CHOROCHRONOS
A RESEARCH NETWORK FOR SPATIOTEMPORAL DATABASE SYSTEMS

The CHOROCHRONOS Participants*

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 [Wol97, Wol98, GBE+98]).

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. Similar efforts that study spatiotemporal information have occurred under the auspices of the National Center for Geographic Information and Analysis in the United States.

The CHOROCHRONOS participants 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.

The remainder of this paper provides an overview of the research pursued within CHOROCHRONOS and presents results of the project. The following section gives a brief outline of the research activities. Section 3 concerns contributions to the area of spatiotemporal data modeling. Section 4 discusses storage structures, indexing, and query processing; and Section 5 covers spatiotemporal database architectures. Section 6 offers conclusions.

2 Overview of Research in CHOROCHRONOS

2.1 Introduction

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. Figure 1 summarizes this view.

* Andrew Frank, Stephane Grumbach, Ralf Hartmut Güting, Christian S. Jensen, Manolis Koubarakis, Nikos Lorentzos, Yannis Manolopoulos, Enrico Nardelli, Barbara Pernici, Hans-Jörg Schek, Michel Scholl, Timos Sellis, Babis Theodoulidis and Peter Widmayer. Contact author: Timos Sellis, Dept. of Electrical and Comp. Engin., National Tech. University of Athens, Zografou 15773, Athens, Greece, timos@dbnet.ece.ntua.gr

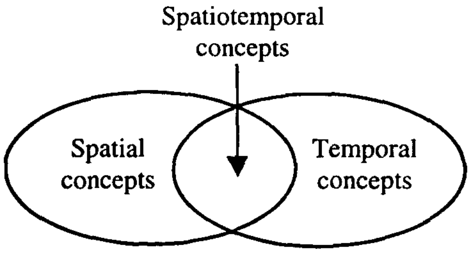

Figure 1

All the individual spatial and temporal concepts (e.g., rectangle or time interval) must be considered. However, attention focuses on the area of intersection, which is challenging, as it represents inherently *spatio-temporal 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.

Rather, the temporal characteristics of spatial objects (i.e., how entities evolve in space) must be investigated in order to produce inherently spatiotemporal concepts such as unified spatiotemporal data structures, spatiotemporal operators (e.g., approach, shrink), and spatiotemporal user-interfaces.

2.2 Research Objectives

The main technical goal of this network 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 descriptions of some of these areas.

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 [Me99b]. 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 [Me99a].

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 notion of *rough location* [Bi99] concerns the location of spatial objects at sets of regions of space that form regional partitions of space. Rough locations are characterized by sets of part location relations that relate parts of objects to parts of partition regions [Bi97]. A rough location can be considered an approximation of an exact location in terms of part location in a regional partition, i.e., as an approximation of the exact region of an spatial object in terms of this region's relations to a set of regions forming a regional partition.

Rough locations can be modeled formally by rough sets or by relationship mappings. Given the approximation view of rough location, operations on relationship mappings can be defined that can be used to approximate operations on regions of space. For example, pairs of minimal and maximal operations on relationship mappings approximate union and intersection operations on regions of space [BS99, Bi98].

The work on ontology is coordinated by the group at the Technical University of Vienna.

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 [EGSV97, EGSV98]. 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 [ESG98]. 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 [GBE+98] 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 [GRS98a, GRSS97]. 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 [GRS99] 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) [GRS98b] developed techniques to represent it with an orthographic dimension 2, thus leading to an efficient evaluation of queries.

Also in the context of constraint database models, researchers at UMIST have studied the role of spatial and temporal constraints in STDBMS [SK98b]. The efforts here concentrate on the development of a new spatiotemporal constraint-based database model. This model is based on the linear constraint database model of DEDALE and the indefinite temporal constraint database model (ITCDB) previously proposed in [SK98b].

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 [BBJ98]. This framework may be generalized to cover also spatial aspects. Second, the core of an SQL-based language, STSQL, has been proposed [TJS98]. 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 [BJS98].

Another approach is pursued by UMIST. This group has developed the TAU system, which offers a formal, integrated framework for developing temporal object databases [KT96a, KT96b, KT96c, KT96d, ST97a, ST97b, ST97c]. 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 [TH97a, TH97b, TJ98, Try98].

