The Data Engineering Conference

The Technical Committee on Data Engineering, in addition to publishing the Bulletin, also sponsors the Data Engineering Conference, referred to as ICDE ("International Conference on Data Engineering"). The most recent conference (for 2003) was held in Bangalore, India. The 2004 conference will be held in Boston, MA. A "Call for Papers" for this conference appears on the back inside cover of this issue.

The Data Engineering Conference is one of three large and prestigous annual database conferences, the others being SIGMOD and VLDB. It is the IEEE Computer Society's flagship conference in the database area. The program for the conference is excellent, the result of a very competitive paper selection process. Because of the quality of the conference, many of the leading researchers in our field regularly attend the conference. Further, any paper published in the conference is included in the SIGMOD Anthology's CD or DVD collection of database papers. So ICDE papers have a very wide readership.

I hope that many Bulletin readers will submit papers to the Data Engineering Conference, not only in 2004, but in subsequent years as well. Perhaps I shall have the pleasure of meeting you at the conference, as I very frequently attend.

The Current Issue

Go to any recent database conference and it is immediately clear that the area of query processing over data streams is a very "hot" area of research. There are usually at least two sessions devoted to streams, sometimes more. Why has this area suddenly blossomed? There seem to me to be two very good reasons, both related to the continuous revolutionary advances in hardware, both increasing performance and decreasing costs. First, users want their traditional data immediately. They do not want to wait for it to be stored, much less processed and squirreled away in a data warehouse. Such immediate real time access to data may result in a very substantial competitive business advantage. Second, much more data is now available from a vast array of devices. Perhaps most spectacular has been the incredible shrinking size and decreasing cost of sensors (sometimes referred to as "smart dust"). These devices produce prodigous amounts of data that need to be captured in real time, aggregated, and, at times, queried in real time as well.

Data stream work is at a very early stage, and is clearly research. But there is already some industrial interest, e.g. Intel. And, as with any technology, solving problems and exploring applications can lead to a dramatic increase in interest and deployment. Johannes Gehrke, our issue editor, works in the area of streams, and knows the streams community well. The current issue is a very nice cross section of the ongoing work on streams, which nicely captures the state of this young field. I want to thank Johannes for his fine job in assembling the issue, and hope that you all enjoy it.