

## Editor's Notes

The technical papers section of this issue is dedicated to a popular research topic -- rule management in expert database systems. The guest editor of this special section is Timos Sellis, and I would like to thank him as well as the authors of the nine technical papers for their contribution (the Guest Editor's Letter elaborates on this special issue).

Following this special section is a special report on the 1988 Object-Oriented Database Workshop which has been compiled by J. Joseph, S. Thatte, C. Thompson, and D. Wells. Although a year has elapsed from that meeting, I am sure that the report will be useful and of interest to many readers. Next, there is a correspondence from Richard Snodgrass concerning Date's proposal for extending SQL to support time. This issue of the RECORD is concluded with calls for papers and participation.



Arie Segev

July 15, 1989

## Chair's Message

The recent SIGMOD election has resulted in a complete changing of the SIGMOD officers. As new Chair, I now have the daunting task of continuing the tradition of excellence and growth that my predecessor Dave DeWitt has so ably performed. With help from my able colleagues Laura Haas as Vice Chair, Michael Carey as Secretary/Treasurer, Arie Segev continuing as Editor of SIGMOD Record, and all of you, I will try my best to keep SIGMOD the foremost forum for the database R/D community.

In my discussions with many database professionals, the single most serious problem facing the database R/D community these days is the haphazard proliferation of database conferences. On the one hand, the proliferation of database conferences has been good for solidifying and popularizing databases as a discipline. On the other hand, it has contributed to the general dilution of the technical contents of most database conferences, and, because of the practical difficulty of anyone being able to attend all of the conferences, to the ineffectiveness of many of the database conferences in fostering rapid advances in emerging areas of research in databases.

I am of the opinion that the database R/D community needs a few major international conferences, and also a few regional conferences around the world. Further, I believe that there need to be a few forums for emerging areas of research. I would like the SIGMOD conference to continue to be one of the major international database conferences. However, I also would like the SIGMOD conference to serve as the forum for emerging areas of database R/D. During the next few months, I will try to develop a concrete plan to realize these goals on the basis of ideas I will solicit from many database professionals and any ideas that I may receive from you.

During the past year, some positive changes have been made to SIGMOD Record. By working with Arie Segev, I would like to help make SIGMOD Record an even more useful publication for you. One idea I have already discussed with Arie is the publication of special issues of very general interest jointly with the IEEE Data Engineering quarterly bulletin for which I have been serving as Editor-in-Chief. One topic is "Directions for Future Database R/D", which may focus on several areas of R/D that are of significant interest to SIGMOD. Another topic is "Status of Database R/D Outside U.S.A.", which may highlight R/D activities in Europe, Asia, and South America. I will work with Arie to identify issue editors and schedules for such issues.

I would like to thank you for the confidence you have shown in me by electing me to serve you for the next two years. I will do my best.

Won Kim  
July 6, 1989  
Austin, Texas

## Letter from the Guest Editor

This issue of the SIGMOD Record is devoted to Rule Management and Processing in expert database management systems. In the last few years we have witnessed an explosion in research focusing on the integration of artificial intelligence and database management. An important aspect of AI and DBMS integration is identifying functional similarities in database processing and reasoning with rules. This allows techniques designed for use in either AI or DBMS technology to be tailored for use in a functionally integrated environment.

The increased interest in the area of AI and DBMS integration has resulted in three very successful international meetings, the First International Workshop, and the First and Second International Conferences on Expert Database Systems, held in 1984, 1986 and 1987, respectively; all three were organized by Larry Kerschberg. Today, panel discussions on expert database systems are held in every major DBMS conference and commercial systems have started appearing in the market. Several new research directions are opening, such as, knowledge/data models, architectures to support both loose- and tight-coupling of knowledge base and database systems, techniques to support the manipulation, indexing, and evolution of large knowledge/data bases, query processing and optimization in knowledge bases, to name a few.

One of the most important issues in the development of expert database systems is the specification and manipulation of rules. The papers in this issue present some of the key research ideas that are currently being pursued.

In the first paper Stonebraker, Hearst and Potamianos suggest modifications to the POSTGRES rules subsystem to increase its usability and function. These include rulesets to structure knowledge bases hierarchically; rulesets can also be used with a versioning facility for customization. The second paper, by Hanson, describes the design and implementation strategy for Ariel; in particular, he focuses on integrating the rules system with transaction processing, and making the rules system powerful and efficient. In the following paper, Chakravarthy discusses the functionality of the HiPAC rule processor in terms of the requirements and how it interfaces with the rest of the DBMS components.

The fourth paper by Bocca and Freytag presents Educe\*. They examine the compilation of logic programs and its subsequent storage, retrieval and execution, and discuss their approach to rule-based query optimization and compilation. Widom and Finkelstein in the next paper present their approach to defining a clean, flexible, and complete semantics for incorporating SQL-like rules into an SQL system. In a different context, the paper by Jarke and Jeusfeld presents ConceptBase, a prototype knowledge base management system for design applications, especially in the software engineering area. They demonstrate how to model software process control information, discuss the efficient access and evaluation of rules, and present their ideas on integrating rule/constraint/query processing with temporal information and the handling of versions and configurations.

The last three papers have a common paradigm of production rules, for expert database systems. The paper by Sellis, Lin and Raschid describes the basic issues behind the DIPS

system. In particular, they focus on the mechanisms needed to efficiently monitor changes in the database which trigger productions, and the concurrent execution of production actions. Delcambre, Waramahaputi and Etheredge, in the next paper, present a similar approach taken in RPL, an SQL-like production language. They show how productions can be written in RPL and discuss the performance of a main-memory based prototype. Finally, in the last paper of this special issue, Kiernan, de Maindreville and Simon present the RDL1 production rule language and discuss the performance of their prototype, especially with respect to temporary relation management.

As with any special issue, only a few of the many current projects and research efforts on rule subsystems could be included. I would like to thank the authors for their contributions and their cooperation in meeting the deadlines during the preparation of this issue. It is my hope that this issue of the SIGMOD Record will help stimulate even more interest in this exciting area.

Timos K. Sellis

July 1989