# Dagstuhl Seminar
# *Integrating Spatial and Temporal Databases*

**22 – 27 November 1998**

**Schloss Dagstuhl-Wadern, Germany**
http://www.dag.uni-sb.de/eng

# Seminar Report

Contents

# Seminar Summary

Spatial databases incorporate the notion of space in order to accommodate the requirements for databases that allow reasoning about 2D and 3D such as geographical applications( GIS). Their study exists for more than twenty years. Lately, is triggered even more by the progress achieved in the power of computers which permits them to accommodate graphics and easily perform geometrical calculations. Spatial databases form an autonomous, active research community and a series of International Conferences are regularly organised (series of Symposium on Spatial Databases and Symposium on Spatial Data Handling). A number of Journals concern with spatial Databases as well(Cartographica, International Journal of Geographic Information Systems). The National Center for Geographic Information and Analysis (NCGIA, USA) is an established body coordinating research in Spatial Databases and their beneficial application in geography. OpenGIS is an International Consortium trying to bring Interoperability into Geographic Information Systems.

Temporal databases incorporate the concept of time to create high-level abstractions useful in database applications. This has been an active area of research for about twenty years. In the last few years the importance of the temporal database area has been recognised by the international scientific community. This recognition came in part in the form of the ARPA/NSF sponsored International Workshop on Temporal Database Infrastructure in 1993, a VLDB-affiliated temporal workshop in 1995, a special section of the IEEE Transactions on Knowledge and Data Engineering on temporal and real time databases published in August 1995, and the incorporation of temporal constructs, proposed by the temporal database community, in the soon-to-be standardised SQL3 language.

The main objective of the seminar was to bring together researchers from the two areas that have been working independently from each other and only recently have started to talk to each other. For example, research work on integration has started appearing on the main conferences and publications of each discipline.

One of the main issues discussed was whether it is feasible and if yes, how the research should be further integrated and if possible, what the mechanisms that the community can define so as to accelerate the process of developing a spatiotemporal infrastructure.

The “Integrating Spatial and Temporal Databases” seminar focused on establishing the foundations of a new discipline and also the future directions of that discipline, with respect to both research issues and the means to incorporate spatiotemporal databases into main-stream application development. A list of topics discussed at this seminar follows:

- Strategic discussions about the future of spatiotemporal databases as a discipline.
  Evaluation of the current state of the art with respect to the current trends in the DBMS tools and standards.

- Research Issues in Spatial and Temporal Databases: What is important?
- Spatiotemporal data models: relational, object-oriented, deductive and hybrid models. Where do the spatial and temporal capabilities fit in?
- Spatiotemporal user interfaces and languages. Update and retrieval languages for various types of temporal data models.
- Implementation issues in spatiotemporal databases. Issues that arise from experience of implementors and users and the agenda for research into these areas and transition to use in practice.
- Issue a "call for action" to the community (academia and vendors alike)

This seminar brought together over sixty researchers from fifteen countries that have dealt with different disciplines (spatial and temporal), as well as developers of databases and users, to conduct a fruitful discussion and evaluation of the activities thus far with a view on establishing the foundations of a new discipline that of spatiotemporal databases. There was a general agreement that there is still work to be done in Spatiotemporal Design, Data Models, Query Languages and Indexing while areas such as Temporal Data Models and Algebras are almost complete. Spatiotemporal Data Mining, Query Processing and Optimisation will produce significant results in the next ten years.

# Organisers

## Oliver Günther

Oliver Günther (Guenther) was born on October 22, 1961 in Stuttgart, Germany. He received his Diplom in Industrial Engineering from the University of Karlsruhe in 1984, and M.S. and Ph.D. degrees in Computer Science from the University of California at Berkeley in 1985 and 1987, respectively. Between 1988 and 1989, he was Postdoctoral Fellow at the International Computer Science Institute in Berkeley and Assistant Professor of Computer Science at U.C. Santa Barbara. From 1989 until 1993 he was Director of the Environmental Information Systems Division at FAW, a computer science research laboratory in Ulm, Germany. Since 1993, he has been Professor and Director of the Institute of Information Systems at Humboldt University in Berlin. He is also Chair of the Berlin-Brandenburg Graduate Program in Distributed Information Systems.

Professor Günther has conducted research projects in the areas of database management, knowledge-based systems, geographic and environmental information systems, as well as distributed information management. His list of publications includes nine books and more than 60 papers on related topics, and he was one of the founders of the SSD symposium series on spatial databases. He is Associate Editor of *GeoInformatica* and the *International Journal on Geographic Information Science*, and he has served on more than 30 conference program committees. Professor Günther held visiting faculty positions at the University of Cape Town, the University of California at Berkeley, and the Ecole Nationale Superieure des Telecommunications in Paris. He serves as a consultant to various government agencies and industrial companies, and he chairs the Supervisory Board (Aufsichtsrat) of POPTEL AG, an Internet telephony company.

## Timos Sellis

Timos Sellis received his B.Sc. degree in Electrical Engineering in 1982 from the National Technical University of Athens, Athens, Greece. In 1983 he received the M.Sc. degree from Harvard University and in 1986 the Ph.D. degree from the University of California at Berkeley, where he was a member of the INGRES group, both in Computer Science. In 1986, he joined the Department of Computer Science of the University of Maryland, College Park as an Assistant Professor, and became an Associate Professor in 1992. Between 1992 and 1996 he was an Associate Professor at the Computer Science Division of the National Technical University of Athens (NTUA), in Athens, Greece, where he is currently a Full Professor. Timos Sellis is also the head of the Knowledge and Database Systems Laboratory at NTUA.

His research interests include extended relational database systems, active database systems, and spatial, image and multimedia database systems. He has published over 100 articles in refereed journals and international conferences in the above areas.

Timos Sellis is a recipient of a Presidential Young Investigator (PYI) award for 1990-1995, and of the **VLDB 1997 10 Year Paper Award** for his paper "The R+-Tree: A Dynamic Index for Multidimensional Objects", by T. Sellis, N. Roussopoulos and C. Faloutsos (which appeared in VLDB 1987).

He is a member of the Editorial Boards of the International Journal on Intelligent Information Systems: Integrating Artificial Intelligence and Database Technologies, and Geoinformatica. Since 1996, he is coordinating the project CHOROCHRONOS: A Research Network for Spatiotemporal Database Systems.

## Babis Theodoulidis

Dr. Theodoulidis holds a Diploma in Computer Engineering and Informatics from the University of Patras, Greece, an M.Sc. in Computer Science from the University of Glasgow, United Kingdom and a Ph.D. in Computation from the University of Manchester Institute of Science and Technology.

He is currently a senior lecturer in the Department of Computation at UMIST where he has been working since January 1989.

He is Member of the British Computer Society, Member of the ACM, Member of IEEE Computer Society, Fellow Member of the Greek Computer Society and Fellow Member of the Engineering Council of Greece.

His research interests lie in the areas of Databases, Requirements Engineering, System Development Methodologies and Computer-Aided Software Engineering. He has extensively published in these areas and his research work has been supported by the Engineering and Physical Sciences Research Council (EPSRC), the ESPRIT programme of the European Union and by British and European industry.

Dr Babis Theodoulidis currently coordinates the activities of the TimeLab research laboratory specialising on intelligent information systems engineering.