3.2.3 ECA rules

Research has been initiated aiming at the modeling of *interactive spatiotemporal configurations* using ECA (Event-Condition-Action) rules [VTS96, VTS98, MPSV98, BMP98]. 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 [KVS97]. 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 [BCPT97]. 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 [BCPT97]. The approaches of TSOS and LaTeR are being extended and applied in CHOROCHRONOS to the design of interactive multimedia scenarios [MPSV98].

Similar issues have been studied at UMIST for the ITCDB model. Calculus and algebra query languages have been defined for ITCDB, and the complexity of query answering has been studied in detail [Kou97a]. There is also considerable theoretical work on tractable classes of linear constraints, and experimental work on backtracking algorithms for disjunctive temporal constraints [Kou96, Kou97b, CJK97, SK98a].

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 Storage Structures, Indexing Techniques and Query Processing

Having given a brief overview of the data modeling efforts undertaken in CHOROCHRONOS, this section concentrates on efforts to develop techniques for the efficient implementation of the proposed data models and languages.

4.1 Storage Structures and Indexing

Substantial efforts have 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, (b) benchmarking issues have been studied, and (c) research on purely spatial or temporal structures has been continued.

Modern DBMSs should be able to efficiently support the retrieval of data based on the spatiotemporal extents of the data. To achieve this, existing multidimensional access methods need to be extended. Work has been done in this area as well; specifically, approaches that extend R-trees and quadtrees are reported in [TVS96] and [TVM98], respectively, along with extensive experiments on a variety of synthetic data sets.

Work on benchmarking issues for spatiotemporal data was initiated in [TSPM98], which introduced basic specifications that a spatiotemporal index structure should satisfy, 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 was proposed. A platform for evaluating spatiotemporal query processing strategies has been designed and implemented and has been already used for evaluating spatial join strategies [PRS99]. The "*A La Carte*" environment also provides benchmarking features for spatial join operations [GOPSS98].

As regards spatial structures, efficient spatial access methods based on hierarchical regular decomposition of the space for images containing multiple non overlapping [NP97] or overlapping [MNPP98] features were proposed. Bulk operations, in particular bulk loading but also bulk insertions and join operations, are inefficient if carried out by repeatedly calling operations for individual items; a generic bulk loading algorithm for spatial and also non-spatial indexes has been proposed in [BSW97]. This method is asymptotically optimal, since it achieves the runtime of external sorting.

Concerning temporal structures, overlapping B^+-trees were used for transaction-time indexing in [TML98]. An asymptotically optimal B-tree supporting transaction time was proposed in [BGO+96]. Taking the R-tree as its outset, the GR-tree [BJSS98a] permits the data regions it indexes and the bounding regions inside the tree to expand with the passing of time and indexes in this way efficiently general bitemporal data (i.e., data with both valid and transaction time support). An Informix DataBlade is available that implements the tree [BSSJ99]. The 4R-tree provides a well-performing, light-weight alternative to the GR-tree [BJSS98b].

A very natural way to operate on a GIS is by direct interaction via a walk-through user interface. We have proposed and implemented the virtual reality geographic information system ViRGIS [Paj98, POS+98] that holds the geographic data on a central database server and interactively loads and selectively displays the graphical information over the internet. In this context, speed limitations, as given by the network, required the use of multi-resolution techniques and new image compression methods [PW96]. The system can be used for many practical applications, e.g., visualizing a flight path of an airplane or for placing antennae on a terrain [BES+98].

With reference to distributed data structures, a solution has been defined for scaleable data structures for multi-dimensional points. Such a structure is able to dynamically enlarge with the insertion of new points distributing the load over the network and supports exact, partial, and range queries [NBP98]. In a dynamic, distributed scalable data structure, the coordination among clients that access and manipulate the data is a fundamental ingredient, without which not even correctness, let alone efficiency, can be achieved. We studied the basics of distributed coordination, pointed out the limits of what can be achieved, and proposed efficient distributed coordination mechanisms [Wat98, WW98a, WW98b].

Distributed networks in which spatiotemporal information is shared among many processors were studied with respect to fault tolerance. The concept of swapping a failing edge with a best possible replacement was introduced in [NPW98a]. This concept has since been studied for a number of important network features, such as the diameter, the radius, or the length of a shortest path [NPW98b]. The distribution of a linear quadtree representing raster data on a disk array towards reduced response times of user queries was addressed in [NPM97].

