

Proceedings of the IASTED  
International Conference on



# DATABASES AND APPLICATIONS

Editor: M.H. Hamza

As part of the 22nd IASTED International Multi-Conference on  
**Applied Informatics**

February 17 – 19, 2004  
Innsbruck, Austria

A Publication of The International Association of  
Science and Technology for Development – IASTED

ISBN: 0-88986-412-8

ISSN: 1027-2666

3-53

A Press

Anaheim • Calgary • Zurich

Proceedings of the IASTED  
International Conference on



# DATABASES AND APPLICATIONS

Editor: M.H. Hamza



As part of the 22nd IASTED International Multi-Conference on  
**Applied Informatics**

February 17 – 19, 2004  
Innsbruck, Austria

A Publication of The International Association of  
Science and Technology for Development – IASTED

ISBN: 0-88986-412-8



E200500274

ISSN: 1027-2666

ACTA Press

Anaheim • Calgary • Zurich

# TABLE OF CONTENTS

## DBA 2004

### MOBILE, DISTRIBUTED, SPATIAL, AND TEMPORAL DATABASES

419-011: A Tool for Transforming Conceptual Schemas of Spatio-Temporal Databases with Multiple Representation <i>M. Minout, C. Parent, and E. Zimányi</i> .....	1
419-031: Logical Query Transformation in Bitemporal Databases <i>B. Stantic, J. Terry, and A. Sattar</i> .....	7
419-056: Database Schema Detection and Mapping on Mobile Applications: An Ontology-based Approach <i>V. Morocho, L. Pérez-Vidal, and F. Saltor</i> .....	13
419-063: Effective Decompositioning of Complex Spatial Objects into Intervals <i>H.-P. Kriegel, P. Kunath, M. Pfeifle, and M. Renz</i> .....	19
419-071: User Profiles in Location-based Services: Make Humans More Nomadic and Personalized <i>S. Yu, S. Spaccapietra, N. Cullot and M.-A. Aufaure</i> .....	25
419-073: Optimistic Concurrency Control for Inverted Files in Text Databases <i>M. Marin</i> .....	31
419-082: Modelling Real-time Database Systems in Duration Calculus <i>D. Van Hung and H. Van Huong</i> .....	37
419-094: Asynchronous Periodic Patterns Mining in Temporal Databases <i>K.-Y. Huang and C.-H. Chang</i> .....	43
419-099: A Locking Model for Mobile Databases in Mobile Environments <i>H.N. Le, M. Nygård, and H. Ramampiaro</i> .....	49

### INFORMATION FILTERING AND RETRIEVAL, ONTOLOGY, AND SEMANTIC WEB

419-020: Extracting Ontologies from Relational Databases <i>I. Astrova</i> .....	56
---	----

419-022: Community Formation Scenarios in Peer-to-Peer Web Service Environments <i>O. Kaykova, O. Kononenko, V. Terziyan, and A. Zharko</i> .....	62
--	----

419-026: Structural Knowledge Graph Navigator for the ICONS Prototype <i>M. Trzaska and K. Subieta</i> .....	68
---	----

419-028: Improving the Retrieval Accuracy by Dynamically Adjusting Metadata for Document Databases <i>X. Chen and Y. Kiyoki</i> .....	74
--	----

419-035: Improving Collaborative Filtering by W-clustering <i>H. Chen, K. Furuse, N. Ohbo, and S. Nishihara</i> .....	81
--	----

419-065: High Performance Associative Coprocessor Architecture for Advanced Database Searching <i>C. Layer and H.-J. Pfleiderer</i> .....	87
--	----

419-074: A New Heuristic Rule for Classification-based Relevance Feedback <i>I. Moghrabi and F. Yamout</i> .....	93
---	----

419-080: Overlapping Clustering Methods for a Japanese Meta Search Engine <i>M. Ohta, H. Narita, K. Katayama, and H. Ishikawa</i> .....	100
--	-----

419-086: Towards the Architecture of P2P-based Information Retrieval <i>H. Ding, P. Küngas, and Y. Lin</i> .....	107
---	-----

419-802: Database Oriented Parsing and Translation <i>S. Rezaei</i> .....	113
--	-----

### DATA WAREHOUSING AND DATA MINING

419-027: A Peculiar Business Intelligence Application: OLAP Applied to Fault Data at Siemens Mobile Communications <i>M. Contini and D. Tagliabue</i> .....	116
--	-----

419-039: Evaluation of the Serial Association Rule Mining Algorithms <i>F. Kovács, R. Iváncsy, and I. Vajk</i> .....	121
---	-----

419-044: Classification of Websites as Sets of Feature Vectors <i>H.-P. Kriegel and M. Schubert</i> .....	127
--	-----

419-046: A Quality Measure for Distributed Clustering <i>E. Januzaj, H.-P. Kriegel, and M. Pfeifle</i> .....	133
---	-----

419-062: The Case for an Agri Data Warehouse: Enabling Analytical Exploration of Integrated Agricultural Data <i>A. Abdullah, S. Brobst, M. Umer, and M.F. Khan</i> .....	139
--	-----

419-091: Anticipatory Clustering <i>M. Goller and M. Schrefl</i> .....	145
---	-----

419-093: Building a Near Real Time Active Data Warehouse <i>D. Johnston-Bell, C. Moore, and B. Garner</i> .....	151
--	-----

419-102: HDDV: Hierarchical Dynamic Dimensional Visualization for Multidimensional Data <i>K. Techapichetvanich, A. Datta, and R. Owens</i> .....	157
--	-----

419-804: Implementation of Temporal Databases <i>A. Gudelj, M. Krcum, and Z. Juric</i> .....	163
---	-----

419-805: Application of the Expert Diagnostic System for Marine Diesel Engine Condition <i>M. Krcum, G. Radica, and Z. Juric</i> .....	169
---	-----

## QUERY LANGUAGES AND DATABASE ALGORITHMS

419-018: The System Architecture of the GOQL Language <i>E. Keramopoulos, T. Ptohos, and P. Pouyioutas</i> .....	174
---	-----

419-040: Methods in Query Languages of Object Oriented Databases <i>K. Gilariski</i> .....	180
---	-----

419-061: Explicit and Implicit LIST Aggregate Function for Relational Databases <i>W. Litwin</i> .....	185
---	-----

419-077: Query and Relevance Feedback in Latent Semantic Index with Reduced Time Complexity <i>F. Yamout, I. Moghrabi, and M. Oakes</i> .....	191
--	-----

419-083: A Query Language Solution for Fastest Flight Connections <i>J.A. Bakker and J.H. ter Bekke</i> .....	197
--	-----

419-084: A Query Language Solution for Shortest Path Problems in Cyclic Geometries <i>J.A. Bakker and J.H. ter Bekke</i> .....	203
---	-----

419-050: A Flexible Database Authorization System <i>B. Blicharski and K. Stencel</i> .....	208
--	-----

## STORAGE AND RETRIEVAL

419-024: Implementation and Performance Evaluation of SOM-based R*-Tree <i>D.-Y. Lee, C.-S. Cheung, and S.-H. Bae</i> .....	214
--	-----

419-049: Image Retrieval using Symbol String <i>A. Amato, T. Delvecchio, and V. Di Lecce</i> .....	220
---	-----

419-055: CIMWOS: A Multimedia Archiving and Indexing System <i>N. Hatzigeorgiu, N. Sidiropoulos, and H. Papageorgiu</i> .....	226
--	-----

419-103: Optimised Hybrid Approach for Efficient Storage and Retrieval of Multidimensional Data using Hilbert Curves <i>T.K. Srivani and K.T. Jagdish</i> .....	232
--	-----

419-106: Extraction of Representative Melodies Considering Musical Composition Forms for Content-based Music Retrievals <i>K.-I. Ku and Y.-S. Kim</i> .....	238
--	-----