# Abstracts of Talks

# What are spatioTemporal Databases

## CHOROCHRONOS: A Research Network for Spatio-temporal Database Systems (TMR Programme)

Timos Sellis

The main objective of CHOROCHRONOS is to allow European researchers working on spatial and temporal databases to achieve a higher understanding of each other's work, integrate their results and methodologies, and advance the state of the art in this area through an intensive three-year research program. This will culminate in the design and partial implementation of an architecture for Spatiotemporal Database Systems (STDBMS). The Participants also cooperate, through intensive workshops, with researchers from other disciplines who are dealing with temporal and spatial information in their research, and would benefit from the development of an STDBMS. CHOROCHRONOS stimulates training and mobility of young researchers working in the areas of spatial and temporal databases and actively pursues dissemination of results throughout European academic institutions and industry. The participants include: The National Technical University of Athens (NTUA) (Project Co-ordinator), Aalborg University, FernUniversität Hagen, Universita Degli Studi di L'Aquila, University of Manchester Institute of Science & Technology (UMIST), Politecnico di Milano, Institut National de Recherche en Informatique et en Automation (INRIA), Aristotle University of Thessaloniki, Agricultural University of Athens, Technical University of Vienna, and the Swiss Federal Institute of Technology, Zurich (ETHZ). The main technical goal of this network is to study the issues involved in the design and implementation of an STDBMS and to propose an STDBMS architecture. To achieve this goal, we propose to carry out the research covered by the following six tasks:

1) Ontology, Structure and Representation for Space and Time.
2) Models and Languages for STDBMS.
3) Graphical User Interfaces for Spatiotemporal Information.
4) Query Processing in Spatiotemporal Databases.
5) Storage Structures and Indexing Techniques for Spatiotemporal Databases.
6) The Architecture of an STDBMS.

# Requirements of Traffic Telematics to Spatial and Temporal Databases

Thomas Brinkhoff

Typical Services in the area of traffic telematics are:
- traffic information services,
- on-board and off-board navigation services,
- breakdown and emergency call services,
- information and booking services,
- fleet services.

Such services are based on techniques like GSM (for voice dialogues and data communication) and GPS (for determining a position). In the service centre, different types of geographical data are needed, especially, very detailed street maps containing information like road names, house numbers, points of interest, and traffic restrictions.

A very important task for implementing traffic telematic services, is to determine the location of a customer on the base of imprecise information and to compute routes considering the actual and the future traffic situation. Thus, a database system is required that supports spatial queries, different types of routing algorithms as well as the management of spatio-temporal objects like traffic jams.

Because of maintenance and operating reasons it is necessary to use one standardized database system instead of several special-purpose systems. Other important aspects are the performance, on-line updates, fault tolerance, and the support of a multi-user / multi-application environment.

# Ontology-Based Map Integration

Harry Uitermark

Map Integration is the process of establishing links between similar related representations of features in different spatial databases. The context of map integration is the reuse of update information, that is update propagation between two topographic databases with different spatial and temporal resolution. Crucial in this process is certainty about the equivalence of feature representations. Several kind of conflicts exist between different databases, for example different models, schema's, classes and data structures. With spatial information there are extra conflicts, for example different geometry's (polygons Vs polylines), different segmentation's (roads Vs road segments) and different aggregations (houses Vs building blocks). The emphasis in this research is on the semantics of spatial data. To define this semantics we propose a ontology-based framework for map integration. An ontology is in our definition a limited, structured set of unambiguously defined concepts. Ontologies exists on two levels. The first level is that of the discipline, in our case topographic mapping: the domain ontology. The second level is that of the applications or databases: the application ontologies. The relationships between the domain ontology and the application ontology are determined by the abstraction rules or capture criteria. These rules determine what features are selected, how they are represented, simplified and aggregated. Establishing these relationships between concepts in the real world and concepts in the database make it possible to find corresponding feature representations: if object instances from different databases refer to the same concept in the domain ontology they are semantically related, and if they overlap (that is: same position) they are most likely equivalent feature representations. The formulation of the abstraction rules makes it possible to check corresponding feature representations for consistency. This is part of our six step map integration strategy, which includes the synchronization of the databases as a first step. Our ontology-based approach for map integration creates interoperability between heterogeneous databases. At the moment we are implementing a prototype of a Map Integrator in Mathematica.

## References

[1] H. T. Uitermark, "The integration of geographic databases. Realising geodata interoperability through the hypermap metaphor and a mediator architecture", presented at Second Joint European Conference & Exhibition on Geographical Information (JEC-GI'96), (eds. M. Rumor, R. McMillan, and H. F. L. Ottens), Vol. I, pp. 92-95, Barcelona, Spain, IOS Press, March 27-29, 1996.

[2] H. T. Uitermark, "The integration of maps. How Mathematica? is used in the modelling of geo-objects", presented at Second International Mathematica Symposium IMS'97, Rovaniemi, Finland, June 29-July 4, 1997.

[3] F. A. van Wijngaarden, J. D. van Putten, P. J. M. van Oosterom and H. T. Uitermark, "Map Integration. Update propagation in a multi-source environment", presented at 5th ACM Workshop on Advances in Geographic Information Systems ACM-GIS'97, (ed. R. Laurini), Vol. I, pp. 71-76, Las Vegas, Nevada, USA, ACM, New York, November 13--14, 1997.

[4] H. T. Uitermark, P. J. M. van Oosterom, N. J. I. Mars and M. Molenaar, "Propagating updates: finding corresponding objects in a multi-source environment", presented at 8th International Symposium on Spatial Data Handling (SDH98), (eds. T. K. Poiker and N. Chrisman), pp. 580-591, Vancouver, Canada, International Geographical Union, July 11-15, 1998.

[5] H. T. Uitermark, A. B. M. Vogels and P. J. M. van Oosterom, "Semantic and geometric aspects of integrating road networks", to be presented at 2nd International Conference on Interoperating Geographic Information Systems (Interop'99), Zurich, Switzerland, 1999.

# Ontology of Space and Time

## Rough Location

Thomas Bittner

Spatial objects are located at regions of space. *Exact location* is a relation between an object and the region of space it occupies. Spatial objects and spatial regions have a compositional 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 how parts of spatial objects can be located at parts of regions of space, multiple part location relations are identified and a classification of part location relations is provided.

*Rough location* refers to location of spatial objects at sets of regions of space that form regional partitions of space. Rough location is characterized by sets of part location relations relating parts of objects to parts of partition regions. Rough location can be considered as an approximation of 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 location can be modeled formally by rough sets or relationship mappings. Given the approximation view of rough location then operations on relationship mappings can be defined that can be used to approximate operations on regions of space. It was shown that pairs of minimal and maximal operations on relationship mappings approximate union and intersection operations on regions of space.

The notion of rough location is important in the context representation of knowledge about spatial objects in physical reality. It was shown that by empirical means, i.e., from observation and measurement, only knowledge about rough location can be known. The observing or measuring process creates the regional partition in which rough location is observed or measured. Since only rough location can be observed and measured data representing those observations and measurements refer to rough location. Consequently operations performed on those data need to operations on rough location.

## Spatio-Temporal Reasoning About Identity and Visibility

Max Egenhofer

Current data models lack the power to express the semantics associated with basic changes to objects-changes that may affect their object identity such as creation, destruction, or being left intact. We have developed a language that describes chance based on object identity, i.e., that trait that uniquely distinguishes one object from another. Based on a small set of identity primitives

and through a systematic derivation of their combinations, we capture the semantics of change operations. These operations and their combinations are fundamental to the types of change commonly experienced by geographic phenomena and modeled by researchers who study spatio-temporal change. The operations reveal that the most basic forms of spatio-temporal change require few spatial considerations, such as those involved in the formation and dissolving of aggregates, and few temporal considerations, such as the ordering and coincidence of events. When analyzing a temporal sequence of a spatial configuration, we find scenarios that are based on similar properties: existence is replaced by visibility, and non-existence by invisibility. The combination of both types of spatio-temporal information provides new opportunities for spatio-temporal analysis and reasoning, capturing such interesting scenarios as hidden objects (existing, but invisible) or "visible but non-existing with history" (i.e., former objects that left traces). Such combination of existence and visibility form the basis for describing more complex scenarios of change and reasoning about spatio-temporal chance, and they serve as the basis for extending future spatio-temporal query languages

## Geospatial Lifelines

David Mark

A geospatial lifeline is here defined to be the continuous set of positions occupied by an object in geographic space over some time period. Geospatial lifeline data consist of discrete space-time observations of a geospatial lifeline, describing an individual's location in geographic space at regular or irregular temporal intervals. Geospatial lifelines appear to represent an important natural class of geospatiotemporal phenomena. After a brief discussion of a general typology of geospatiotemporal phenomena, the paper presented some examples of applications of the lifelines concept, including investigation of alibis of criminals, patterns of credit card and telephone use, emergency calls to the police, tracking of birds and mammals in the wild, and residential life histories of cancer patients. Some components of the technical research program include investigating cross-scale knowledge-discovery methods for geospatial lifelines that are captured at various levels of spatial and temporal detail, designing and prototyping computational models that can deal with large sets of geospatial lifelines, assessing the computational models by examining real-world applications, and developing methods for visualization of data and analytical results. Some methods of analysis will be based on HŠgerstrand's "Time Geography" model from the 1960s. The paper also mentioned the need for research on privacy implications of such systems. Technological advances, especially GPS, mean that the monitoring and recording of location and movement will intensify, providing unprecedented potention for surveillance and control of populations, as well as major societal benefits.

