New Technical Committee Chair and Executive Board

The current issue of the Bulletin includes the results of the recent Technical Committee on Data Engineering election for TC Chair. Betty Salzberg of Northeastern, a prior Associate Editor of the Bulletin, won election and has appointed a new executive committee. Her committee appointments are listed in her letter, which precedes this letter, and are included on the front inside cover of this and future issues of the Bulletin. I want to congratulate Betty on her election to TC Chair and welcome her and her new appointments to the Technical Committee's Executive Committee.

Changing Editorial Staff

This issue also marks on-going changes in the Bulletin's Editorial Board. Joe Hellerstein, who editted the December, 1997 issue on "Data Reduction Techniques" has completed his two year term. Joe has done a fine job both on the December issue and his prior issue on "Query Processing for Non-standard Data". As I have frequently written during such transitions, the Bulletin depends on the hard work and technical expertise of our Associate Editors, and I want to thank Joe for contributing both of those to the Bulletin.

With Joe's departure, I have now appointed Amr El Abbadi as a Bulletin Associate Editor. Amr has been at the University of California, Santa Barbara since 1987 and is currently an Associate Professor in the Computer Science Department. He has a Ph.D. in Computer Science from Cornell University. Amr's main research interests are in developing basic mechanisms for supporting distributed information management systems, including databases, digital libraries, workflow systems, and geographic information systems. Earlier work concentrated on the development of protocols and algorithms that ensure high availability and fault tolerance in such systems. I welcome Amr to the Bulletin and look forward to working with him in the future.

This Issue

As even a cursory check will reveal, data mining has become *the* hot area of database research. After years of focusing on query processing in which the idea is to make a given user query perform well, a good part of the database community has now refocused on how to extract useful information without the presence of a user query.

This new data mining activity has a longer history in the AI community. Two factors make it an interesting area for the database community.

1. Much of the data being mined was at least originally captured in a database system, and indeed, a fair bit of it is mined directly from a database system.

2. Our database community has unique expertise in the area of scalable algorithms. Given the vast amounts of data being mined, controlling the number and parallelism of the data scans is essential to the data mining enterprise.

This issue, brought together by Daniel Barbará, captures some of the excitement, controversy and diversity of this increasingly important subfield. I want to thank Daniel for his efforts, which I know directly have involved considerable difficulties. I'm sure you will find the issue both interesting and useful.