## XML INDEXING, QUERY PROCESSING, AND WEB-BASED DATABASES

419-023: Evaluation of a Document Database Description by Different XML Schemas <i>N. Kenab, T. Ould Braham, and P. Bazex</i> .....	244
--	-----

419-029: Combining XML and Databases — Perceptions in Six Finnish Organizations <i>J. Kontio</i> .....	250
---	-----

419-032: XML Update Processing for Incremental Refresh of XML Cache <i>S. Han, D.H. Hwang, and H. Kang</i> .....	256
---	-----

419-033: XML View Indexing using an RDBMS-based XML Storage System <i>D. Park, Y.-S. Kim, J. Park, and H. Kang</i> .....	262
---	-----

419-036: Integrating Web Information with XML Concrete Views <i>T.-T. Dang-Ngoc, H. Kou, and G. Gardarin</i> .....	268
---	-----

419-085: Indexing XML Data with a Schema Graph <i>O. Luoma</i> .....	274
419-092: SSUI: Semantic Search User Interface for Intelligent Information System <i>J. Zhou, M. Liu, and Y. Liang</i> .....	280
419-104: Design and Implementation of XML-RL Query System <i>I. Qi and M. Liu</i> .....	286

#### **PAPERS FROM OTHER *IASTED* CONFERENCES**

440-063: The Implementation of an ID System using a Cell-Phone <i>J. Kim and T. Kobayashi</i> .....	293
440-070: PN-based Security Design for Data Storage <i>L. Wang and B. Hirsbrunner</i> .....	299

<b>AUTHOR INDEX</b> .....	305
---------------------------	-----

# A Tool for Transforming Conceptual Schemas of Spatio-Temporal Databases with Multiple Representation

Mohammed Minout  
Department of Informatics & Networks  
Université Libre de Bruxelles  
1050, Brussels, Belgium  
mminout@ulb.ac.be

Christine Parent  
INFORGE  
Université de Lausanne  
CH 1015 Lausanne, Switzerland  
christine@lbd.epfl

Esteban Zimányi  
Department of Informatics & Networks  
Université Libre de Bruxelles  
1050, Brussels, Belgium  
ezimanyi@ulb.ac.be

## Abstract

Nowadays, classical conceptual models (such as ER or UML) are used for designing database applications. These classical conceptual models usually come with associated CASE tools assisting the user from the creation of the conceptual schema until the generation of a physical schema for a relational or an object-relational DBMS. However, when developing spatial or temporal databases such classical models are inadequate since they do not consider spatio-temporal concepts. Although spatial and/or temporal extensions have been proposed for these models, such extensions do not cope with the requirements of advanced geographical applications, in particular, they do not cope with multiple representations of the same real-world phenomenon. In the context of the European project MurMur we developed a conceptual model called MADS coping with spatio-temporal information having multiple representations. Typical examples of multiple representation arise with multi-scale or time-varying information. The MADS model is supported by a set of associated tools allowing both the definition of the schema and the queries of the application at a conceptual level. Such conceptual specifications are then automatically translated into the language supported by the target GIS and/or DBMS software (e.g., SQL schema definitions or queries for an Oracle database). In this paper we describe a Translator module that allows to implement conceptual MADS schemas into target platforms. It is composed of a Transformation module that translates conceptual MADS schemas into a more simple schema using only concepts supported by the target software, and a Wrapper module that generates the data structures in the language of the GIS and/or DBMS software.

**Key Words:** Conceptual Modeling, Spatio-Temporal Databases, Logic Design, Geographic Information Systems.

## 1 Introduction

Applications manipulating geodata are difficult to model due to the particularity and complexity of the spatial and temporal components. More facets of real-world entities have to be considered (location, form, size, time valid-

ity), more links are relevant (spatial and temporal links), several spatial abstraction levels often need to be represented. Thus, modeling spatio-temporal databases requires advanced facilities, such as the following:

- Objects with complex structure (e.g., non-first-normal form), generalization links, composition/aggregation links, at least equivalent to those supported by current object-oriented models. This should achieve full representational power in terms of data structures.
- Spatial objects with one or several geometries associated to different resolutions or user viewpoints;
- Alternative geometry features to support both discrete and continuous views of space.
- Temporal objects with complex lifecycle that allow users to create, suspend, reactivate, and eventually delete objects.
- Timestamped attributes that record their past, present, and future values.
- Spatio-temporal concepts for describing moving and deforming objects.
- Explicit relationships to describe structural links as well as spatial (such as adjacency, inclusion, spatial aggregation, . . .), and synchronization links (such as before, during, . . .). The knowledge of the topological links between real-world entities is an essential requirement for applications.
- Causal relationships describing the causes and effects of changes that happen in the real world.

The model must also allow defining schemas that are readable and easy to understand. A key element for achieving this double objective is the orthogonality of the structural, temporal, and spatial dimensions of the model (and more generally of the concepts of the model). Thus, whatever the concept of the model (e.g., object, relationship, attribute, aggregation), the spatial and temporal dimensions may be associated to it.

This work focuses on logical and physical design for spatio-temporal databases. This is one step of a complete design methodology for applications coping with spatial and/or temporal data having multiple representations. The contribution of this work is two-fold. (1) The presentation of MADS, a conceptual spatio-temporal model with mul-

multiple representations. The MADS model includes the advanced modeling facilities described previously. Such facilities have been derived from a generic analysis of the requirements of typical spatio-temporal applications. (b) The translation of these concepts into a logical model, called MADSLog, which is platform independent, and then into a physical model targeted to specific implementation platforms (DBMSs or GISs).

The remainder of the paper is organized as follows. Section 2 presents the MADS data model. In Section 3, a small set of transformation rules for the multi-representation, spatial, and temporal dimensions is given. Section 4 presents the architecture and design of the Translator tool. Section 5 concludes and points to directions for future research.

## 2 The MADS Data Model

MADS is a conceptual spatio-temporal model based on an extended ER model. MADS stands for Modeling of Application Data with Spatio-temporal features. In [2] the authors analyze different spatio-temporal data models along the axes of expressiveness, simplicity and comprehensiveness, formalism, and user friendliness, making the conclusion that none of the existing models satisfied all the criteria.

MADS includes a set of predefined spatial and temporal Abstract Data Types (ADTs) that are used for describing the spatial and temporal characteristics of schemas. Similarly, MADS provides a set of ADTs supporting multi-representation features. In MADS the structural, spatial, temporal, and multi-representation modelling dimensions are orthogonal, meaning that spatial, temporal and multi-representation features can be freely added to any construct of the schema (object and relationship types, aggregation links, attributes, ...).

### 2.1 Spatial and Temporal Dimensions

The concept of spatiality covers the notions of shape and location. Shape describes the geometric form associated with the representation, generally a point, a line or an area. The location allows to situate this form in space. To describe the spatiality of real-world phenomena represented in the database, MADS provides several Spatial Abstract Data Types (SADTs) [4] organized in a generalization hierarchy [3]: generic (Geo), simple (Point, Line, Simple Area, Oriented Line), and complex types (Point Set, Line Set, Complex Area, Oriented Line Set). The generic type Geo means only "this type is spatial". The definition of the precise type of each occurrence will be done at the time of its creation. Each spatial type has an associated set of methods to define and handle the instances of this type.

MADS also enables to describe continuous fields with the concept of space-varying attribute, i.e., attributes whose values are defined by a function having as domain the set

of points of any non-punctual spatial extent, i.e., any SADT value but a unique point. Spatial phenomena described by space-varying attributes can be continuous (like elevation or temperature), stepwise (like the type of crop in a cultivated area), or discrete (like mines in a minefield). Each type of function defines the kind of interpolation, if any, that can be done to compute the value(s) of the attribute.