# Models and Query Languages

## A Spatio-Temporal Model for Integrated Information Management

Peggy Agouris

Current geographic information systems provide little or no support for modeling dynamic phenomena. More broadly, they lack the ability to effectively model change. Change occurs all the time but in varying increments such that it can be seen variously as gradual or abrupt. For many applications it is change that is of direct interest and yet in current GIS this information is not directly available. We propose a model that captures and stores elements of change and makes these available for direct query and analysis. Our project is funded by the US National Imagery and Mapping Agency (NIMA) and principal investigators are Kate Beard and Peggy Agouris. Anthony Stefanidis is cooperating investigator.

Our model works from the perspective that there are many possible representations of spatial phenomena stored implicitly in information resources and that the information on change lies in combinations of these representations. We arrive at information about change in spatial phenomena through multiple observations of phenomena over time. Analysis of these observations may reveal very different temporal behaviors from quite dynamic to essentially static. The types of change that may be observed include existence (phenomena appear and disappear), changes in shape, changes in location, changes in non-spatial characteristics and combinations of these. Furthermore patterns of change built from multiple observations over time can lead to estimates or predictions of unobserved change

We propose a spatial-temporal model that models the types of change listed above and makes these available for direct query and analysis through a spatial-temporal gazetteer. The content of the gazetteer is built and maintained from a library of spatial information resources called the multimedia information store which can include maps, imagery, video, and text as well as other possible media. The gazetteer represents only a small fraction of possible changes. The vast number of changes remain implicitly held in the multimedia information store but are available for extraction as the need or interest arises. Queries which invoke extraction operations against the multimedia store can be used to capture various types of change and post them to components of the gazetteer. Operations over the gazetteer can then be used to construct scenarios of change.

The advantage of this approach is that we extract and explicitly maintain only changes which are of direct interest. The information content of the gazetteer need not be uniformly developed for all phenomena. The variations in fact should reflect information request patterns of users showing greater depth for objects queried most frequently and intensely.
Our research approach first proceeds to formalize components of the model. The model forms the conceptual framework for specifying a set of operations between components of the model and users interactions with and across components of the model. These specifications are then to provide the basis for prototype development.

## Querying Temporal and Spatial Constraint Databases

Manolis Koubarakis

We review the scheme of indefinite constraint databases and the problem of query evaluation in this scheme. Because query evaluation is in the general case NP-hard, we try to discover tractable cases of query evaluation. We start with the assumption that we have a class of constraints C with satisfiability and variable elimination problems that can be solved in PTIME. Under this assumption, we show that there are several general classes of databases and queries for which query evaluation can be done with PTIME data complexity. We then search for tractable instances of C in the area of temporal and spatial constraints. Classes of constraints with tractable satisfiability problems can be easily found in the literature. The largest class that we consider is the class of Horn disjunctive linear constraints over the reals. Because variable elimination for Horn disjunctive linear constraints cannot be done in PTIME, we try to discover subclasses with tractable variable elimination problems. The class of UTVPI+Disequations constraints is the largest class that we show to have this property. Finally, we restate the initial general results with C ranging over the newly discovered tractable classes. Tractable query answering problems for indefinite temporal and spatial constraint databases are identified in this way. Two of them are significant extensions of problems studied previously by other researchers while all others are new.

Joint work with my student Spiros Skiadopoulos.

## Management of Nested Tables with Temporal Data

Nikos Lorentzos

The presentation concerns the definition of a temporal extension to QBEN, a query language for the management of data that does not satisfy First Normal Form. To this end, QBEN supports the relevant classical relational algebra operations (Union, Except, Nest, Unnest etc) and also two more operations, Unfold and Fold. Complicated queries are formulated step-wise, as a sequence of simple queries. In terms of execution, it is expected that these simple queries do not run sequentially; instead, an optimizer compiles them into one and produces optimal code. The characteristics of QBEN can be summarized as follows.

1. Management of non-temporal data: Some of the known algebraic operations had to be extended further.

2. Management of temporal data: QBEN also supports some special type of periodic temporal data, by supporting two valid time columns in the same table.

3. Management of interval data: Except time intervals, QBEN also supports all other possible interval data types, such as intervals of integers, of strings etc. Hence, it can be applied to a wide range of relevant applications.

## Conceptual Modeling of Spatio-temporal Data

Stefano Spaccapietra
with
C. Parent, Université de Lausanne, HEC S. Spaccapietra, EPFL Lausanne
E. Zimanyi, INFODOC, Université Libre de Bruxelles

The benefits of a conceptual modeling approach to database design have been extensively demonstrated in the domain of traditional databases. On the contrary, both practice and research regarding modeling of spatial/geographic or temporal data are mostly implementation oriented. In this presentation we addressed the issues of characterizing the features that a model should support to conceptually describe space and time related phenomena. This provides a check-list that allows to assess if a data model qualifies for use at the conceptual level. Next we developed the specification of a conceptual data model, named MADS, that we have specifically designed to meet the requirements we had identified. MADS is strongly rooted in the orthogonality principle, which commands that structural, spatial and temporal features may be independently defined. For instance, MADS allows space and time to be associated with objects, relationships or attributes, as appropriate, depending only on application requirements. By combining space and time, MADS also supports the description of moving objects. Continuos field may be described associated to semantic objects.

## On the Integration of Object-Oriented and Temporal Databases

Holger Riedel

Using the object-oriented approach for modelling and querying spatio-temporal databases seem compromising, because several concepts present in spatio-temporal databases are well-analyzed and and integrated in current OODBs.In this presentation we concentrate on the integration of temporal aspects into standard OODB models and query languages. We made the observation that the most important concept is orthogonality. This applies to data modelling. Additional concepts for temporal modelling have to go along with widely accepted standards for object modelling like the ODMG or Java. Moreover, orthogonality is important for temporal querying, because additional temporal constructors have to support the existing features and add further functionality. Also the implementation needs orthogonality using independence between the logical level and the physical level. Then additional physical concepts like new index structures or algorithms can easily be added and, on the other hand, extensions to the logical model (like support for spatial objects or imprecise querying) can be added without complications.

In our approach we formalized the ODMG model and OQL using the COCOON approach. Then valid-time and Temporal aspects are added by introducing a temporal domain used to define temporal type constructors for valid-time and transaction-time which can be used to build arbitrarily nested type structure with flexible support of temporal information. Moreover, the (valid-time or transaction-time) lifespan of an object in a class can be recorded. In contrast to other temporal object models proposed so far, we do not enforce any further temporal constraints on the type/class hierarchy, referential integrity and the relationship between the lifespan of an

object and its attribute values, because we made the observation that such constraints make it impossible to model many applications in an adequate way.
We extended OQL using the ideas of TSQL2. Moreover, temporal information can be accessed and manipulated similar to collection-valued attributes. Because the usual group-by clause of OQL allows to build arbitrary groupings, temporal aggregations can be described easily in this approach.

Another interesting aspect of our approach is the independence of the physical and logical level which allows either the transformation of temporal data to usual object-oriented storage techniques, including fragmentation, replication, and specialized indexes, either the extension to temporal access structures which can be added flexibly .

## Temporal Schema Versioning for OODB

Federica Mandreoli

The problem of supporting temporal schema versioning has been extensively studied in the context of the relational model. In the object-oriented environment, previous works were devoted to the study of the different aspects of schema evolution or (non-temporal) versioning in branching models, due to the traditional origination of the object-oriented model fromCAD/CAM and CIM.

Nowadays, the common adoption of the object-oriented model for a wide class of applications, extends temporal versioning requirements and expectations also to this model.

