Overview of the Special Section on Temporal Database Infrastructure

Richard T. Snodgrass

Department of Computer Science
University of Arizona
Tucson, AZ 85721
`rts@cs.arizona.edu`

The area of temporal databases has been active for the last fifteen years, with a corpus nearing 800 papers. Most applications need to store time-varying data, yet there are no widely used commercial temporal databases. A primary reason for the absence of technology transfer from research to practice is the lack of a commonly accepted consensus data model or query language. Even the terminology is inconsistent.

On June 14–16, 1993, the ARPA/NSF International Workshop on an Infrastructure for Temporal Databases was held in Arlington, Texas. Forty-four prominent researchers from ten countries, representing ten industrial labs and companies and over twenty universities, participated in this workshop.

This special section presents some of the results of that workshop and associated activities. The first paper is an abbreviated version of the workshop report. The conclusions of the four working groups that met during the workshop are documented there.

The second paper provides a condensed version of a glossary of temporal database concepts, ironed out through a long dialogue, first via electronic mail over nine months, then at the workshop, and then back to email, again over several months.

The third report provides the initial specification for a consensus temporal extension to SQL-92, termed *TSQL2*. This is the first definition of a temporal query language to draw on the expertise of a substantial portion of the temporal database community.

All of these reports are also available on-line, in the `tsql` directory at `FTP.cs.arizona.edu`. More complete versions are also available there via anonymous FTP. In particular, there is the complete version of the workshop report (at 59 pages), the extended version of the glossary, with more discussion of the terms (at 55 pages), a separate consensus temporal query language test suite (at 45 pages), and twenty-three commentaries discussing the major design decisions of the TSQL2 language (totaling 350 pages). Hard copies are available upon request to Robyn Austin, Department of Computer Science, University of Arizona, Tucson, AZ 85721, `robyn@cs.arizona.edu`. The commentaries, in a binder, cost $60; the other documents are free while supplies last. The proceedings of the workshop (43 position papers totaling 624 pages) is also available, at $35. Both large documents are sent via airmail in the U.S. and via ground book rate elsewhere (which can take several months). Foreign air mail book rate is available for an additional $20.

In sum, the last two years have seen a blossoming of consensual efforts within the temporal database community. The workshop and its report, the glossary, the temporal query language test suite and the TSQL2 language definition are tangible artifacts of the community's willingness to work together to produce a realistic, coherent, articulated infrastructure on which future research and industrial development can proceed.

I wish to extend a sincere thanks to the large number of dedicated individuals who participated in one or more of these efforts. Of these, several were particularly helpful. Gio Wiederhold has pushed for an infrastructure for temporal databases for many years. Maria Zemankova was instrumental in shepherding the initial workshop proposal through an NSF review. Without their diligence the workshop would never have occurred. Niki Pissinou served ably as general chair for the workshop. The members of the workshop steering committee actively organized the meeting, then captured the discussions as authors of the workshop report. Christian S. Jensen so capably lead the glossary and test suite efforts. It has been thoroughly enjoyable working with such talented people.