Temporality associated to object or relationship types concerns the existence of instances in their type. Thus, temporality describes their lifecycle: objects are created, can be temporarily suspended, then reactivated, and finally removed. The lifecycle is described by a particular time-varying attribute, Status, which can take one of the four values: *not-yet-existing*, *active*, *suspended*, or *disabled*. An attribute is said to be time-varying if its values change over time, while keeping track of the changes. Like space-varying attributes, time-varying attributes are defined as a function of the time to the value domain. Keeping track of rainfall values in a given area over time is an example of information varying simultaneously in both space and time. The different values of the attribute over time are conserved, and each value is associated with a temporal element that describes its validity as seen by the application (valid time) or known to the database (transaction time). This temporal element is described by using one of the Temporal Abstract Data Types (TADTs) in MADS, which are also organized in a hierarchy.

### 2.2 Multiple Representations

There is no natural unique way to look at some phenomenon and there is no natural unique way to represent it. Usually, the same database must serve many different purposes, each of them requiring different data or the same data but with different representations. The multiple representation facilities supported by the MADS data model, rely on a simple idea [8]: the possibility to stamp any element of a schema with a representation stamp (or r-stamp) that identifies for which scale, time frame, viewpoint, etc. the element is valid. Object types, relationship types, methods, attributes (including geometry and lifecycle) can be r-stamped and therefore have different representations.

Representations may vary according to different criteria. In the MurMur project, we focused on two criteria. The first, viewpoint, is the materialization of user's needs. For example different user profiles see different information in the database. The second, resolution that specifies the level of detail of a representation. However, any other criteria can be used for defining the representations. Stamping may apply on data, whether it is object and relationship instances, or attribute values, and on meta-data, whether object and relationship type definitions or attribute definitions. Stamping allows users to personalize object and relationship types by changing their attribute composition according to their stamps, and to keep several values for the same attribute.

## 2.3 Constrained Relationships

MADS constrained relationship types are relationship types that convey spatial and temporal constraints on the objects they link. MADS includes topological and synchronization relationships as built-in constrained relationship types. For example, a topological relationship type **Inside** may be defined to link object types **Station** and **River**, expressing that the geometry of a **Station** is within the geometry of the related **River**. There are at least three constrained types for spatial relationship: topological, orientation, and metric. Each type defines a spatial constraint between the geometry of the objects linked. The MADS model proposes a range of predefined topological relationships, such as disjunction, adjacency, intersection, cross, inside, and equal.

Synchronization relationships allow specifying constraints on the lifecycles of the participating objects. They convey useful information even if the related objects are not timestamped. They allow in particular to express constraints on schedules of processes. For the semantics of the definition of synchronization relationships, MADS uses the predefined operators proposed by Allen [1], namely: before, equals, meets, overlaps, during, starts, and finishes.

We illustrate the concepts explained in the previous paragraphs with the example shown in Figure 1. It shows

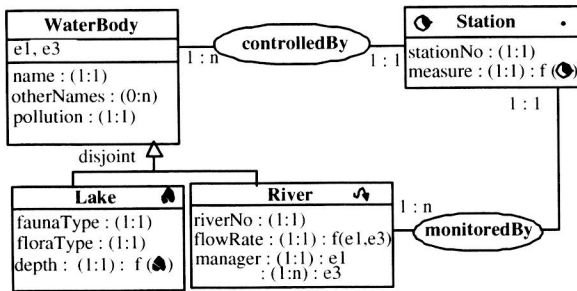


Figure 1. An example of a MADS schema.

a MADS schema for a river monitoring application. In MADS schemas, the spatial, temporal, and representation characteristics are visually represented by icons. We refer to [8] for the complete list of the spatial and temporal types available in MADS.

Temporal and spatial icons are embedded into object or relationship types. The temporal icon (symbolizing a clock) on the left-hand side of the object/relationship type expresses that the lifecycle information is to be kept. Spatial icons are shown on the right-hand side. The example schema shows the object types **River** and **Lake** that are subtypes of **WaterBody**. The **Station** type keeps water quality information about the rivers. The stations may not operate continuously: their lifecycle information allows to records periods of operation. Each station provides several measures, whose validity period is recorded. Each lake

records space-varying information of depth. The **WaterBody** object type has two representations **e1** and **e3**. The **manager** attribute has two definitions, one for **e1** and another for **e3**. Both definitions of **manager** have the same domain but have different cardinalities. One value is kept for the stamp **e1** and several values for stamp **e3**.

## 3 Transformation Rules

MADS is a spatio-temporal data model explicitly obeying the orthogonality principle in adding space, time, and multiple representation features to data structures. Therefore, the transformations of an element of the schema (for example a spatio-temporal object type) according to each modeling dimension are independent of the characteristics of the element on other dimensions. Thus, the transformation rules could be defined for each dimension, independently of the others.

The following paragraphs present some translation rules for each modeling dimension. As the rules for translating the structural dimension are well known (e.g., [5]), we do not discuss them in this paper.

### 3.1 Rules for Multiple Representation Transformations

In a database with multiple representations, each element of the schema (object type, relationship type, attribute, and method) has an r-stamp defining its visibility. This r-stamp is defined either explicitly by the designer of the database, or implicitly as being the same as that of the element to which it belongs (e.g., an object type belongs to its schema, a component attribute belongs to its composite attribute).

There is a set of rules for transforming multiple representations. We restrict in this paper to one example of transformation, the rest of these rules are given in [3].

The purpose of this rule is to transform an attribute or method having several definitions (each one associated with an r-stamp) into several attributes (one attribute by definition). Indeed, target systems (e.g., Oracle 9i, Arc View, etc.) do not allow neither a type having two attributes or methods of the same name, nor an attribute or a method having two definitions. Figure 2 shows the application of this transformation on the **manager** attribute that has two definitions in the initial schema: monovalued for the stamp **e1**, multivalued for the stamp **e3**.

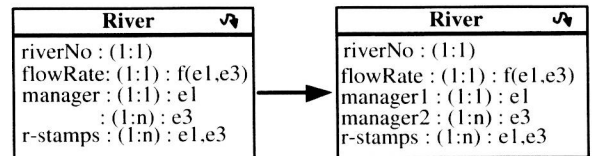


Figure 2. Transformation of **manager** attribute.

## 3.2 Rules for Spatial Transformations

MADS enables the designer to attach spatiality both to object or relationship types and to attributes. As already said, MADS provides a set of Spatial Abstract Data Types, organized in a generalization hierarchy.

### 3.2.1 Spatial Types

An spatial object or relationship type represents phenomena whose spatial reference is relevant for the application. For the target systems that do not have the concept of spatial type, but have spatial attributes, such as Oracle 9i, a spatial type is transformed by creating a monovalued spatial attribute, called **geometry**, having the same spatiality as the initial object type (see Figure 3).

### 3.2.2 Space-Varying Attributes

Although many geographical applications use space mainly in a discrete way, there is often a need to describe continuous fields, such as elevation or land use. A continuous field may be perceived in two different ways. The usual view is that it describes the values of a variable over the whole space. This is the approach chosen by [6, 7] and by all GISs offering a continuous view of the space (raster GIS). Another view is that a continuous field describes the values of an attribute over the geometry of an object type. This is the approach chosen in MADS, both for spatial objects and for spatial attributes with non-zero dimension. Space-varying attributes are represented using the function icon  $f(\cdot)$ . There are three types of functions as mentioned in Section 2.1. In this rule, the function used is continuous.

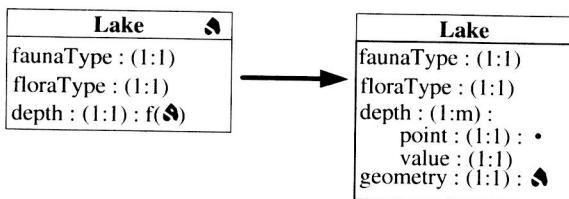


Figure 3. Transformation of spatial object type Lake and the space-varying attribute depth.

For systems that do not support the concept of space-varying attribute, the transformation consists in replacing the attribute by a multivalued complex attribute composed of a spatial element and a value. Figure 3 shows this transformation for the space-varying attribute **depth**, composed of a point and the **value** of the attribute at that point.