We propose a formal model for the management of temporal schema versioning in object-oriented databases. Its de_nition is partially based on the ODMG Release 2.0 Object Model and partially introduces new concepts. The proposed model supports all the schema changes which are usually considered in the OODB literature for which the full semantics and correctness proofs are provided. Semantic issues arising from the introduction of temporal schema versioning (like different notions of consistency and referential integrity) are also addressed on a formal basis. We are currently considering the integration of the branching approach in the temporal schema versioning framework. To this end, we are working on an extension of our model to also accommodate parallel versions.

## Representation and Manipulation of Moving Points on RDBMS：An Extended Data Model for Location Estimation

José Moreira

The fields of application of spatio-temporal systems, i.e., systems that must operate with time-varying spatial properties, are vast and heterogeneous. Since it would be difficult to treat such diversity as a whole, we introduce a classification for spatio-temporal systems based on the properties of the represented objects. Building on this classification, we also claim that features of some complex objects can be derived from those of simpler ones, suggesting an evolutionary approach, starting with the study of simple objects and progressing by enriching them with new features. The presentation focuses on the definition of a data model for representation of moving points. The model is based on the decomposition of the trajectory of moving points into sections. The movement within each section of a trajectory is described by a variability function. Since, for most systems, it is not possible to store the exact knowledge about the movement of a mobile, the answers to queries may be imprecise. We propose two additional approaches to deal with imprecision, the superset and the subset semantics, based on a maximum value for the variability function, and a smooth technique to integrate them in the model. Finally, we analyse some functional aspects of the implementation of the data model on relational and spatial database systems.

## Integrity Constraints for Interactive Multimedia Scenarios

Barbara Pernici

When authoring multimedia scenarios, and in particular scenarios with user interaction, where the sequence and time of occurrence of interactions is not predefined, it is difficult to guarantee the consistency of the resulting scenarios. As a consequence, the execution of the scenario may result in unexpected behavior or inconsistent use of media. The present paper proposes a methodology for checking the temporal integrity of Interactive Multimedia Document (IMD) scenarios at authoring time at various levels. The IMD flow is mainly defined by the events occurring during the IMD session. Integrity checking consists of a set of discrete steps, during which we transform the scenario into temporal constraint networks representing the constraints linking the different possible events in the scenario. Temporal constraint verification techniques are applied to verify the integrity of the scenario, deriving a minimal network, showing possible temporal relationships between events given a set of constraints. A set of query categories is also defined to examine the properties of a scenario according to the possible temporal relationships: such queries can allow the refinement of a scenario, the verification of given properties, and support scenario revisions.

## Review of SpatioTemporal Data Models

Achilleas Pavlopoulos

Currently, there are many efforts to integrate spatial and temporal database technology into spatio-temporal database systems. A number of new theory and concepts have emerged and a number of spatiotemporal data models have been proposed. This paper investigates the different types of spatio-temporal data models proposed in the literature. It provides an overview of previous achievements within the field and critically evaluates the different approaches through the use of a case study and the construction of a comparison framework.

## 3D Topological Consistency and the Time Factor

Rolf A. de By

In the first part of this two-tiered presentation, we discuss the typical characteristics of urban growth problems in developing countries, which leads to the identification of technical requirements for spatio-temporal data support for urban planners. In the technically minimalist approach we are taking such requirements are 3D representations of urban building infrastructure, and the successive stages that urban areas go through.

We then discuss a 3D vector data representation based on work by Pigot and Bresson, address issues of our prototype implementation, as well as extensions to support topological singularities. Our implementation is based on standard RDBMS technology, with a number of spatial query operators added. Future work will address valid time interval extensions as well as relevant indexing techniques.

In the second half of the talk, we discuss a number of elementary differences (amounting to worries) between thematic and spatio-temporal databases:
(1) spatio-temporal databases seem not to fully support (yet) the notion of data independence, crucial to so many appreciated functions of standard database systems;
(2) spatial data acquisition is a much more autonomous process than that of thematic data acquisition, and hence, data independence is much more difficult to achieve;
(3) consequently, a well-founded understanding of spatial database design is currently lacking;
(4) moreover, many language/model extensions are currently being proposed, but not always with a clear understanding of their (formal) semantics, which makes them particularly dangerous tools to apply correctly in application development.

# Representing and Querying Moving Objects in Databases

Ralf Hartmut Güting

Spatio-temporal databases deal with geometries changing over time. In general, geometries cannot only change in discrete steps, but continuously, and we are talking about moving objects. If only the position in space of an object is relevant, then "moving point" is a basic abstraction; if also the extent is of interest, then the "moving region" abstraction captures moving as well as growing or shrinking regions. We propose a new line of research where moving points and moving regions are viewed as three-dimensional (2D space + time) or higher-dimensional entities whose structure and behaviour is captured by modeling them as abstract data types. Such types can be integrated as base (attribute) data types into relational, object-oriented, or other DBMS data models; they can be implemented as data blades, cartridges, etc. for extensible DBMSs. In the talk we explain the approach, discuss the problem of selecting the appropriate abstraction level in modeling, and finally describe a specific design of types and operations for moving objects.

# Models and Query Interfaces

## Uncertainty for Spatial and Temporal Relationships

Michalis Vazirgiannis

Spatial and temporal information bear uncertainty, especially as regards the relations between spatial and temporal facts. It is recognized that the existing spatial relations do not cover the areas of metrics (i.e. to what degree are two objects A, B overlapping, or to what degree they meet), and the area of uncertainty (i.e. objects A, B overlap significantly, object A is far away from B etc.) [Alt94]. If such concepts where supported, in spatial models, a new era for spatial decision making and subsequently for query processing would arise. So far, spatial data modeling and spatial reasoning has been based on the hypothesis that the extent and the boundary of objects are determined by the Boolean decision of whether a point belongs to the objects or not. The objects in this case are defined by their boundary that encloses the interior. All points that do not belong to either of them are external to the object. Spatial relations are currently designed as binary predicates yielding a Boolean and thus strict decision whether a certain relationship holds for two spatial objects or not. Well-known examples are topological predicates like overlap, meet, equal, disjoint, and inside. There is a great amount of knowledge that resides in the spatial and temporal relations among objects in a relevant framework. We are interested in modeling this knowledge and reason about this taking in account the uncertainty related to the Spatio-temporal relations among objects. We are interested in representing the uncertainty of spatial information related to the position, and the shape of an object and also the uncertainty related to spatial relations features (topology, direction, metrics) Fuzzy set theory is an extension or generalization of classical Boolean set theory and aims at representing the degree to which an object is classified to a set. This paper aims at definition of a model to represent and handle the uncertainty inherent in the spatial and temporal aspects of related contexts. As for the temporal dimension we model the uncertainty of the proximity of a fact to a desired time point. As regards the spatial domain we define a model to represent and reason on the uncertainty inherent to the metrics of the various aspects of spatial relations, namely topology, direction and metrics. We apply this scheme in the context of collaborative work sessions where a user searches for sessions which, among other features, are characterized by temporal and spatial aspects.

# Modelling Scenarios in VRML Worlds using Nets

Isabelle Mirbel

Lately a new generation of application domains is emerging. Such applications, heavily dynamic and interactive, include interactive multimedia applications, virtual reality worlds, digital movies or 3D animations. Such applications deal with an increased complexity, due to the kind and the number of objects involved. They also encapsulate rich interaction modes, external and internal. And they also deal with intensive spatio-temporal dependencies between the participating objects with motion becoming a central issue. Indeed in the aforementioned application domain, much of the information conveyed and manipulated has spatial and/or temporal aspects. Synthetic worlds (like VRML worlds, digital movies, interactive multimedia scenarios etc.) can be started at any point in time, and there is a multitude of events that may occur (i.e. 3D object collisions) which may trigger other actions (i.e. change of the motion of the objects that collided) provided that some constraints hold. The session concept provides a new spatio-temporal context that has a different temporal origin and perhaps a different evolution according to the internal/external interaction that takes place. In this case the concept of scenario is an important entity representing the spatio-temporal course of the context (i.e. spatio-temporal actions possibly related) in conjunction with the occurring interaction (in terms of events) and potentially with the validity of conditions/constraints.