The support of multi-user access in a distributed system creates severe efficiency problems already with classical data structures. An approach to solve these problems by interleaving the operations was presented by [LSW97, SW97]. In the context of spatial data structures, particular theoretical problems related to space-filling curves [ARR+97], binary space partitions [BFG+96], and spatial join scheduling [NW97b] have been solved satisfactorily.

Two lines of research, the application-oriented discipline of geographic information systems and the technical discipline of geometric computation, in particular geometric algorithms and spatial data structures, will have significant practical impact in the near future. A state-of-the-art compendium on algorithmic foundations of geographic information systems was presented by [KNRW97], and a recent survey of spatial data structures can be found in [NW97a]. As a first result, a new basic data structure for GIS that is based on Voronoi diagrams for moving points [AGMR98] has been proposed in [GRR98].

4.2 Query Processing

Work on *query processing and optimization* has focused 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.

In [TPSS98] it was argued that expressing spatial operations, required by different application domains, is possible through a set of window searches, so that their execution could be supported by the available spatial indexing techniques. When the availability of index structures is not guaranteed, incremental algorithms have been proposed to support join operations for time-oriented data [PJ98]. Regarding the execution of inherently spatiotemporal operations, the basic classes of spatiotemporal operations required by different application domains involving the representation and reasoning on a dynamic world were defined [VTS98].

For query optimization purposes, analytical models that estimate the cost (in terms of node or actual page accesses) of join [TSS98], nearest-neighbor [PF98, PF99], and similarity [PM97a] queries involving R-tree indexed spatial data were introduced. Also, earlier analytical models for selection queries [TS96] were used as a platform to support direction relations (e.g., north, northeast) between two-dimensional objects [TPSS98], and the problem of *parallel and distributed similarity search* was studied in a shared-nothing multicomputer architecture, where an R-tree is declustered among the sites of the network [PM97b, PM98, PM96].

Finally, in the DEDALE system, we are developing query optimization techniques based on query rewriting that exploit the constraint representation, which allows to alternate between various modes of evaluation: lazy, symbolic, arithmetic, geometric, etc. The orthographic dimension, is used to rewrite queries in any dimension to a combination of queries in 2 dimensions.

5 Architectures

In database technology, many different system architectures have been proposed and are used in systems dealing with spatial and temporal data. Each has its particular strength and weaknesses. The most obvious system architectures are the ones proposed as being specially useful in spatial only and temporal only systems. Five possible architectures are presented next.

5.1 Standard Relational, Object Oriented or Object Relational DBMS with Additional Layer

A first possibility is to implement a spatiotemporal layer on top of a standard relational, object oriented or object relational DBMS. One can distinguish between two different approaches.

The 'thin layer' approach focuses on exploiting, as much as possible, the DBMS data model in order to handle spatial and temporal data. Concepts such as abstract data types are used whenever possible.

The 'thick layer' approach emphasizes the middleware aspects. The DBMS is primarily used as a persistent object store.

While query processing in the thin layer approach is done in the generic component using generic concepts, in the thick layer approach, tailored application specific query processing is done in the additional layer. This is similar to the support of application object services known from middleware layers. Both approaches make it possible to offer data models and query languages that are different from those supported by the underlying generic component.

5.2 Combination Architecture

Very similar to the layered architecture is the combined architecture that also uses a standard DBMS as its basis. In addition, other storage components (such as a file system) are used to store the spatial and temporal data and indexes. While this approach allows to *exploit different pre-existing components* possibly specialized for spatial or temporal data handling, it also introduces the *problem of coordinating different independent subsystems*. In particular, the query optimization task across different subsystems is non-trivial. As in the layered architecture, the combination architecture can distinguish between a thick-layer and a thin-layer approach.

5.3 Extensible DBMS

Instead of only allowing implementing 'on top' of the DBMS, extensible systems permit the *integration of new system components* such as data types, access methods, storage structures, and low-level query processing into the DBMS kernel. Extensible DBMS have been developed and implemented in research prototypes [HCL+90, SRH90, SPSW90], and they have become commercially available.