## 3.3 Rules for Temporal Transformations

MADS allows to associate temporality to any concept of the model (object type, relationship type, and attribute).

Temporality expresses the lifecycle of object or relationship types, and the validity of values for time-varying attributes.

### 3.3.1 Lifecycle of Temporal Types

In MADS, a temporal object (or relationship) type keeps track of the lifecycle of its instances, defined by the events of creation, suspension, reactivation, and destruction. For each instance, this information remains available after its destruction as its value (current value for nontemporal attributes, history of values for temporal attributes), so it can be able to associate information on the lifecycle of the object type corresponding of the real world.

In a similar way to the transformations for spatiality, the lifecycle of temporal types (Figure 4), is transformed by creating a monovalued time-varying attribute, called **status**, which can take as a value one of the following four states: **not-yet-existing**, **active**, **suspended**, or **disabled**.

This rule also generates a set of temporal integrity constraints associated with the states. For example, an instance cannot be active and suspended at the same time. Another integrity constraint specifies the temporal type associated with the validity of the active value: an instant for a temporal type describing an instantaneous event (and identified as such by the instant icon for its lifecycle), an interval for a temporal type having a continuous lifecycle, a set of intervals for a temporal type that can be suspended.

### 3.3.2 Time-Varying Attributes

This transformation contains two rules. The first rule works

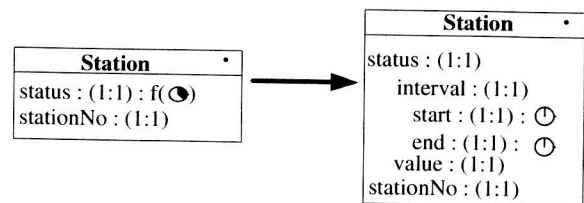


Figure 4. Transformation of attribute interval.

as the transformation of a space-varying attribute in changing the variable attribute **depth** by **status** and the point by **interval**. This rule generates an integrity constraint expressing that the intervals must be disjoint and the values must belong to the domain of the status, namely one of the states **not-yet-existing**, **active**, **suspended**, or **disabled**. However, most of existing systems have a domain of the instant type (for example the DATE type), but not a domain of the interval type of time. The second rules thus transforms an interval into a complex attribute having the same name, the same cardinality, and whose component attributes (**start** and **end**) are of instant type, as shown in Figure 4. This rule generates a integrity constraint expressing

that for each value of the attribute `interval`, `interval.start` must be less than or equal to `interval.end`.

## 4 Translator Tool

### 4.1 Design and Implementation

The Translator tool allows to translate conceptual MADS schemas into physical schemas that can be implemented on operational systems. Since the target implementation platforms have different expression power (e.g., relational DBMSs do not support multivalued attributes while object-relational DBMSs do), this translation process uses an intermediate logical model, called MADSLog, containing all concepts of MADS and a few logical concepts such as the reference attribute. The use of the MADSLog model allows making the translator extensible, minimizing the effort for adding new target platforms.

For each target GIS and DBMS, a subset of the MADSLog data model is defined containing all the concepts of MADSLog that have an equivalent in the target data model. For instance, in Oracle 9i using an object-relational model, multivalued attributes may be represented using the NESTED TABLE and VARRAY concepts. However, in Oracle 9i using a classical relational model a multivalued attribute has to be transformed into a new object type and a new relationship type. Therefore, for each target GIS and DBMS it must be determined which rules will be applied while generating the corresponding MADSLog schemas.

The Translator tool is composed of two modules. First, a Transformation module that transforms a MADS schema into a MADSLog schema adapted to the target GIS and DBMS (e.g., Oracle 9i). Second, a set of specialized modules, called Wrappers, one for each target GIS and DBMS. A Wrapper rewrites the MADSLog schema provided by the Transformator, and generates a physical schema expressed in the language of the target system (e.g., SQL scripts for Oracle 9i).

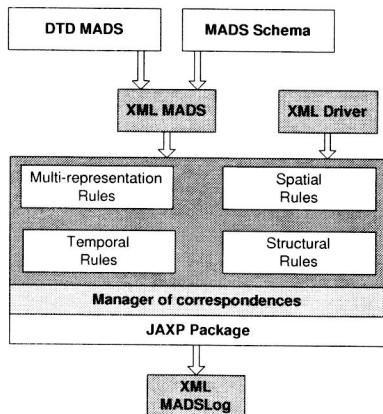


Figure 5. Architecture of the Transformation module.

The architecture of the Transformation module is shown in Figure 5. The core of the Transformation module is composed of all transformation rules for the structural, spatial, temporal, and multiple representation dimensions. Associated to each target GIS and DBMS, there is a Driver defining which rules will be applied while generating MADSLog schemas. Finally, another module keeps the correspondence between each element of the original MADS schema and its translation in the MADSLog schema. Such correspondences are used for translating queries.

We use XML as an exchange format for expressing schemas. Each MADS schema produced by a graphical schema editor is exported in XML format respecting a Data Type Description (DTD). The different drivers selecting the set of transformation rules to be activated are also expressed in XML.

The Translator has been implemented using the Java platform. A screen capture is shown on Figure 6. We used the JAXP (Java API for XML Parsing) package to implement the transformation rules. The JAXP package enables applications to parse and transform XML documents using an API that is independent of a particular XML processor implementation.

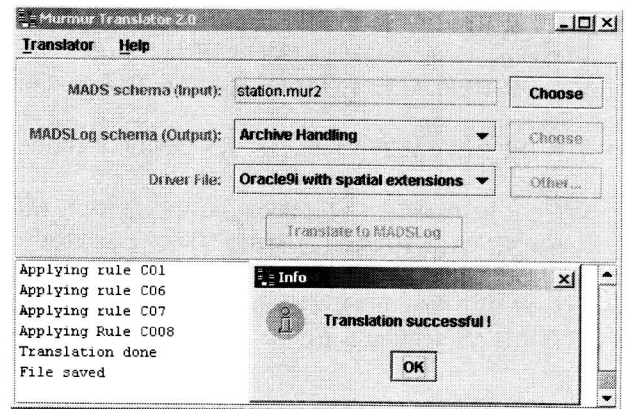


Figure 6. Interface of the MurMur Translator

The MADSLog schema is produced by applying successively the transformation rules to each element of the schema. As shown in Figure 4, the same element can be successively transformed by several rules. The total order for applying the rules is as follows: first, rules for multiple representation, then those relating to spatial and temporal dimensions, and finally the structural rules.

The tool works recursively: it applies to each node of the schema all the rules defined in the driver. If a rule is activated (is applied to the node), the process restarts again from the root of the schema. The process stops when no transformation rule is applicable to the target schema.

## 4.2 Example of Use

We take an excerpt of the example in Figure 1 (the object type Station). The transformation of this conceptual MADS schema into a MADSLog schema targetting an object-relational model is given in Figure 7. The MADSLog schema defines the Station object type with attributes station, statusNo, and geometry, as well as the DStatus domain, representing the lifecycle.

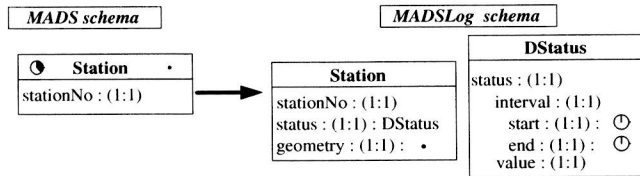


Figure 7. MADS and MADSLog schemas.