Therefore, we propose to model and represent Interactive Spatio-Temporal (IST) configurations in terms of active rules. Active rules describe reactions in response to particular events under given conditions. They provide an easy way to capture constraints during the design phase of an application, and they can also be useful to maintain the integrity of the database; they also can be a good means of implementing some of the constraints defined at the conceptual level. Active rules also represent an easy and homogeneous way to take into account all the kinds of integrity constraints. They can for example be helpful when managing objects in a scene to check some integrity constraints. They can also be helpful when used inside a multimedia application, to link an action the fact that an object is out of the screen, or too far away from another one.

# Data Mining

## Applying TEMPOS to Geographical Data Analysis

Marlon Dumas

We begin by presenting an application involving the analysis of the use of resources and space over time, in a ski resort located in the French Alps. This application aims at contributing to the studies on the reorganization and development of the resort's infrastructures.

The temporal aspects of this application are studied on the basis of the TEMPOS model, which is a temporal database framework integrating the main concepts and facilities required to manage the data historical dimensions on top of an object DBMS.
In its current stage, TEMPOS comprises:

1) A model for time and histories formalized by a hierarchy of ADT which allow to manipulate basic and complex temporal values (instants, durations, sets of instants) observed at multiple levels of granularities, as well as timestamped object properties.
2) Upward compatible extensions to ODMG schema definition and query languages, respectively ODL and OQL.
3) A language for describing patterns of histories based on regular expression operators with time constraints.

All three components have been implemented on top of the O2 DBMS and their functionalities have been tested on concrete applications such as the one presented in this talk.
In this talk we also introduce a taxonomy of algebraic operators on histories which is used to guide the presentation of the TEMPOS model, as well as the queries formulated in the context of the aforementioned application.

## Spatio-Temporal Data Mining

John Roddick

This talk outlines the issues and the current state of temporal and spatio-temporal data mining and knowledge discovery. Important issues, particularly pertaining to temporal and spatio-temporal data mining, will be illustrated including:

- Mining methods and Architectures;
- Complexity
- Interestingness criteria
- Inferencing ability

Current research topics are discussed and the results of experiments reported. Of particular interest is the practice of mining from the results of previous data mining operations - meta-mining. Results from experiments are reported.

# Performing Sequence Analysis with Data Mining Techniques

Myra Spiliopoulou

Modern organizations record all their business transactions and any other potentially useful information into repositories, the size of which increases at a tremendous pace. Many of those data are inherently temporal, such as stock trades, evolutionary trends of populations, lifelines of project contracts. These enormous archives will never be inspected directly by human eye. For them, we need techniques that extract and aggregate information to form interesting patterns. The notion of "interestingness" [ST96] depends on the application. In some applications, important patterns have been recognized and can be described formally, such as the "panic reversal" pattern in stock trade. For those applications, techniques have already been developed to detect patterns in time series, mostly by pattern matching (see e.g. [BC96, FRM94]). In other applications, interesting patterns cannot be specified per se. This is due to two reasons First, the events across the time sequence are not comparable in terms of a quantifiable property, i.e. they do not form trends. This holds for a sequence of telecommunication signals or of web page accesses by a user. Here, we need to aggregate sequences and study their joint statistical properties such as frequency of occurrence. The other reason is the lack of knowledge about expected patterns For example, what are the behavioural patterns of bank customers who will later close their accounts? For those applications, the first step of sequence analysis should focus on the extraction of a preliminary set of possibly interesting patterns. The data analysis tool employed for this purpose can exploit (i) the statistical properties of the data, (ii) the intuition and background knowledge of the expert and, occasionally, (iii) phenomena observed for specific data instances. For this type of applications we propose WUM, our web utilization miner, initially designed to assist in pattern discovery in web usage analysis, but based on a theory widely useful for mining sequential patterns. WUM extracts and aggregates sequences into a match for a "template". Differently from conventional templates, WUM templates are comprised of named variables subject to statistical constraints, and of unnamed variables optionally subject to structural constraints [Spi98].For the specification of those constraints, a declarative mining language, MINT, has been developed, adhering to the style of SQL and DMQL [SF98]. In SQL-jargon, MINT supports the "aggregation" and "grouping" of sequences; the statistical constraints of the template variables correspond to predicates in a "having"-clause. Thus equipped, WUM can be used for the prediction of events with given statistical confidence and for the identification of patterns occurring in a statistically significant number of sequences. Our work is only one step in the establishment of a supportive infrastructure for complex pattern discovery problems. Much activity is needed at several levels, including:

the combination of theoretical foundations and supportive tools from the areas of trend discovery in time series and sequence discovery;
index structures appropriate formining queries, leaning on the existing structures used in conventional queries over spatiotemporal data;
extensions of the temporal query languages with "mining" semantics; and, not least, methods for formalizing the notion of importance for patterns and metrics for measuring it.

# Applications

## The Tiger Temporal DBMS Prototype

Michael Bohlen

Tiger is an interval-based bitemporal database system prototype. Tiger demos ATSQL, a temporal extension of SQL-92. ATSQL uses statement modifiers to differentiate upwards compatible (UC), temporally upwards compatible (TUC), sequenced (SEQ), and nonsequenced (NSEQ) statements. UC ensures that legacy statements on legacy databases keep being valid with the exact same semantics. TUC allows to migrate data structures and applications independently. Essentially, legacy applications will not be aware that (a part of) the database has been migrated to become temporal. Advanced built-in and user-controlled temporal functionality is offered through SEQ and NSEQ statements. SEQ statements enforce a semantics that is consistent with viewing a database as a sequence of nontemporal databases. NSEQ statements, on the other hand, give the user full control over timestamps through built-in functions and predicates. Tiger is online accessible over the WWW via URL http://www.cs.auc.dk/~tigeradm/. An applet provides a shell-like interface that is suitable for distance learning. The main concepts of ATSQL are presented in a collection of books. Books can be browsed, modified, created, and deleted. Commands are evaluated on the remote server, which uses an Oracle database for data storage, manipulation, and retrieval.

## Geologic Hypermaps are more than Clickable Maps!

Marwan Abu-Khalil and Agnes Voisard

Geologic maps are {\it interpretations} of 3-D phenomena. Geologic hypermaps handle objects such as explanations, legends, geologic profiles, photos, videos and base data as well as various types of links among these objects. The major aspects to be modeled in such maps are uncertainty, fuzziness and complex relationships among the underlying data. In addition, map making is an incremental process which asks formultidimensional versioning on geospatial components, time and assumptions. In this talk, we first identify the requirements to represent and manipulate geologic hypermaps efficiently and we define a framework to support designers in the map making process. We then present a model to describe the structure of such maps. This model allows the explicit formulation of theories and assumptions leading to a particular map version.

# Query processing

## Point-based Spatio-Temporal Indexing

Yannis Theodoridis

An efficient benchmarking environment for spatiotemporal access methods should include a wide set of synthetic and real datasets for extensive experimentation purposes. The first part of my talk presents and evaluates three temporal extensions of the R-tree, the 3D R-tree, the 2+3 R-tree and the HR-tree, which are capable of indexing spatiotemporal data. Our experiments have shown that the while the HR-tree was the larger structure, its query processing cost was over 50% smaller than the ones yielded by the 3D R-tree and the 2+3 R-tree. As for generating data, several algorithms have been implemented in the past to generate static spatial (point or rectangular) data, for instance, following a predefined distribution in the workspace. However, by introducing motion, and thus temporal evolution in spatial object definition, generating synthetic data tends to be a complex problem. In the second part of my talk, I discuss the parameters to be considered by a generator for such type of data, present an algorithm for spatiotemporal data generation following an extended set of distributions, and visualize some of the results also giving hints for possible applications.

## Processing and Optimisation of Multi-way Spatial Joins Using R-trees

Yannis Theodoridis

One of the most important types of query processing in spatial database management systems (SDBMS) and geographic information systems (GIS) is the spatial join, an operation that selects object pairs from two relations that satisfy some spatial predicate. A multi-way join combines data originated from more than two relations. Although several techniques have been proposed for pair-wise spatial joins, currently there does not exist a method for multi-way spatial join processing by utilising existing indices on the relations to be joined. In this talk we present the close correspondence between multi-way joins and constraint satisfaction problems (CSPs) to solve multi-way spatial joins by applying systematic search algorithms that exploit R-trees to efficiently guide search. In addition to general methodologies, we present cost models and an optimization algorithm, and evaluate them through extensive experimentation.