Figure 2 shows a general extensible DBMS architecture; the shaded boxes represent new components. One possible spatiotemporal system architecture is to use one of these systems and extend it with spatiotemporal components. The project participants at Hagen and ETH have own prototype implementations of extensible DBMS with special properties that make them particularly interesting in a spatiotemporal context.

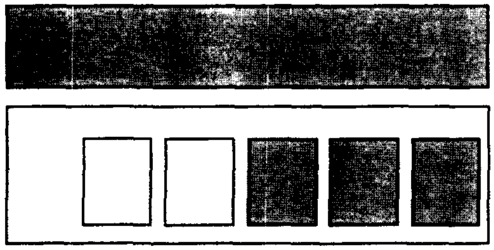

Figure 2: The General Extensible Architecture

5.4 The Secondo Architecture

Secondo [GFB+97, DG99] represents an approach to build extensible database systems based on a formalism for describing such systems called *second-order signature* (SOS) [Gut93]. In SOS, one can define a "logical" DBMS data model and a query algebra, a representation model with a query plan algebra ("physical algebra"), and optimization rules to translate from the former to the latter. As a result, it is possible to construct a generic DBMS system frame that can be filled with implementations of a wide range of data models. Another main idea is to structure the execution system as a collection of algebras. For example, a relation representation with access operations, an index structure with search operations, a package of join algorithms, or a set of data types for a special application would each be implemented as an "algebra module." Secondo strives to offer simple and well-structured interfaces for "users" wishing to implement such algebra packages. Hence, Secondo could be a good environment for implementing packages of spatiotemporal data types, or spatiotemporal access methods. The shaded boxes in Figure 3 represent such algebra packages.

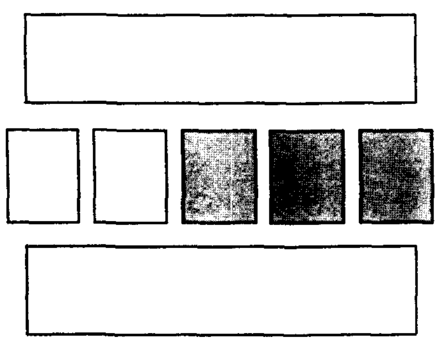

Figure 3: The Secondo Approach

5.5 The CONCERT Architecture

CONCERT [BRS96] focuses *on physical database design* aspects in the context of extensible systems. One of the main difficulties in implementing a new component for a standard extensible system is its integration with the built-in query processing engine. Instead of allowing arbitrary extensions, CONCERT provides a set of basic physical design possibilities with their corresponding query processing strategies. Extensions are built by defining relationships to the built-in physical design concepts.

A useful aspect of CONCERT in the context of spatiotemporal data is its ability to integrate non-database repositories such as legacy application systems or WWW servers into its query processing architecture. This makes it possible to exploit existing subsystems in a way similar to the combination architecture. On the other hand, the effort to be made coordinating the subsystems complicates the system design. Figure 4 gives an overview of CONCERT.

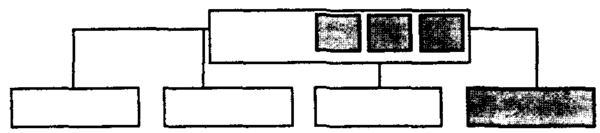

Figure 4: The CONCERT Approach

6 Conclusions

In this paper we presented CHOROCHRONOS, a European research project in spatiotemporal databases. Throughout the two and a half 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 is being sponsored in part by the European Commission funded TMR project, "CHOROCHRONOS: *A Research Network for Spatiotemporal Database Systems*", contract number ERBFMRX-CT96-0056.

7 References

[AGMR98] G. Albers, L. J. Guibas, J. S. B. Mitchell, and T. Roos, *Voronoi Diagrams of Moving Points*, Int. Journal of Computational Geometry and Applications, 8(3), 1998.

[ARR+97] T. Asano, D. Ranjan, T. Roos, E. Welzl, and P. Widmayer, *Space Filling Curves and Their Use In The Design of Geometric Data Structures*, Theoretical Computer Science, 181, 1997.

[BBJ98] 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.

[BCPT97] Brusoni V., L. Console, B. Pernici, and P. Terenziani, *LaTeR: An efficient, General Purpose Manager of Temporal Information*, IEEE Expert (12) 4, 1997.