Finally, the Wrapper translates this MADSLog schema into an SQL script for Oracle 9i as follows:

```
CREATE OR REPLACE TYPE DStatus_Int AS OBJECT
(start DATE, end DATE);
CREATE OR REPLACE TYPE DStatus_Value AS OB-
JECT (interval DStatus_Int, value FLOAT)
CREATE OR REPLACE TYPE Status_Seq AS OBJECT
(status DStatus_Value, sequence INTEGER);
CREATE OR REPLACE TYPE DStatus AS TABLE OF
Status_Seq;
CREATE OR REPLACE TYPE DId AS OBJECT (value
INTEGER);
CREATE OR REPLACE TYPE DStation AS OBJECT (id
DId, stationNo VARCHAR2(12), status DStatus, geometry
MDSYS.SDO_GEOMETRY);
CREATE TABLE Station OF DStation NESTED TABLE
status STORE AS StationStatus;
```

## 5 Conclusion and Future Work

Schema translation is important for providing integration and interoperability in multiple database systems. Two shortcomings of many of the existing approaches to the problem are that they are not easily automated and that they lack a formal basis for evaluating the correctness of the resulting translations. We have discussed an approach based on structural, spatial, temporal, and multiple representations transformations that overcomes both of these problems.

This paper presents a tool for translating conceptual schemas of spatio-temporal databases with multiple representation. The ultimate goal of the tool is to be a universal translator from the MADS conceptual model into any operational data model of GIS or DBMS. The translation is made in two steps. First, translation of a MADS schema into a less expressive MADS schema, called MADSLog,

involving only the concepts available at the target data model. This translation generates a set of spatio-temporal integrity constraints. Second, translation of this MADSLog schema into a schema expressed with the syntax of the target system. This two-steps strategy allows to reuse most of the code when adding a new target data model.

We continue to study the problem of schema translation, in particular the translation of integrity constraints, implemented using triggers in target DBMSs. Further, we are also implementing a tool allowing automatic translation of visual queries expressed over the conceptual MADS schema into queries of the target DBMS or GIS (e.g., SQL queries for Oracle).

## References

- [1] J.F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, November 1983.
- [2] C. Parent, S. Spaccapietra, and E. Zimányi. Spatio-temporal conceptual models: Data structures + space + time. In *Proc. of the 7th ACM Symposium on Advances in Geographic Information Systems, GIS'99*, pages 26–33, Kansas City, USA, 1999.
- [3] C. Parent, E. Zimányi, M. Minout, and A. Aissaoui. Implantation d'un modèle conceptuel avec multi-représentation. In A. Ruas, editor, *Traité IGAT sur la représentation multiple et la généralisation*, chapter 7, pages 131–147. Hermès, 2002.
- [4] P. Rigaux, M. Scholl, and A. Voisard. *Introduction to Spatial Databases: Applications to GIS*. Morgan Kaufmann, 2000.
- [5] H. Tardieu, A. Rochfeld, and R. Colletti. *La méthode Merise: Principes et outils*. Editions d'Organisation, 2000.
- [6] N. Tryfona and T. Hadzilacos. Logical data modeling for spatio-temporal applications: Definitions and a model. In *Proc. of the Int. Database Engineering and Applications Symposium, IDEAS'98*, pages 14–23, Cardiff, U.K., 1998. IEEE Computer Society.
- [7] N. Tryfona, D. Pfoser, and T. Hadzilacos. Modeling behavior of geographic objects: An experience with the object modeling technique. In *Proc. of the 8th Int. Conf. on Advanced Information Systems Engineering, CAiSE'97*, pages 347–359, Barcelona, Spain, 1997. LNCS 1250, Springer-Verlag.
- [8] C. Vangenot. Supporting decision-making with alternative data representations. *Journal of Geographical Information and Decision Analysis*, 5(2):66–82, 2001.

# Logical Query Transformation in Bitemporal Databases

Bela Stantic, Justin Terry, Abdul Sattar  
Institute for Integrated and Intelligent Systems  
Griffith University, Brisbane Australia

email: {B.Stantic, J.Terry, A.Sattar}@griffith.edu.au

## ABSTRACT

Most database applications contain some amount of time-dependent data because the majority of database applications are temporal in nature. While a lot of work has focussed on temporal data models and query languages, only a little work has addressed access methods. Research has proved the benefits of adding time dimensions to data and keeping the history of changes, however, the informational benefits of temporal data can, without effective management, be easily outweighed by the costs of poor access times and difficulties in formulating queries. This paper proposes a logical query transformation that relies on the *POINT* representation of current time. Logical query transformation enables off-the-shelf Spatial indexes to be used. Testing results indicate a significant computational advantage with this approach. Further we propose to combine this approach with the previously introduced layered approach for temporal DBMS implementations to simplify the formulation of queries and to achieve the full benefit.

## KEY WORDS

Temporal databases, Indexing methods, Point approach, Layered architecture

## 1 Introduction

The focus of research in temporal databases has concentrated on two distinct time lines: transaction time and valid time [8]. The valid time line represents the time when a fact is valid in the modelled reality and the transaction time line represents the time when a transaction was performed. A bitemporal database records the database states with respect to both valid and transaction time. It represent reality more accurately, however as they are "partially persistent data structures", being append only, they are usually very large in size. Thus the informational benefits of bitemporal data can be easily outweighed by the costs of poor access times. The need for efficient access methods is even more crucial in bitemporal databases than in conventional databases.

Despite the amount of published research devoted to temporal databases and applications [4], [13], [7], access methods for bitemporal data have only lately come into focus as a research topic, despite being a crucial element in the criteria governing their application. Existing research demonstrates that regular indices such as

$B^+$ -trees are unsuited for temporal data [11]. However due to the similarities between bitemporal and spatial data, because valid and transaction time dimensions are considered orthogonal [12], time can be treated as a region in two-dimensional space. So spatial indices can be adapted for indexing bitemporal data.

Recent research has shown that it is extremely unlikely that industry will give up on the current relational model, as it represents huge investments in developed code and expertise. As disk storage price is dropping significantly more and more applications are adding temporal dimensions to their database systems and are thus facing deteriorating response times.

*Now-relative* data is data that is current at the present time. Despite it being a natural and meaningful part of bitemporal databases only a few index structures acknowledge their existence. Even though it is expected that they will be accessed the most often. To represent that a fact is valid *now* or that a tuple belongs to the current database state special variables may be used to represent current time. Problems with having variables to represent current time in index structures have previously been identified in the literature, so a different way to represent current time is needed. In the remainder of this paper we will refer to the *MAX* approach when current time is represented using the maximum-timestamp of the underlying database management system, and to the *POINT* approach when current tuples are represented as points, with end times equal to start times [15].

In this paper we present a method to logically divide the regions of bitemporal data and additionally to employ the geometry types feature of spatial indexes. The query transformations that rely on the *POINT* approach to represent current time, enable usage of off-the-shelf spatial R-tree indexes. Logical query transformations can be done in a layer which sits on top of the relational database system and can be combined with the previously proposed layered approach to temporal DBMS implementation. Further, the earlier hypothesis that a straightforward *MAX* approach is inefficient in the handling of *now-relative* bitemporal data [1] is given its first empirical demonstration by comparing these approaches in a real environment. The *POINT* approach is conversely shown to be significantly more efficient.

This paper is organized as follows, we first examine more closely temporal data concepts. In section 3

$ID$	$Name$	$Position$	$[Vt^+]$	$Vt^-)$	$[Tt^+]$	$Tt^-)$
1	Megan	DBA	21.08.2002	10.05.2004	<b>16.02.2002</b>	<b>UC</b>
2	Stephan	Teacher	23.07.2000	30.01.2001	01.07.2001	26.10.2002
3	Mark	Admin	<b>22.03.2001</b>	<b>now</b>	<b>01.07.2001</b>	<b>UC</b>
4	Steven	Officer	<b>21.02.2002</b>	<b>now</b>	13.04.2000	03.02.2001
...	...	...	...	...	...	...
...	...	...	...	...	...	...

Table 1. Sample *now-relative* Bitemporal data

we briefly introduce some existing methods to index *now-relative* bitemporal data and highlight their limitations and disadvantages. In section 4 we explain the logical query transformations. Section 5 presents experiment, results and analysis and introduces Layered approach. Finally, in section 6, we present our conclusions and discuss possible extensions and future work.

