Temporal Databases

(A Tutorial Proposal)

Jan Chomicki and David Toman

1 Synopsis

Considerable effort in the database research is currently focused on the limitations of the classical relational model and on ways to overcome these limitations. The fruits of the research are slowly finding its place in the mainstream commercial systems, e.g., the object-relational or deductive extensions of SQL. Many of these extensions are also being incorporated in the new SQL standards.

In this tutorial we concentrate on another way of extending the capabilities of the standard relational model: the *temporal* extensions that require an a-priori interpreted data type of (linearly) ordered *time instants*. We discuss the techniques used for efficiently incorporating temporal data into a relational system and describe the additional functionality needed to support these extensions.

While the tutorial mainly concentrates on a single data type for representing time, the general techniques can be extended to other interpreted domains of data (e.g., spatial databases). We address the following issues:

1. The formal basis for introducing a new data type into a data model and the issues connected with such an extension.

2. The impact of general interpreted types on existing query languages. Mainly, we show that a careful choice of the basic operations leads to much cleaner and simpler languages while preserving (and often improving) the expressive power of the *kitchen-sink* based approaches[1].

3. The need for new storage and access methods for the new data types; mainly due to a much richer set of operations applicable to the elements of the new type.

All the above points are discussed in the setting of *Temporal Databases* and *Temporal Query Languages*. We briefly survey the classical results. However, the main emphasis is on recent developments in the area of temporal databases. In the light of several recent results [Abiteboul et al., 1996, Chomicki and Kuper, 1995, Chomicki et al., 1996, Kanellakis et al., 1995, Toman and Niwinski, 1996, Toman, 1996, Toman, 1997] many of the above issues have taken an unexpected turn and the correctness of several mainstream approaches has to be reevaluated.

[1]In many temporal data models, new features are added without considering their interaction with existing features.

2 Detailed Outline of the Tutorial

1. Introduction
   * temporal domains
   * temporal dimensions of data
2. Temporal Databases and Query languages
   * abstract vs. concrete databases
   * temporal connectives, temporal logic
   * two-sorted first-order logic
   * space-efficient encodings, intervals
   * evaluation by compilation
   * practical query languages, including TSQL2 et al.
3. Query processing for temporal databases
   * temporal indexing
   * temporal joins
4. More Powerful Languages
   * multidimensional connectives
   * higher-order features
5. Incomplete temporal information
6. Temporal integrity constraints
   * transaction-time constraints
   * dependencies and normal forms

3 Necessary Background

Only a basic familiarity with the relational model (relations, relational calculus and algebra, basic SQL) and common knowledge of elementary logic (propositional connectives, first-order quantification) are needed to follow the tutorial. All the remaining concepts will be introduced during the tutorial.

4 Target Audience and Why is this of Interest to EDBT?

The tutorial is designed to attract both the theocratically minded and the practically oriented participants of EDBT. For the theoreticians it offers an overview of recent results in the area of temporal databases, how they fit together, and what are the open issues. For the more practically oriented participants the tutorial addresses (at least) the following two points of interest:

1. It shows that naive approaches to temporal databases *cannot* work correctly, that there is little hope of fixing them, and why, and
2. It provides an feasible alternative for implementation of temporal query languages.

It also provides many directions and challenges for future research, especially in the areas of efficient implementation.

5 Brief Resume of Author(s)

Jan Chomicki
received his M.S. from Warsaw University, Warsaw, Poland, and his Ph.D. from Rutgers University, New Brunswick, New Jersey, both in Computer Science. He is currently an Assistant Professor at the Department of Computer Science, Monmouth University, Long Branch New Jersey. He has published around 25 papers on temporal, deductive and constraint databases, and has served on a number of program committees of major database conferences.

David Toman
received his M.S. degree in Computer Science from the Masaryk University, Brno, Czech Republic and his Ph.D. from Kansas State University. He was a NATO/NSERC Postdoctoral Fellow at the University of Toronto in 1996-97. Currently he is with the BRICS center for Basic Research in Computer Science, Aarhus University, Denmark. His research interests include temporal, deductive, and constraint database systems, logic programming, and programming languages.

Contact address:


```
David Toman
BRICS, Dept. of Comp. Sci., U. of Aarhus    tel: (45) 89423473
Ny Munkegaede, Bygn 540, B4.19              fax: (45) 89423255
DK-8000 Aarhus C, Denmark                   david@brics.dk
```


References

[Abiteboul et al., 1996] Abiteboul, S., Herr, L., and den Bussche, J. V. (1996). Temporal Versus First-Order Logic to Query Temporal Databases. In *ACM Symposium on Principles of Database Systems*, Montreal, Canada.

[Chomicki et al., 1996] Chomicki, J., Goldin, D., and Kuper, G. (1996). Variable Independence and Aggregation Closure. In *ACM Symposium on Principles of Database Systems*, pages 40–48, Montreal, Canada.

[Chomicki and Kuper, 1995] Chomicki, J. and Kuper, G. (1995). Measuring Infinite Relations. In *ACM Symposium on Principles of Database Systems*, pages 78–85, San Jose, California.

[Kanellakis et al., 1995] Kanellakis, P., Kuper, G., and Revesz, P. (1995). Constraint Query Languages. *Journal of Computer and System Sciences*, 51(1):26–52.

[Toman, 1996] Toman, D. (1996). Point vs. Interval-based Query Languages for Temporal Databases. In *ACM Symposium on Principles of Database Systems*, Montreal, Canada.

[Toman, 1997] Toman, D. (1997). Point-based temporal extensions of SQL. In *International Conference on Deductive and Object-Oriented Databases*.

[Toman and Niwinski, 1996] Toman, D. and Niwinski, D. (1996). First-order queries over temporal databases inexpressible in temporal logic. In *International Conference on Extending Database Technology*, Avignon, France.