[BES+98] M. Beck, S. Eidenbenz, C. Stamm, P. Stucki, P. Widmayer, *A Prototype System for Light Propagation in Terrains*, Invited Contribution to Computer Graphics Int'l., 1998.

[BFG+96] B. Becker, P.G. Franciosa, S. Gschwind, S. Leonardi, T. Ohler, P. Widmayer, *Enclosing a Set of Objects by Two Minimum Area Rectangles*, Journal of Algorithms, 21, 1996.

[BGO+96] B. Becker, S. Gschwind, T. Ohler, B. Seeger, P. Widmayer, *An Asymptotically Optimal Multiversion B-tree*, The VLDB Journal, 5(4), 1996.

[Bi97] Bittner, T. *A Qualitative Coordinate Language of Location of Figures within the Ground.* Proc. of COSIT'97, 1997.

[Bi98] Bittner, T. *Towards a Model Theory for Figure Ground Location.* Proc. 6th Symp. on Mathematics and AI, Fort Lauderdale, Florida, 1998.

[Bi99] Bittner, T. *Rough Location,* TR, Department of Geoinformation. Tech. Univ. of Vienna, 1999.

[BJS98] 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.

[BJSS98a] Bliujute, R, C. S. Jensen, S. Saltenis, and G. Slivinskas, *R-Tree-Based Indexing of Now-Relative Bitemporal Data,* Proc. 24th Int'l. Conf. on Very Large Databases, New York, 1998.

[BJSS98b] Bliujute, R, C. S. Jensen, S. Saltenis, and G. Slivinskas, *Leight-Weight Indexing of General Bitemporal Data,* TimeCenter TR-30, Aalborg University, 1998.

[BMP98] 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.

[BRS96] Blott, S., L. Relly, and H.-J. Schek, *An Open Abstract-Object Storage System*, Proc. ACM SIGMOD Conf. on Management of Data, 1996.

[BS99] Bittner, T. and J. Stell. *A Boundary-Sensitive Approach to Qualitative Location.* Annals of Mathematics and Artificial Intelligence. To appear.

[BSSJ99] Bliujute, R, S. Saltenis, G. Slivinskas, and C. S. Jensen, *Developing a DataBlade for New Index,* Proc. 15th IEEE Int'l. Conf. on Data Engineering, Sidney, Australia, 1999.

[BSW97] J. Van den Bercken, B. Seeger, P. Widmayer, *Generic Approach to Bulk Loading Multidimensional Index Structures*, Proc. 23rd Int'l. Conf. on Very Large Databases, Athens, Greece, 1997.

[CJK97] Cohen D., P. Jeavons and M. Koubarakis, *Tractable Disjunctive Constraints*, Proc. 3rd Int'l. Conf. on Principles and Practice of Constraint Programming, LNCS 1330, R. and A. Varzi. *The Structure of Spatial Localization.* Philosophical Studies 82(2), 1995.

[DG99] Dieker, S., and R.H. Güting, *Plug and Play with Query Algebras: SECONDO. A Generic DBMS Development Environment.* Fernuniversität Hagen, Informatik-Report 249, 1999.

[EGSV97] 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.

[EGSV98] 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.

[ESG98] 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.

[GBE+98] 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.

[GFB+97] Güting, R.H., C. Freundorfer, L. Becker, S. Dieker, and H. Schenk, *Secondo/QP: Implementation of a Generic Query Processor*, CHOROCHRONOS TR CH-97-09.

[GOPSS98] Günther, O., V. Oria, P. Picouet, J. M. Saglio, and M. Scholl, *Benchmarking Spatial Joins A La Carte*, Proc. of the 10th International Conference on SSDM, Capri, Italy, 1998.

[GRR98] C. Gold, P. Remmele, and T. Roos, *Fully Dynamic and Kinematic Voronoi Diagrams in GIS*, Algorithmica, 1998 (special issue on Cartography and GIS).

[GRS98a] 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.

[GRS98b] Grumbach, S., P. Rigaux, and L. Segoufin, *Spatio-Temporal Data Handling with Constraints*, Proc. ACM Symp. on Geographic Information Systems, Washington, DC, 1998.

[GRS99] 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.