## 2 Temporal Data Concepts

A temporal database is one that supports some aspect of time, as distinct from user-defined time. While being ever changing, time is an important aspect of all real world phenomena. Each event bears a time attached to it, sometimes in more than one form. Time marks the starting and ending of an event and establishes the validity of data. Facts, i.e. data valid today may have had no meaning in the past and may hold no identity in the future. Some data, on the other hand may hold a historical significance or may continue to be valid up to a predefined period in time. This relationship of time and data adds a temporal identity to most data and in this light it would be hard to identify applications that do not require or would not benefit from database support for time-varying data.

*Now-relative* data are facts that are valid *now* and/or tuples that belong to the current database state. In the literature to represent that a fact is valid *now* or that a tuple is current, special variables are used, such as “now”, “UC” (until changed) or “ $\infty$ ”. The introduction of variables into temporal databases leads to the under researched area of *Variable databases* [3]. It is our intention to avoid such variables in our approach.

In bitemporal databases when an object is first inserted, its transaction time temporal attribute has the form  $[Tt^+, now)$  indicating that tuple is current and ending time is unknown. Updates are allowed to be made to the objects of the most recent version only. No modification to the past is allowed and deletion is only logical. When an object is deleted, its transaction temporal attribute is changed from  $[Tt^+, now)$  to  $[Tt^+, Tt^-)$  where  $Tt^-$  is actually the current time when the transaction is executed.

Because bitemporal databases are basically append only, they are usually very large in size [5], for that reason efficient access methods are even more important in bitemporal databases than in conventional databases.

## 3 Temporal Access Methods

A lot of multidimensional access structures have been proposed and some of them have been recommended for temporal databases [9]. The effectiveness of the majority of the proposed models have only been analysed on theoretical calculations based on worst case scenarios [11]. In this section we briefly review the relevant access methods for indexing *now-relative* bitemporal data.

### 3.1 Indexing *now-relative* data

In the literature it is generally agreed that usage of the spatial indexes based on R-trees for bitemporal data is possible and that the maximum-timestamp approach is a straightforward solution, i.e. the *MAX* approach. Further it is usually suggested as obvious that *now-relative* facts in the *MAX* approach are represented using very large rectangles, and that the resulting search performance should be poor due to excessive dead space in the index nodes and overlap between nodes [1].

The GR-Tree [1] and the 4R-Tree [2] support both *now-relative* valid and transaction time. Essentially both seek to improve the efficiency of the underlying structures by reducing dead space in the tree formed by using the *MAX* approach in a standard R\*-Tree. The GR-Tree extends the R\*-Tree to offer storage for both static tuples (i.e. tuples with closed valid and transaction time ranges) and growing tuples (i.e. tuples with valid and/or transaction time end unknown, *now-relative* data). It introducing two Flags, *Hidden* and *Rectangle*, in the non-leaf nodes where *Hidden* is used to represent growing geometries that may be placed together with other regions in a MBR (Maximum Bounding Regions), while *Rectangle* is used to indicate bounding areas that are rectangles, thus helping reduce dead space and overlap.

In our approach, we reduce the dead space by using the *POINT* approach to represent current time. We logically divide the regions and additionally employ the geometric power of spatial indexes to perform query transformations, so off-the-shelf spatial R-tree indexes can be used. Additionally we propose an extension to a previously proposed layered approach, recommended for temporal query languages. This layer can do the query transformations using predefined rules as explained in this paper.

## 4 Query Transformation

A number of queries are chosen to perform query transformations on *now-relative* bitemporal data. The queries are spatial representations of the bitemporal timeslice queries proposed in [16] and used in [15] to measure the performance of the *MAX* and the *POINT* approach. Fig 1 represents the temporal semantics of the queries. In this work we considered only query 1 and query 2, which are of most importance in bitemporal databases.

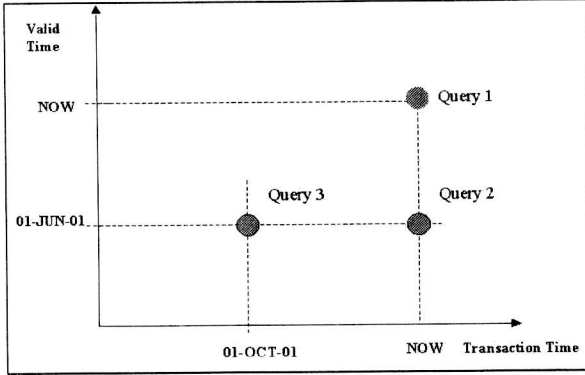


Figure 1. Time Slice Queries Q1, Q2, and Q3

**Query 1** retrieves the current state in both transaction and valid time. Semantically, it retrieves our current belief about tuples valid in the current state of the modelled reality. It selects tuples with valid time intervals that overlap with *now* and transaction times that meet *now*. It therefore selects tuples that have valid time intervals starting at or before *now* and ending after *now* and transaction time end equal to *now*.

For Query 1 using the *MAX* approach it represents the domain of tuples that meet following criteria:

$$[Vt^+ \leq now \wedge Vt^- > now \wedge Tt^- = T_{max}]$$

Spatially transformed data from table 1 are shown in Figure 2, values for  $Vt_1 \dots Vt_6$  represent the valid time values from table 1 (values  $Vtn$  associate with some specific date from table 1, similarly for transaction time values  $Tt_1 \dots Tt_5$ ). A geometry that would have  $[Vt^+ \leq now \text{ and } Vt^- \geq now \text{ and } Tt^- = T_{max}]$  would have to intersect the point  $[now, T_{max}]$ .

For Query 1 using the *POINT* approach it represents the domain of tuples that meet the following criteria:

$$[(Vt^+ \leq now \wedge (Vt^- > now \vee Vt^- = Vt^+)) \wedge Tt^- = Tt^+]$$

The Spatial transformation of the *POINT* representation for Query 1 as shown in Figure 3 is a little more complex. Since we look for all tuples where transaction time is current  $Tt^- = Tt^+$ , which means that tuples of interest for this query can be represented as points or lines, only rectangles would not be part of the result set.

Tuples where  $[Vt^+ \leq now \text{ and } Vt^- = Vt^+]$  would be points (since  $Tt^- = Tt^+$ ) and would lie below the line  $[VT = now]$ . Also, geometries where  $[Vt^+ \leq now \text{ and } Vt^- > now]$  would be line geometries, parallel to the VT-Axis since  $Tt^- = Tt^+$ , which intersect the line  $[VT=now]$ . The *POINT* query thus returns all tuples representing point geometries that intersect with the area under the line  $[VT=now]$  and line geometries parallel to the VT-Axis intersecting with  $[VT=now]$ . For the *POINT* approach the current time slice query represents a union set of the points and lines from the previously mentioned regions.

