporal modeling constructs; temporal querying facilities; and a temporal modeling, analysis, and design methodology. TAU offers a set of tools that can extend any ODMG-compliant ODBMS with support for the concepts defined in TAU and provide temporal object query services in the syntax of TOQL. The main components of the TAU system are: the TAU Library that provides a set of built-in temporal literal types, the TODL compiler for specifying temporal database schemas in the syntax of TODL, and the TOQL compiler for supporting TOQL queries. TAU ensures design and source code portability across multiple platforms.

**3.2.2 Conceptual Modeling.** The work on spatial and temporal conceptual modeling extends the participants' previous work on temporal and spatial data modeling.. Spatial modeling aspects, e.g., the representation of objects' "position" in space, as well as temporal modeling aspects, e.g., the capture of the valid time of objects' properties, have been studied; and resulting new modeling constructs have been applied to existing conceptual models such as the ER model. Furthermore, the structure and behavior of so-called spatiotemporal phenomena (e.g., a "storm") have been investigated, and a formal framework with a small set of new modeling constructs for capturing these during conceptual design, has been defined [17, 19].

**3.2.3 ECA Rules.** Research has been initiated aiming at the modeling of *interactive spatiotemporal configurations* using ECA (Event-Condition-Action) rules [4, 16, 20]. The ECA rule model can represent efficiently all the features of novel application domains (VRML data, video data, etc.) such as interaction, constraints/conditions, and spatiotemporal actions to be taken [14]. The Event part of a rule represents a piece of interaction in terms of event(s) (simple or composite) that will trigger the actions included in the Action part of the rule, given that the constraints in the Condition part of the rule hold.

**3.2.4 Imprecision and Uncertainty in Spatiotemporal Information.** Work in this area has focused on the problem of representing *indefinite (or imprecise) temporal information*. In previous research, researchers from Politecnico di Milano have proposed the TSOS model for representing imprecise temporal information, emphasizing in particular temporal constraints on data, defined at schema level; applications have been proposed in the area of office information systems and traffic control. Subsequently, in cooperation with the University of Turin, the LaTeR approach has been proposed to specify imprecise temporal data (both quantitative and qualitative) and to reason about such data by using the STP constraint reasoning approach [2]. In recent work, the LaTeR temporal reasoner has been integrated with a commercial relational database, in which imprecise times are stored, and extensions of relational algebra operations are defined to formulate queries on the temporal relational database. The semantics and properties of the extended relational algebra with temporal operators have been studied [2]. The approaches of TSOS and LaTeR are being extended and applied in CHOROCHRONOS to the design of interactive multimedia scenarios [16].

Finally, modeling issues related to uncertain spatiotemporal data have been examined at the National Technical University of Athens. By adopting fuzzy set methodologies, a general spatial data model has been extended to incorporate the temporal dimension of geographic entities and their uncertainty. In addition, the basic data interpretation operations for handling the spatial dimension of geographic data have been extended to also support spatiotemporal reasoning and fuzzy reasoning.

## 4. Other CHOROCHRONOS Work

Substantial efforts within the project have also been devoted to the study of *storage structures and indexing*. In particular, (a) efficient extensions of spatial storage structures to support motion have been proposed, and (b) benchmarking issues have been studied.

In order to efficiently support the retrieval of data based on the spatiotemporal extents of the data, existing

multidimensional access methods need to be extended. Work has already initiated in this area. For example, approaches that extend R-trees and quadtrees have been devised, along with extensive experiments on a variety of synthetic data sets.

Work on benchmarking issues for spatiotemporal data first introduced basic specifications that a spatiotemporal index structure should satisfy, then evaluated existing proposals with respect to the specifications, and illustrated issues of interest involving object representation, query processing, and index maintenance. As a second step, a benchmarking environment that integrates access methods, data generation, query processing, and result analysis should be developed. The objective is to obtain a common platform for evaluating spatiotemporal data structures and operations that are connected to a data repository and a synthetic data set generator. Work on this is already under way as well as work on generating spatiotemporal data in a controlled way so that benchmarks can be run.

Work on *query processing and optimization* has focused thus far on (a) the development of efficient strategies for processing spatial, temporal, and inherently spatiotemporal operations, (b) the development of efficient cost models for query optimization purposes, and (c) the study of temporal and spatial constraint databases.

## 5. Conclusions

In this paper we presented briefly some of the research going on within CHOROCHRONOS, a European research project in spatiotemporal databases. Throughout the three years that have passed since the beginning of the project, significant progress has been achieved in several areas. These include the understanding of the requirements of spatiotemporal applications, data models, indexing structures, query evaluation, and architectures for STDBMSs.

Although CHOROCHRONOS has made significant progress, much work remains to be done before an STDBMS may become a reality. Until the end of our project, we intend to concentrate on efficient implementations of the proposed data models, more work on indexing and query optimization, and on experimentation with the alternative architectures proposed. We welcome other database researchers to join us in these exciting efforts!

## Acknowledgement

The research presented has been performed under the project "CHOROCHRONOS: A Research Network for Spatiotemporal Database Systems", contract number ERBFMRX-CT96-0056, which is sponsored by the European Commission under the Training and Mobility of Researchers Programme. The input from all participating institutions is gratefully acknowledged.