[GRSS97] 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)

[Gut93] Güting, R.H., *Second-Order Signature: A Tool for Specifying Data Models, Query Processing, and Optimization*, Proc. ACM SIGMOD Conf. on Management of Data, 1993.

[HCL+90] Haas, L. M., W. Chang, G. M. Lohman, J. McPherson, P. F. Wilms, G. Lapis, B. Lindsay, H. Pirahesh, M. Carey, and E. Shekita, *Starburst mid-flight: As the dust clears*, IEEE TKDE, 2(1), 1990.

[KNRW97] M. van Kreveld, J. Nievergelt, T. Roos, and P. Widmayer (Eds.): *Algorithmic Foundations of Geographic Information*, Proc. Algorithmic Foundations of Geographic Information Systems, Udine, LNCS 1340, 1997.

[Kou96] Koubarakis M., *Tractable Disjunctions of Linear Constraints*, Proc. of the 2nd Int. Conf. on Principles and Practice of Constraint Programming, LNCS 1118, 1996.

[Kou97a] Koubarakis M., *The Complexity of Query Evaluation in Indefinite Temporal Constraint Databases*, Theor. Computer Science, 171, 1997.

[Kou97b] Koubarakis M., *From Local to Global Consistency in Temporal Constraint Networks*, Theoretical Computer Science, 173, 1997.

[KT96a] Kakoudakis, I., and B. Theodoulidis, *The TAU Temporal Object Model*, TR-96-4, TimeLab, Dept. of Computation, UMIST, United Kingdom, 1996.

[KT96b] Kakoudakis, I., and B. Theodoulidis, *The TAU Time Model*, TR-96-5, TimeLab, Dept. of Computation, UMIST, United Kingdom, 1996.

[KT96c] Kakoudakis, I., and B. Theodoulidis, *The TAU Temporal Object Definition Language*, TR-96-6, TimeLab, Dept. of Computation, UMIST, United Kingdom, 1996.

[KT96d] Kakoudakis, I., and B. Theodoulidis, *The TAU System Architecture and Design*, TR-96-9, TimeLab, Dept. of Computation, UMIST, United Kingdom, 1996.

[KVS97] 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.

[LSW97] Larsen K., E. Soisalon-Soininen, and P. Widmayer, *Relaxed Balance Through Standard Rotations*, Proc. 5th Int'l Workshop on Algorithms and Data Structures, LNCS 1272, 1997.

[Me99a] Medak, D. *Lifestyles - A new Paradigm in Spatio-Temporal Databases*. TR, Department of Geoinformation, Tech. Univ. of Vienna, 1999.

[Me99b] Medak, D. *Lifestyles - Life and Motion of Socio-Economic Units*. TR, Department of Geoinformation, Technical University of Vienna, 1999.

[MNPP98] Manolopoulos, Y., E. Nardelli, A. Papadopoulos, and G. Proietti, *MOF-Tree: a Spatial Data Structure to Manipulate Multiple Overlapping Features*, Information Systems, 1998.

[MPSV98] Mirbel I, B. Pernici, T. Sellis, and M. Vazirgiannis, *Integrity Constraints on Interactive Multimedia Scenarios*. Subm. for publication, 1998.

[NBP98] Nardelli, E., F. Barillari, and M. Pepe, *Distributed Searching of Multi-Dimensional Data: A Performance Evaluation Study*, Journal of Parallel and Distributed Computing (49), 1998.

[NP97] Nardelli, E. and G. Proietti, *Time and Space Efficient Secondary Memory Representation of Quadtrees*, Information Systems (22) 1, 1997.

[NPM97] Nanopoulos A., A. Papadopoulos, and Y. Manolopoulos. *On the Impact of Area Balancing of Declustered Linear Quadtrees*, Proc. 2nd Intern. Conf. for Geographical Information Systems in Urban, Environmental and Regional Planning, Athens, October 1997.

[NPW98a] Nardelli, E., G. Proietti, and P. Widmayer, *Finding All the Best Swaps of a Minimum Diameter Spanning Tree Under Transient Edge Failures*, Proc. ESA'98, 1998.

[NPW98b] Nardelli, E., G. Proietti, and P. Widmayer, *Finding a Detour-Critical Edge of a Shortest Path Between Two Edges*, Inf. Proc. Letters, 1998.