# Properties of Poset Based Representations for Spatial Data

Enrico Nardelli

Formal methods based on the mathematical theory of partially ordered sets (i.e., posets) have been used for the description of topological relations among spatial objects since many years.
In particular, the use of the lattice completion (or normal completion) of a poset modelling a spatial subdivision has been shown by Kainz, Egenhofer and Greasley to be a fundamental technique to build meaningful representations for topological relations.
In fact, they proved that the new elements introduced by the normal completion process can (and have to) be interpreted as being the intersection of spatial objects. This is fundamental, from a mathematical point of view, since it means that the lattice resulting from the normal completion is the closure of the given set of spatial objects with respect to the intersection operation.

This result, however, leaves it open the question of the closure of the set of spatial objects with respect to the other fundamental operator to manipulate spatial subdivisions, namely the union operator.

In this talk we first precisely clarify the limitations for the use of lattices as models for spatial subdivisions.

Then we show that a technique already known in lattice theory, namely the construction of the maximal antichain lattice of a given poset, can be used to define another completion operator that builds the closure of the given set of spatial objects with respect to the union operation.

We also show that this new completion operator commutes with the normal completion and the lattice obtained from the application of both completion operators is minimal and unique up to isomorphism.

Finally we show how to apply the introduced operations when working on a subset (i.e. a view) of the spatial subdivision so that the computation only consider objects relevant to the view itself.

Our result gives further theoretical motivations to the use of lattices built on simplicial complexes as a model for spatial regions and relations, since this kind of lattices are, by construction, closed with respect to both the union and the intersection operations.

## Experience Building a Rose-Algebra Based Spatial Dood

Norman Paton

This talk presented experience in the design and implementation of a spatial extension [2] to the ROCK & ROLL [1] deductive object-oriented database system. The resulting system provides:

1. A rich object data model, with the vector spatial types of the ROSE Algebra [3] supported as spatial literals.
2. A deductive query and rule language for deriving information from the facts stored in the object model.
3. An imperative data manipulation language for creating and manipulating object model constructs.
4. The resulting system is available on the WWW at
5. http://www.cee.hw.ac.uk/Databases/rnr.html

References

[1] Barja, M.L., Fernandes, A.A.A., Paton, N.W., Williams, M.H., Dinn, A., and Abdelmoty, A.I., Design and Implementation of ROCK & ROLL: A Deductive Object-Oriented Database System, Information Systems, Vol 20, No 3, 185-211, 1995.

[2] Fernandes, A.A.A., Dinn, A., Paton, N.W., Williams, M.H. and Liew, O.,Extending a Deductive Object-Oriented Database System with Spatial Data Handling Facilities, to be published in Information and Software Technology, 1999.

[3] R.H. Guting and M. Schneider, Realm-Based Spatial Data Types: The ROSE Algebra, VLDB J., 4(2), 243-286,1995.

## I/O complexity for range queries on region data stored using an R-tree

Guido Proietti

In this talk we analyze the node distribution of an R-tree storing region data, like for instance islands, lakes or human-inhabited areas.We show that real region datasets are packed in minimum bounding rectangles (MBRs) whose area distribution follows the same power law, named {\em REGAL} (REGion Area Law), as that for the regions themselves. Moreover, these MBRs are packed in their turn into MBRs following the same law, and so on iteratively, up to the root of the R-tree. Based on this observation, we are able to accurately estimate the search effort for range queries, the most prominent spatial operation, using a small number of easy-to-retrieve parameters.

Experiments on a variety of real datasets (islands, lakes, human-inhabited areas) show that our estimation is extremely accurate, enjoying a maximum geometric average relative error within 30%.

## Similarity Search in Spatial Databases

Thomas Seidl

For modern database systems that manage complex objects, similarity search is an important task. Examples are the support of reducing the number of parts in CAD repositories, the docking problem in biomolecular databases, recall of similar cases in medical X-ray records, exploration of pictorial and image archives, and mining multimedia databases in general. In these applications, the objects of interest can hardly be specified by exact parameters such as unique key values, characteristic keywords, or selective attributes. Instead, the users provide examples, raw sketches, or other weak descriptions of what they want to obtain from the database. The system is expected to report all objects that as much as possible fulfill that description. Several similarity models for complex objects have been developed for a variety of spatial and spatiotemporal database applications. Examples include feature transforms, geometric approximations, 2D section coding, and methods for total matching as well as subsequence matching in time-series databases. Our new concept of shape histograms provides an intuitive and quite general approach for shape similarity search. The object space is partitioned into disjoint bins for each of which the fraction of space is measured that is occupied by a given object. Thus, a shape histogram is a discrete approximation of a spatial object. Since histograms are high-dimensional vectors whose dimension is equal to the number of bins, they are a typical example of a geometric feature transform approach. Even the concept of shape histograms was primarily developed for spatial objects, it is extended to spatiotemporal objects easily by considering time as an additional dimension. This works well for time series and for versioned objects but also for moving objects for which time proceeds continuously. Similarity search then may address the behavior of objects in addition to their spatial structure.

## Processing of Spatio-Temporal Queries in Image Databases

Michael Vassilakopoulos

Overlapping Linear Quadtrees is a structure suitable for storing consecutive raste according to transaction time (a database of evolving images). This structure saves considerable space without sacrifcing time performance in accessing every single image. Moreover it can be used for answering efficiently window queries for a number of consecutive images (spatio-temporal queries). In this report, we present five such temporal window queries (strict containment, border intersect, general intersect, cover and fuzzy cover). Besides, based on two methods of producing synthetic pairs of evolving images (random trees and random images with specified aggregation) we present empirical results on the I/O performance of these queries)

# List of Participants

**Marwan Abu-Khalil**

Freie Universität Berlin
FB Mathematik und Informatik
FG Datenbanken
Takustr. 9
D-14195 Berlin
D
e-Mail: marwan@inf.fu-berlin.de

Peggy Agouris

University of Maine
Dept. of Spatial Information Science & Engineering & NCGIA
348 Boardman Hall
ME 04469-5711 Orono
USA
Tel: +1-207-581-2180
Fax: +1-207-581-2206
e-Mail: peggy@spatial.maine.edu
URL: http://www.spatial.maine.edu/~peggy/peggy.html

**Tom Bittner**

Technische Universität Wien
Inst.für Geoinformatik
Gusshausstraße 27-29
A-1040 Wien
A
Tel: +43-1-58801-3786
Fax: +43-1-504-3535
e-Mail: bittner@geoinfo.tuwien.ac.at

**Thomas Brinkhoff**

Mannesmann AUTOCOM GmbH
Niederkasseler Lohweg 20
40547 Düsseldorf
D
Tel: +49-211-5368-311
Fax: +49-211-5368-304
e-Mail: tbrinkhoff@ac,.org

**Jean Brodeur**

Centre for Topographic Information
King 2144-010
QC JIJ 2E8 Sherbrooke
CDN
Tel: +1-418-819-564-4889
Fax: +1-418-819-564-4892
e-Mail: brodeur@rneau.gc.ca

Michael Böhlen

Aalborg University
Dept. of Comp. Science
Fredrik Bajers Vej 7E
DK-9220 Aalborg Ost
DK
Tel: +45-96 35 89 19
Fax: +45-98 15 98 89
e-Mail: boehlen@cs.auc.dk
URL: http://www.cs.auc.dk/~boehlen/

**Antonio Corral**

Aristotle University - Thessaloniki
Dept. of Informatics
GR-54006 Thessaloniki
GR
Tel: +30-31-996363
Fax: +30-31-996360
e-Mail: antonio@skyblue.csd.auth.gr

**José Antonio Cotelo Lema**

FernUniversität-GH-Hagen
FB Informatik
LST IV Prof. Güting
D-58084 Hagen
D
Tel: +49-2331-987-4281
e-Mail: jose-antonio.cotelo-lema@fernuni-hagen.de