## 6. References

- [1] Böhlen M. H., R. Busatto, and C. S. Jensen, Point-Versus Interval-Based Temporal Data Models, *Proc. 14th Int'l. Conf. on Data Engineering*, Orlando, Florida, 1998.
- [2] Brusoni V., L. Console, B. Pernici, and P. Terenziani, LaTeR: An efficient, General Purpose Manager of Temporal Information, *IEEE Expert* (12) 4, 1997.
- [3] Böhlen M. H., C. S. Jensen, and B. Skjellaug, Spatio-Temporal Database Support for Legacy Applications, *Proc. 1998 ACM Symposium on Applied Computing*, Atlanta, Georgia, 1998.
- [4] Brisaboa, N. R., I. Mirbel, and B. Pernici, Constraints in Spatio-Temporal Databases: A Proposal of Classification, *Proc. 3rd Int'l. Workshop on Evaluation of Modeling Methods in Systems Analysis and Design*. Pisa, Italy, 1998.
- [5] Erwig, M., R.H. Güting, M. Schneider, and M. Vazirgiannis, Spatio-Temporal Data Types: An Approach to Modeling and Querying Moving Objects in Databases. FernUniversität Hagen, Informatik-Report 224, 1997, to appear in *GeoInformatica*. Also CHOROCHRONOS TR. CH-97-08.
- [6] Erwig, M., R.H. Güting, M. Schneider, and M. Vazirgiannis, Abstract and Discrete Modeling of Spatio-Temporal Data Types. *6th ACM Symp. on Geographic Information Systems*, 1998.
- [7] Erwig, M., M. Schneider, and R.H. Güting, Temporal Objects for Spatio-Temporal Data Models and a Comparison of Their Representations. *Intl. Workshop on New Database Technologies for Collaborative Work Support and Spatio-Temporal Data Management (NewDB)*, 1998.
- [8] Güting, R.H., M.H. Boehlen, M. Erwig, C.S. Jensen, N.A. Lorentzos, M. Schneider, and M. Vazirgiannis, A Foundation for Representing and Querying Moving Objects. FernUniversität Hagen, Informatik-Report 238, 1998, submitted for publication. Also CHOROCHRONOS TR. CH-98-03.
- [9] Grumbach, S., P. Rigaux, and L. Segoufin, The DEDALE System for Complex Spatial Queries, *Proc. of the Fifteenth ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, Seattle, Washington, 1998.
- [10] Grumbach, S., P. Rigaux, and L. Segoufin, Spatio-Temporal Data Handling with Constraints, *Proc. ACM Symp. on Geographic Information Systems*, Washington, DC, 1998.
- [11] Grumbach, S., P. Rigaux, and L. Segoufin, On the Orthographic Dimension of Constraint Databases, *Proc. Intern. Conf. on Database Theory (ICDT 99)*, Jerusalem, Israel, 1999.
- [12] Grumbach, S., P. Rigaux, M. Scholl, and L. Segoufin, Dedale: A spatial constraint database, *Proc. of the Workshop on Database Programming Languages*, Boulder, 1997 (Also in Proc. 13e Journées de Bases de Données Avancées)
- [13] Kakoudakis, I., and B. Theodoulidis, The TAU System Architecture and Design, TR-96-9, TimeLab, Dept. of Computation, UMIST, United Kingdom, 1996.
- [14] Kostalas I., M. Vazirgiannis, and T. Sellis, Spatiotemporal Specification and Verification for Multimedia Scenarios, First CHOROCHRONOS Intensive Workshop on Spatio-Temporal Database Systems, Austria, 1997.
- [15] Medak, D. Lifestyles - A new Paradigm in Spatio-Temporal Databases. TR, Department of Geoinformation, Technical University of Vienna, 1999.
- [16] Mirbel I, B. Pernici, T. Sellis, and M. Vazirgiannis, Integrity Constraints on Interactive Multimedia Scenarios. Submitted for publication, 1998.
- [17] Tryfona, N. and C. S. Jensen, Conceptual Data Modeling for Spatiotemporal Applications, CHOROCHRONOS TR CH-98-08.
- [18] Tsostras, V., C. S. Jensen, and R. T. Snodgrass, An Extensible Notation for Spatiotemporal Index Queries, *ACM SIGMOD Record*, 27(1), 1998.
- [19] Tryfona, N., Modeling Phenomena in Spatiotemporal Applications: Desiderata and Solutions, *Proc. 9th Int. Conf. on Database and Expert Systems Applications*, LNCS, 1998.
- [20] Vazirgiannis M., Y. Theodoridis, and T. Sellis, Spatio-Temporal Composition and Indexing for Large Multimedia Applications, *ACM Multimedia Systems*, 6(5), 1998.
- [21] O. Wolfson, Modeling and Querying Moving Objects, *Proc. 13th Int. Conf. on Data Engineering*, Birmingham, UK, 1997.
- [22] O. Wolfson, Moving Objects Databases: Issues and Solutions. *Proc. of the 10th Int. Conf. on Scientific and Statistical Database Management*, Capri, Italy, 1998.