[NW97a] Nievergelt, J and P. Widmayer, *Spatial data structures - concepts and design choices*, In: Proc. Algorithmic Foundations of Geographic Information Systems, eds. van Kreveld et al, LNCS 1340, 1997.

[NW97b] G. Neyer, P. Widmayer, *Singularities Make Spatial Join Scheduling Hard, Algorithms and Computation*, LNCS 1350, 1997.

[Paj98] R. Pajarola, Access to Large Scale Terrain and Image Databases in Geoinformation Systems, Dissertation No. 12729, ETH Zurich, 1998.

[PF98] Proietti, G. and C. Faloutsos, *Selectivity Estimation of Spatial Queries for Line Segment Datasets*, Proc. 7th ACM Intern. Conf. on Information and Knowledge Management (CIKM'98), Washington, DC, Nov. 1998.

[PF99] Proietti, G. and C. Faloutsos, *I/O Complexity for Range Queries on Region Data Stored Using an R-tree*, Proc. 15th IEEE Intern. Conf. on Data Engineering (ICDE'99), Sidney, Australia, March 1999.

[PJ98] Pfoser, D. and C. S. Jensen,*Incremental Join of Time-Oriented Data*, TimeCenter TR-34, Aalborg University, 1998.

[PM96] Papadopoulos A. and Y. Manolopoulos, *Parallel Processing of Nearest Neighbor Queries in Declustered Spatial Data*, Proc. 4th ACM Workshop on Advances in Geographical Information Systems, 1996.

[PM97a] Papadopoulos A. and Y. Manolopoulos, *Performance of Nearest Neighbor Queries in R-trees*, Proc. 4th Int. Conf. on Database Theory Delphi, Greece, 1997.

[PM97b] Papadopoulos A. and Y. Manolopoulos, *Nearest Neighbor Queries in Shared-Nothing Environments*, Geoinformatica, 1(4), 1997.

[PM98] Papadopoulos A. and Y. Manolopoulos, *Similarity Query Processing using Disk Arrays,* Proc. ACM SIGMOD Conf., ACM SIGMOD Record, 27(2), 1998.

[POS+98] R. Pajarola, T. Ohler, P. Stucki, K. Szabo, P. Widmayer, *The Alps at your Finger Tips: Virtual Reality and Geographic Information Systems,* Proc. 14th Int'l. Conf. on Data Engineering, Orlando, Florida, 1998.

[PRS99] A. Papadopoulos, P. Rigaux, and M. Scholl, *A Performance Evaluation of Spatial Join Processing Strategies*, Proc. 6th Intern. Symp. on Spatial Databases (SSD 99), Hong Kong, 1999.

[PW96] Pajarola R., and P. Widmayer, *Spatial Indexing into Compressed Raster Images*, Proc. Int'l. Workshop on Multimedia Database Management Systems, 1996.

[SK98a] Stergiou, K. and M. Koubarakis, *Backtracking Algorithms for Disjunctions of Temporal Constraints*, Proc. AAAI-98, Madison, WI, 1998.

[SK98b] Skiadopoulos, S., and M. Koubarakis, *Tractable Query Answering in Indefinite Linear Constraint Databases*, CHOROCHRONOS TR CH-98-08.

[SPSW90] Schek, H.-J., H.-B. Paul, M.H. Scholl, and G. Weikum, *The DASDBS project: Objectives, Experiences, and Future Prospects*, IEEE TKDE, 2(1), 1990.

[SRH90] Stonebraker, M., L. Rowe, and M. Hirohama, *The Implementation of POSTGRES*, IEEE TKDE, 2(1), 1990.

[ST97a] Svinterikou, M. and B. Theodoulidis, *The Temporal Unified Modeling Language (TUML)*, TR-97-1, TimeLab, UMIST, UK, 1997.

[ST97b] Svinterikou, M. and B. Theodoulidis, *Mapping TUML to TAU TODL*, TR-97-2, TimeLab, UMIST, UK, 1997.

[ST97c] Svinterikou, M., and B. Theodoulidis, *The TUML process*, TR-97-3, TimeLab, UMIST, UK, 1997.

[SW97] Soisalon-Soininen E., and P. Widmayer, *Relaxed balancing in Search Trees*, Invited contribution to the Festschrift for the 60th birthday of Ron Book, Advances in Algorithms, Languages, and Complexity, 1997.