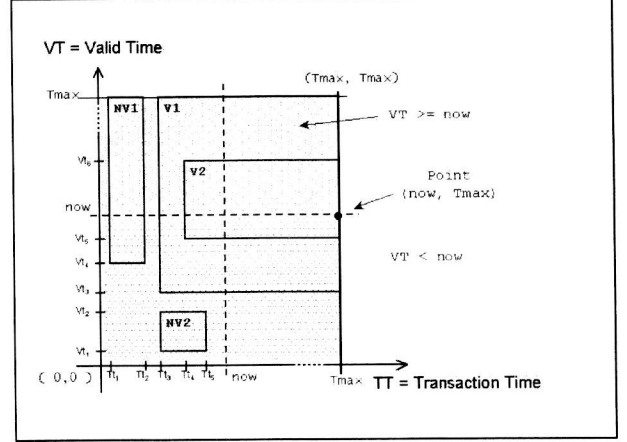


Figure 2. Query 1: *MAX* Approach

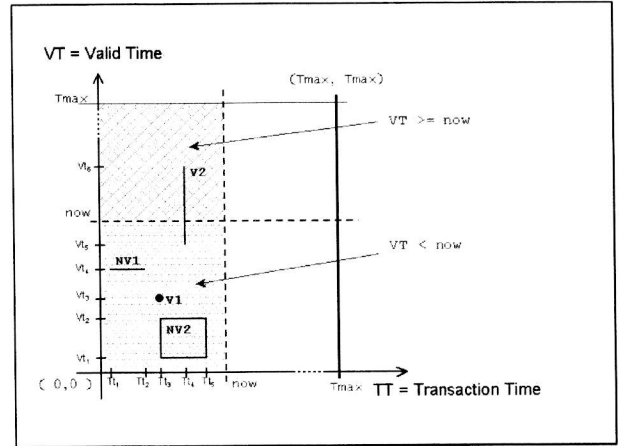


Figure 3. Query 1: *POINT* Approach

Figures 2 and 3 represent the semantics of Query 1 with tuples representing geometries  $[V1, V2]$  being returned as part of the result set while tuples representing geometries  $[NV1, NV2]$  are rejected.

**Query 2** timeslices the relation as of current time in the transaction time domain and as of a past time in the valid time domain. Semantically, it retrieves our current belief about a past state of the world. In the non-Spatial ap-

proach, the query condition for the *MAX* approach is represented by:

$$[Vt^+ \leq history \wedge Vt^- > history \wedge Tt^- = T_{max}]$$

while for the *POINT* approach query condition is:

$$[Vt^+ \leq history \wedge (Vt^- > history \vee Vt^- = Vt^+) \wedge Tt^- = Tt^+]$$

The Spatial representation for Query 2 is somewhat similar to Query 1 and geometries satisfying the query in the *MAX* approach would have to intersect the point (history,  $T_{max}$ ).

Similarly, for the *POINT* approach query which would be satisfied by all *Point* geometries that intersect with the area under the line  $[VT=history]$  and all *Line* geometries parallel to the VT-Axis intersecting with, but not ending at,  $[VT=history]$ .

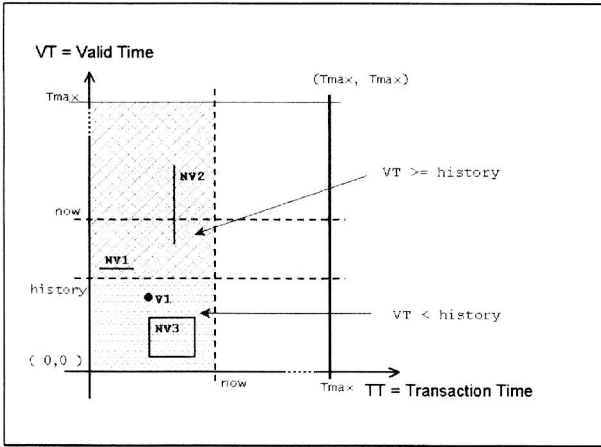


Figure 4. Query 2: *POINT* Approach

Figure 4 represent the semantics of Query 2 with the tuple representing geometry  $[V1]$  being returned as part of the result set while tuples representing geometries  $[NV1..NV3]$  are rejected.

## 5 Experiment

In order to evaluate our logical query transformation we decided to empirically compare the efficiency of this strategy, which relies on *POINT* approach, with the straight forward *MAX* timestamp approach. For each approach we created three bitemporal relations with one million randomly generated tuples, having 10%, 20% and 40% of the tuples overlapped with the current time in both transaction and valid time domains, sample data is presented in Table 1. Spatial indexes have been created on *now-relative* bitemporal data using both the *MAX* and *POINT* approaches to represent *now*. The implementation was carried out on a four 450MHZ CPU - SUN UltraSparc II processor machine, running Oracle 9.2.0.1.0 RDBMS, with a database block

size of 8K using Oracle Spatial 9i Release 2 [10].

Sample code for Query 1 and *MAX* approach is provided below:

```
select
  count(position) from max10 s
where
  (sdo_filter(get_max_geo(
    s.Vts, s.Vte, s.Tts, s.Tte),
    get_max_ge(sysdate, sysdate,
      '31-DEC-9999', '31-DEC-9999'),
    'querytype = WINDOW' )='TRUE'
  );
```

This query finds all geometries that intersect with the point (sysdate, '31-DEC-9999'). It uses the function get\_max\_geo which takes coordinates  $(Vt^+, Vt^-, Tt^+, Tt^-)$  as input returns the corresponding geometry of sdo\_geom type.

Sample code for Query 1 and *POINT* approach is:

```
select
  count(position) from point10 s
where
  ((get_same_geo(s.Vts, s.Vte,
    s.Tts, s.Tte).get_gtype()=1
  AND sdo_filter(get_same_geo(
    s.Vts, s.Vte, s.Tts, s.Tte),
    get_same_geo(
      0, sysdate, 0, sysdate),
    'querytype = WINDOW' )='TRUE')
  OR
  (get_same_geo(s.Vts, s.Vte,
    s.Tts, s.Tte).get_gtype()=2
  AND sdo_filter(get_same_geo(
    s.Vts, s.Vte, s.Tts, s.Tte),
    get_same_geo(
      sysdate, sysdate, 0, sysdate),
    'querytype = WINDOW' )='TRUE'
  );
```

This query finds all point geometries (type 1) that are within the region (0, sysdate, 0, sysdate) and all line geometries (type 2) that intersect with the line defined by (sysdate, sysdate, 0, sysdate). It uses the function get\_same\_geo, which takes coordinates  $(Vt^+, Vt^-, Tt^+, Tt^-)$  as input and returns the corresponding geometry type while taking care that different geometry types (i.e. points, lines and rectangles) are explicitly defined. Rectangle geometries (type 3) are not of interest for Query 1. Full query statements and scripts used for the population of data and index creations can be found in [14].

The Query Performance results for Query 1, (see Figures 5 and 6), support the earlier hypothesis that the straightforward maximum-timestamp (*MAX*) approach is inefficient in the handling of *now-relative* data. This inefficiency increases as the amount of *now-relative* data grows. Conversely, the *POINT* approach proved to be extremely

efficient for the handling of *now-relative* data, with little or no effect from growing *now-relative* data, to the extent of outperforming the Spatial *MAX* approach by more than a factor of 20 in terms of the number of disk accesses. According to the Theory of Indexability [6] I/O complexity cost, measured by the *number of disk accesses*, is one of the most important factors for measuring query performance.

The better performance of the *POINT* approach is achieved by reducing the dimensions of the spatial geometries. Typically, rectangles representing *now-relative* data are represented as lines or points, which take up less room in the leaf nodes than the original rectangles, leading to higher fan-out. The smaller dimensions of these geometries means there is less overlap between the maximum bounding regions, which form the boundaries of each node. In addition, because the representation of *now-relative* data in general does not extend beyond the current time, the amount of dead space (nodes that are searched but do not contribute to the answer) is reduced. Another explanation for the better performance of the *POINT* approach is its reduction of the search space, due to logically dividing the total space to the areas of interest and identifying only the geometry types of interest, which improves the performance significantly.

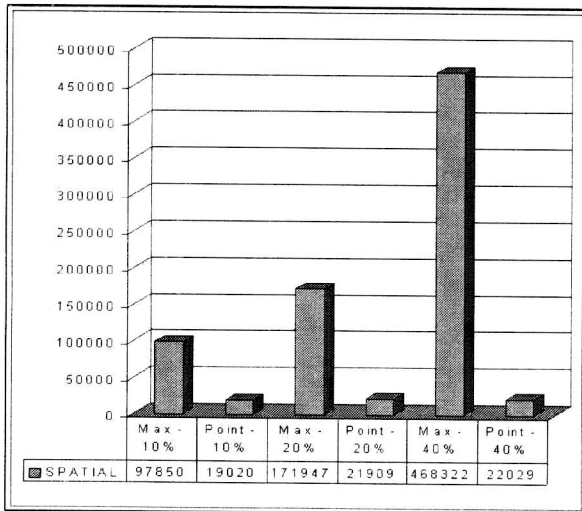


Figure 5. Query 1: Physical Disk Accesses

The large rectangles of current tuples in the *MAX* approach cause index entries to be spread across many nodes. In contrast the *POINT* approach can store many current tuples index entries in the same node, further it has a smaller total area of MBR needed to cover the same tuples and needs relatively fewer leaf nodes, which reduces the total number of disk accesses required to answer the query resulting in better performance.