Marlon Dumas

Université Joseph Fourier
Lab. LSR-IMAG
B.P. 72

F-38402 St. Martin d'Hères
F
Tel: +33-4-76 82 72 31
Fax: +33-4-76 82 72 87
e-Mail: marlon.dumas@imag.fr
URL: http://www-lsr.imag.fr/Les.Personnes/Marlon.Dumas/

Max J. Egenhofer

University of Maine
Dept. of Spatial Information Science & Engineering & NCGIA
348 Boardman Hall
ME 04469-5711 Orono
USA
Tel: +1-207-581-2114
Fax: +1-207-581-2206
e-Mail: max@spatial.maine.edu
URL: http://www.spatial.maine.edu/~max/

Marie-Christine Fauvet

LSR-IMAG Laboratory
Bureau 303b, 3. Etage, Bat. D - Group STORM
BP 72
F-38402 St. Martin d'Hères
F
Tel: +33-4-76 82 72 83
Fax: +33-4-76 82 72 87
e-Mail: marie-christine.fauvet@imag.fr
URL: http://www_lsr.imag.fr/Les.Personnes/Marie-Christine.Fauvet

Andrew Frank

Technische Universität Wien
Inst.für Geoinformatik
Gusshausstraße 27-29
A-1040 Wien
A
e-Mail: frank@geoinfo.tuwien.ac.at

Fabio Grandi

Universitá di Bologna
CSITE-CNR - DEIS
Viale Risorgimento 2
I-40136 Bologna
I

Tel: +39-51-644-3555
Fax: +39-51-644-3540
e-Mail: fgrandi@deis.unibo.it

Oliver Günther

Humboldt Universität
Institut für Wirtschaftsinformatik
Spandauer Str. 1
D-10178 Berlin
D
Tel: +49-30-2093-5743
Fax: +49-30-2093-5741
e-Mail: guenther@wiwi.hu-berlin.de
URL: http://www.wiwi.hu-berlin.de/~guenther/

Ralf Hartmut Güting

FernUniversität-GH-Hagen
Praktische Informatik IV
Informatikzentrum
D-58084 Hagen
D
Tel: +49-2331-987-4279
Fax: +49-2331-987-4278
e-Mail: gueting@fernuni-hagen.de
URL: http://www.fernuni-hagen.de/inf/pi4/

Klaus Hinrichs

Westfälische Wilhelms-Universität Münster
FB 15 Informatik
Einsteinstr. 62
D-48149 Münster
D
Tel: +49-251-833-3752
Fax: +49-251-833-3755
e-Mail: hinrichs@math.uni-muenster.de
URL: http://www.uni-muenster.de/math/inst/info/u/khh/

Christian Jensen

Aalborg University
Dept. of Comp. Science
Fredrik Bajers Vej 7E
DK-9220 Aalborg Ost
DK

Tel: +45-96 35 89 00
Fax: +45-98 15 98 89
e-Mail: csj@cs.auc.dk
URL: http://www.cs.auc.dk/~csj/

**Spiros Kaloudis**

Agricultural University of Athens
Informatics Laboratory
Iera Odos 75
GR-11855 Athens
GR
Tel: +30-1-529-4201
Fax: +30-1-529-4199
e-Mail: kaloudis@auadec.aua.gr

**Georgios Kollios**

Polytechnic University
Computer & Information Science Dept.
6 Metrotech Center
NY 11201 Brooklyn
USA
Tel: +1-909-787-4871
e-Mail: gkollios@cs.ucr.edu

Manolis Koubarakis

UMIST
P.O. Box 88
M60 1QD Manchester
GB
Tel: +44-161-200-3305
Fax: +44-161-200-3324
e-Mail: manolis@co.umist.ac.uk
URL: http://www.co.umist.ac.uk/~manolis/

Nikos Lorentzos

Agricultural University of Athens
Iera Odos 75
GR-11855 Athens
GR
Tel: +30-1-529-4175
Fax: +30-1-529-4199
e-Mail: lorentzos@auadec.aua.gr
URL: http://www.aua.GR/tmhmata/geniko/cv/lorentzos.html

Federica Mandreoli

Universitá di Bologna
CSITE-CNR - DEIS
Viale Risorgimento 2
I-40136 Bologna
I
Fax: +39-51-6443040
e-Mail: fmandreoli@deis.unibo.it
URL: http://www-db.deis.unibo.it/~mandreol

David Mark

University of New York at Buffalo
Dept. of Geography
301 Wilkeson - Ellicott Complex
North Campus
NY 14260-0023 Buffalo
USA
Tel: +1-716-645-25 45 48
Fax: +1-716-645-5957
e-Mail: dmark@geog.buffalo.edu
URL: http://www.geog.buffalo.edu/~dmark/

**Isabelle Mirbel**

Université de Nice - Sophia-Antipolis
Dept. Informatique
Parc Valrose
F-06108 Cedex 2 Nice
F
Tel: +33-4-92 07 66 50
Fax: +33-4-92 07 66 55
e-Mail: mirbel@mezzo.unice.fr

José Moreira

ENST
Computer and Network Science Department
46 Rue Barrault
F-75634 Paris Cedex 13
F
Tel: +33-1-45 81 80 63
Fax: +33-1-45 81 31 19
e-Mail: moreira@inf.enst.fr
URL: http://www.upontu.pt/~jmoreira

**Enrico Nardelli**

Universita degli Studi di l'Aquila
Dipartimento di Matematica pura e Apllicata
Via Vetoio Coppito
I-67010 Aquila
I
Tel: +39-862-433-133
Fax: +39-862-433-180
e-Mail: nardelli@univaq.it

**Christine Parent**

Université de Lausanne
Ecole des HEC
INFORGE
CH-1015 Lausanne-Dorigny
CH
Fax: +41-21-6935195
e-Mail: christine@lbd.epfl.ch

Norman Paton

The University of Manchester
Dept. of Computer Science
M13 9PL Manchester
GB
Tel: +44-161-275-6910
Fax: +44-161-275-6236
e-Mail: norm@cs.man.ac.uk
URL: http://www.cs.man.ac.uk/users/norm/

**Achilleas Pavlopoulos**

Manchester University
P.O. Box 88
M60 1QD Manchester
GB
e-Mail: achilleas.pavlopoulos@stud.umist.ac.uk

Barbara Pernici

Politecnico di Milano
Dipartimento di Elettronica e Informazione
Piazza Leonardo da Vinci 32
I-20133 Milano
I

Tel: +39-02-2399-3526
Fax: +39-02-2399-3411
e-Mail: pernici@elet.polimi.it
URL: http://www.elet.polimi.it/people/pernici

Dieter Pfoser

Aalborg University
Dept. of Comp. Science
Fredrik Bajers Vej 7E
DK-9220 Aalborg Ost
DK
Tel: +45-96 35 89 19
Fax: +45-98 15 98 89
e-Mail: pfoser@cs.auc.dk
URL: http://www.cs.auc.dk/~pfoser/

**Guido Proietti**

Universita degli Studi di l'Aquila
Dipartimento di Matematica pura e Apllicata
Via Vetoio Coppito
I-67010 Aquila
I
Tel: +39-862-433-727
Fax: +39-862-433-180
e-Mail: proietti@univaq.it

Lukas Relly

ETH Zürich
Institut für Informationssysteme
IFW
CH-8092 Zürich
CH
Tel: +41-1-632-7248
Fax: +41-1-632-1172
e-Mail: relly@inf.ethz.ch
URL: http://www-dbs.inf.ethz.ch/~relly/

Jochen Renz

Universität Freiburg
Institut für Informatik
Am Flughafen 17
D-79110 Freiburg
D

Tel: +49 (761) 203-8226
Fax: +49 (761) 203-8222
e-Mail: renz@informatik.uni-freiburg.de
URL: http://www.informatik.uni-freiburg.de/~renz/

**Peter Revesz**

University of Nebraska
Dept. of Computer Science
NE 68588-0115 Lincoln
USA
Tel: +1-402-472-3488
Fax: +1-402-472-7769
e-Mail: revesz@tamana.unl.edu