[TH97a] Tryfona, N., and Th. Hadzilacos, *Logical Data Modeling of SpatioTemporal Applications: Definitions and a Model*, Proc. Int'l. Database Engineering and Applications Symp.. Also CHOROCHRONOS TR CH-97-03.

[TH97b] Tryfona, N. and T. Hadzilacos, *Evaluation of Database Modeling Methods for Geographic Information Systems*, CHOROCHRONOS TR CH-97-05.

[TJ98] Tryfona, N. and C. S. Jensen, *Conceptual Data Modeling for Spatiotemporal Applications*, CHOROCHRONOS TR CH-98-08.

[TJS98] Tsotras, V., C. S. Jensen, and R. T. Snodgrass, *An Extensible Notation for Spatiotemporal Index Queries*, ACM SIGMOD Record, 27(1), 1998.

[TML98] Tzouramanis T., Y. Manolopoulos and N. Lorentzos, *Overlapping B^+trees - an Implementation of a Transaction Time Access Method*, Data and Knowledge Engineering (in press).

[TPSS98] Theodoridis Y., D. Papadias, E. Stefanakis, and T. Sellis, *Direction Relations and Two-Dimensional Range Queries: Optimisation Techniques*, Data and Kn. Engineering, 27(3), 1998.

[Try98] Tryfona, N., *Modeling Phenomena in Spatiotemporal Applications: Desiderata and Solutions*, Proc. 9th Int. Conf. on Database and Expert Systems Applications, LNCS, 1998.

[TS96] Theodoridis Y. and T. Sellis, *A Model for the Prediction of R-tree Performance*, Proc. 15th ACM Symposium on Principles of Database Systems (PODS'96), Montreal, June 1996.

[TSPM98] Theodoridis Y., T. Sellis, A. Papadopoulos, and Y. Manolopoulos, *Specifications for Efficient Indexing in Spatiotemporal Databases*, Proc. 10th Int. Conf. on Scientific and Statistical Database Management, Capri, Italy, 1998.

[TSS98] Theodoridis Y., E. Stefanakis, and T. Sellis, *Cost Models for Join Queries in Spatial Databases*, Proc. 14th IEEE Conf. on Data Engineering, Orlando, Florida, 1998.

[TVM98] Tzouramanis, T., M. Vassilakopoulos, and Y. Manolopoulos, *Overlapping Linear Quadtrees: a Spatiotemporal Access Method*, Proc. 6th ACM Int. Workshop on Geographical Information Systems, 1998.

[TVS96] Theodoridis Y., M. Vazirgiannis, and T. Sellis, *Spatio-Temporal Indexing for Large Multimedia Applications*, Proc. 3rd IEEE Conf. on Multimedia Computing and Systems, ICMCS'96, Hiroshima, Japan, 1996.

[VTS96] Vazirgiannis M., Y. Theodoridis, and T. Sellis, *Spatio-Temporal Composition in Multimedia Applications*, Proc. Int. Workshop on Multimedia Software Development, Berlin, Germany, 1996.

[VTS98] Vazirgiannis M., Y. Theodoridis, and T. Sellis, *Spatio-Temporal Composition and Indexing for Large Multimedia Applications*, ACM Multimedia Systems, 6(5), 1998.

[Wat98] R. Wattenhofer, *Distributed Counting: How to Bypass Bottlenecks*, Dissertation No. 12826, ETH Zurich, 1998

[Wol97] O. Wolfson, *Modeling and Querying Moving Objects*, Proc. 13th Int. Conf. on Data Engineering, Birmingham, UK, 1997.

[Wol98] O. Wolfson, *Moving Objects Databases: Issues and Solutions.* The Proc. of the 10th Int. Conf. on Scientific and Statistical Database Management, Capri, Italy, 1998.

[WW98a] R. Wattenhofer and P. Widmayer, *An Inherent Bottleneck in Distributed Counting*, Journal of Parallel and Distributed Computing, 49, 1998.

[WW98b] R. Wattenhofer and P. Widmayer, *The Counting Pyramid: An Adaptive Distributed Counting Scheme*, Proc. 5th Int. Colloquium on Structural Information & Communication Complexity, 1998.