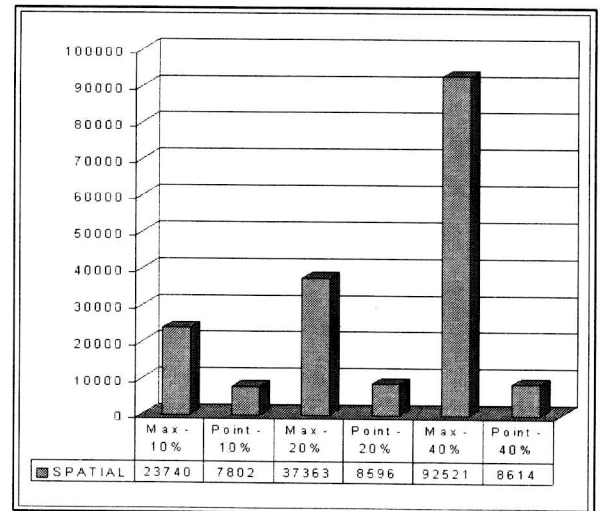


Figure 6. Query 1: CPU Usage

## 5.1 Layered approach

To take a full advantage of logical query transformation and at the same time to simplify the query statements we propose a layered approach on top of the relational DBMS. All required query transformations will happen in this layer according to the predefined rules, as identified in the previous section. This is basically an extension to the previously proposed stratum approach to Temporal DBMS Implementation [16], which can efficiently support bitemporal access methods.

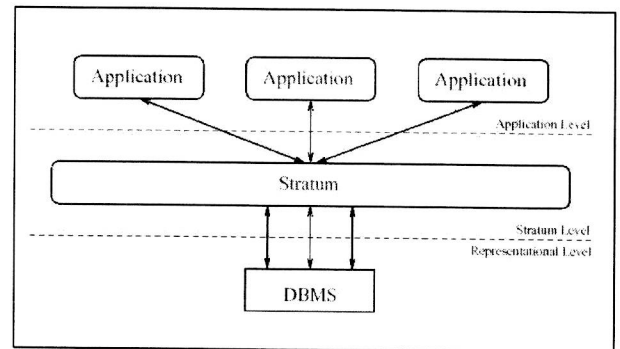


Figure 7. Stratum Approach to Temporal DBMS Implementation

This extension to the layered approach, which will do the logical query transformations to achieve a computational advantage, will enable usage of off-the-shelf spatial indexes that are part of commercial RDBMS. This layer can be developed to handle existing or any future index structures that can be developed for bitemporal *now-relative* data.

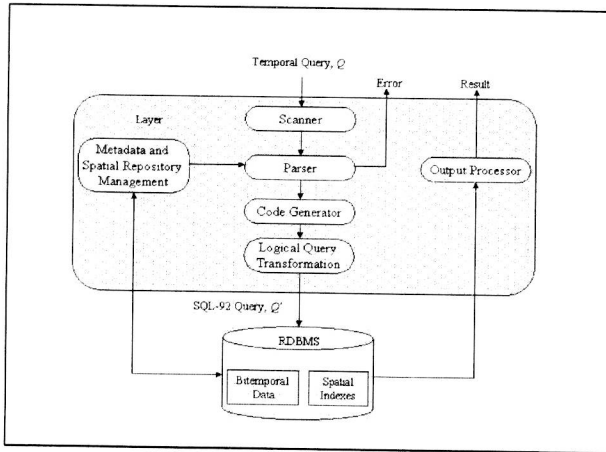


Figure 8. Layered Approach for indexing Bitemporal data

## 6 Conclusion

This study makes the following contributions to the field:

- we proposed a method, which is based only on logical transformations of queries, which reduces the dimensions of the Spatial geometries. Typically rectangles reduce to lines or points, this impacts positively as the geometries will take up less room in the leaf nodes leading to higher fanout and better performance;
- we identified the rules for logical query transformations for popular types of timeslice queries, which are the most common bitemporal queries;
- we proposed a layer on top of the RDBMS where the query transformation can be performed according to predefined rules;

By combining the Stratum approach for temporal DBMS implementation with logical query transformation, which was proposed in this paper, bitemporal data can be efficiently and simply accessed.

It would be interesting to investigate logical query transformations for other types of queries such as range queries or interval queries.

## References

- [1] R. Bliujute, C. S. Jensen, S. Saltenis, and G. Slivinskas. R-Tree Based Indexing of Now-Relative Bitemporal Data. In *VLDB'98, Proceedings of 24th International Conference on Very Large Data Bases, New York, USA*, pages 345–356, 1998.
- [2] R. Bliujute, C. S. Jensen, S. Saltenis, and G. Slivinskas. Light-Weight Indexing of General Bitemporal Data. In *Statistical and Scientific Database Management*, pages 125–138, 2000.
- [3] J. Clifford, C. Dyreson, T. Isakowitz, C. S. Jensen, and R. T. Snodgrass. On the semantics of “Now” in databases. *ACM Transactions on Database Systems (TODS)*, 22(2):171–214, 1997.
- [4] C. Date, H. Darwen, and N. Lorentzos. *Temporal Data and the Relational Model*. Morgan Kaufmann, 2002.
- [5] C. E. Dyreson, R. T. Snodgrass, and M. Freiman. Efficiently Supporting Temporal Granularities in a DBMS. Technical Report TR 95/07, 1995.
- [6] J. Hellerstein, E. Koutsoupias, and C. Papadimitriou. On the Analysis of Indexing Schemes. *16th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, 1997.
- [7] C. S. Jensen. Introduction to temporal databases, research. <http://www.cs.auc.dk/csj/Thesis/pdf/chapter1.pdf>, 2000.
- [8] C. S. Jensen and R. Snodgrass. Temporal Data Management. *IEEE Transactions on Knowledge and Data Engineering*, 11(1):36–44, 1999.
- [9] A. Kumar, V. J. Tsotras, and C. Faloutsos. Designing access methods for bitemporal databases. *University of Maryland at College Park; Report No. UMIACS-TR-97-24*, 1997.
- [10] C. Murray. Oracle Spatial User’s Guide and Reference, Release 9.2 Part No. A96630-01, 2002.
- [11] B. Salzberg and V. J. Tsotras. Comparison of Access Methods for Time Evolving Data. *ACM Computing Surveys*, 31(1), 1999.
- [12] R. Snodgrass and I. Ahn. Temporal databases. *IEEE Computer*, 19(9):35–42, 1986.
- [13] R. T. Snodgrass. *Developing Time-Oriented Database Applications in SQL*. Morgan Kaufmann, 2000.
- [14] B. Stantic and S. Khanna. The *POINT* representation of *now* in Temporal Databases. <http://miami.int.gu.edu.au/POINT/>, 2003.
- [15] B. Stantic, J. Thornton, and A. Sattar. A Novel Approach to Model NOW in Temporal Databases. In *Proceeding of the 10th International Symposium on Temporal Representation and Reasoning (TIME-ICTL 2003)*, Cairns, pages 174–181, 2003.
- [16] K. Torp, C. S. Jensen, and R. T. Snodgrass. Stratum Approaches to Temporal DBMS Implementation. *Proc. IDEAS Conf*, pages 4–13, 1998.