Holger Riedel

Universität Konstanz
Fakultät für Mathematik u. Informatik
Hausfach D188
D-78457 Konstanz
D
Tel: +49-7531-88 44 34
Fax: +49-7531-88 35 77
e-Mail: holger.riedel@uni-konstanz.de
URL: http://www.informatik.uni-konstanz.de/dbis/

**José Ramon Rios Viqueira**

Agricultural University of Athens
Informatics Laboratory
Iera Odos 75
GR-11855 Athens
GR
Fax: +30-1-529-4181
e-Mail: rios@auadec.aua.gr

John Roddick

University of South Australia
School of Computer & Information Science
Database Laboratory
Mawson Lakes Boulevard
SA 5095 The Levels
AU
Fax: +61-8-8302-3381
e-Mail: roddick@cis.unisa.edu.au

URL: http://www.cis.unisa.edu.au/~cisjFr/

**Miguel Rodriguez Luaces**

FernUniversität-GH-Hagen
FB Informatik
LSTV IV Prof. Güting
D-58084 Hagen
D
Tel: +49-2331-987-4282
e-Mail: Miguel.Rodriguez-Luaces@fernuni-hagen.de

Jean-Marc Saglio

ENST
Computer and Network Science Department
C 201-6
46 Rue Barrault
F-75634 Paris Cedex 13
F
Tel: +33-1-45 81 80 62
Fax: +33-1-45 81 31 19
e-Mail: saglio@inf.enst.fr
URL: http://www.inf.enst.fr/~saglio/

Simonas Saltenis

Aalborg University
Dept. of Comp. Science
E1-202
Fredrik Bajers Vej 7E
DK-9220 Aalborg Ost
DK
Tel: +45-96 35 98 31
Fax: +45-98 15 98 89
e-Mail: simas@cs.auc.dk
URL: http://www.cs.auc.dk/~simas/

Hanan Samet

Univ. of Maryland at College Park
Dept. of Computer Science
MD 20742 College Park
USA
Tel: +1-301-405 1755
Fax: +1-301 314 9115
e-Mail: hjs@umiacs.umd.edu

URL: http://www.cs.umd.edu/~hjs/

**Maria Rita Scalas**

Universitá di Bologna
CSITE-CNR - DEIS
Viale Risorgimento 2
I-40136 Bologna
I
Tel: +39-51-644-35 44
Fax: +39-51-644-35 40
e-Mail: mrscalas@deis.unibo.it

Hans-Jörg Schek

ETH Zürich
Institut für Informationssysteme
ETH Zentrum
CH-8092 Zürich
CH
Tel: +41-1-632-7240
Fax: +41-1-632-1172
e-Mail: schek@inf.ethz.ch
URL: http://www-dbs.inf.ethz.ch/

Markus Schneider

FernUniversität-GH-Hagen
Praktische Informatik IV
Informatikzentrum
D-58084 Hagen
D
Tel: +49-2331-987-4285
Fax: +49-2331-987-4278
e-Mail: markus.schneider@fernuni-hagen.de

Michel Scholl

INRIA
Domaine de Voluceau
B.P. 105
F-78153 Le Chesnay
F
Tel: +33-1-39 63 53 29
Fax: +33-1-39 63 56 74
e-Mail: michel.scholl@inria.fr

Thomas Seidl

Ludwig-Maximilians-Universität München
Institut für Informatik
Oettingenstr. 67
D-80538 München
D
Tel: +49-89-2178-2227
Fax: +49-89-2178-2192
e-Mail: seidl@dbs.informatik.uni-muenchen.de
URL: http://www.dbs.informatik.uni-muenchen.de/~seidl

Timos Sellis

National Technical University of Athens
Dept. of Elec. & Comp. Eng.
Division of Computer Science
Zografou
GR-157-73 Athens
GR
Tel: +30-1-7721-601
Fax: +30-1-7721-659
e-Mail: timos@cs.ntua.gr
URL: http://www.dbnet.ece.ntua.gr/~timos/

Spiros Skiadopoulos

UMIST
P.O. Box 88
M60 1QD Manchester
GB
e-Mail: spiros@co.umist.ac.uk

Giedrius Slivinskas

Aalborg University
Dept. of Comp. Science
E1-202
Fredrik Bajers Vej 7E
DK-9220 Aalborg Ost
DK
Tel: +45-96 35 98 31
Fax: +45-98 15 98 89
e-Mail: giedrius@cs.auc.dk
URL: http://www.cs.auc.dk/~giedrius/

Stefano Spaccapietra

EPFL
Dept. d`Informatique - LBD
IN-Ecublens
CH-1015 Lausanne
CH
Tel: +41-21-693-5210
Fax: +41-21-693-5195
e-Mail: stefano.spaccapietra@epfl.ch
URL: http://lbdwww.epfl.ch/

Myra Spiliopoulou

Humboldt Universität
Institut für Wirtschaftsinformatik
Spandauer Str. 1
D-10178 Berlin
D
Tel: +49-30-2093-5730
Fax: +49-30-2093-5741
e-Mail: myra@wiwi.hu-berlin.de
URL: http://www.wiwi.hu-berlin.de/~myra/

Yannis Theodoridis

National Technical University of Athens
Dept. of Elec. & Comp. Eng.
Zografou
GR-157-73 Athens
GR
Tel: +30-1-7721-402
Fax: +30-1-7721-442
e-Mail: theodor@dblab.ece.ntua.gr
URL: http://www.dbnet.ece.ntua.gr/~theodor/

Babis Theodoulidis

UMIST
P.O. Box 88
M60 1QD Manchester
GB
Tel: +44 161 200 3309
Fax: +44 161 200 3324
e-Mail: babis@co.umist.ac.uk
URL: http://timelab.co.umist.ac.uk

Kristian Torp

Aalborg University
Dept. of Comp. Science
Fredrik Bajers Vej 7E
DK-9220 Aalborg Ost
DK
Fax: +45-98 15 98 89
e-Mail: torp@cs.auc.dk
URL: http://www.cs.auc.dk/~torp

Nectaria Tryfona

Aalborg University
Dept. of Comp. Science
Fredrik Bajers Vej 7E
DK-9220 Aalborg Ost
DK
Tel: +45-96 35 89 19
Fax: +45-98 15 98 89
e-Mail: tryfona@cs.auc.dk
URL: http://www.cs.auc.dk/~tryfona

Harry Uitermark

Dutch Cadastre
P.O. Box 9046
NL-7300 GH Apeldoorn
NL
Tel: +31-555-28 58 06
Fax: +31-555-76 12 66
e-Mail: uitermark@kadaster.nl
URL: http://ooa.kadaster.nl

Alex Vakaloudis

Manchester University
U13/USS Bldg.
P.O. Box 88
M60 1QD Manchester
GB
Tel: +44-161-200-3388
Fax: +44-161-200-3324
e-Mail: alvak@co.umist.ac.uk
URL: http://spider.sna.co.umist.ac.uk/alvak/

**Michael Vassilakopoulos**

Aristotle University - Thessaloniki

Dept. of Informatics
GR-54006 Thessaloniki
GR
Tel: +30-31-996363
Fax: +30-31-996360
e-Mail: mvass@computer.org

Michalis Vazirgiannis

National Technical University of Athens
Dept. of Elec. & Comp. Eng.
9 Iroon - Polytechneiou Str.
Zografou
GR-157-73 Athens
GR
Tel: +30-1-7721-602
Fax: +30-1-7721-659
e-Mail: mvazirg@aueb.gr
URL: http://www.dbnet.ntua.gr/~michalis/

Agnes Voisard

Freie Universität Berlin
FB Mathematik und Informatik
FG Datenbanken
Takustr. 9
D-14195 Berlin
D
Tel: +49-30-83875-125
Fax: +49-30-83875-109
e-Mail: voisard@inf.fu-berlin.de
URL: http://www.inf.fu-berlin/~voisard/

**Rolf A. de By**

International Institute for Aerospace Survey & Earth Science
Div. of Spatial InformationTheory & Applied Computer Science
Hengelosestraat 99
NL-7500 AA Enschede
NL
Tel: +31-53-487-4553
Fax: +31-53-487-4335
e-Mail: deby@itc.nl