

Letter from the Editor-in-Chief

Bulletin Matters

Readers will, perhaps, notice that the issue available on the web has subtly changed in format. The individual papers accessed via the Bulletin web sites are now correctly page numbered, and the letters have the letter writers names included. This is an example of teaching an old dog new tricks. The teacher is Sudarshan, and I am the old dog. It seems that the dvipdfm program that I have been using to generate the entire issue can, properly parameterized, also produce individual papers quite nicely. Who knew? Well, Sudarshan knew, and now, thanks to him, I do also.

TCDE and ICDE Activities

Once again, the organizing committee for ICDE have done a fine job in hosting a top-notch conference, this time in Brisbane, Australia. The technical program was first-rate and the social calendar was a real plus.

An interesting new event at the conference was a TCDE Reception, initiated by Kyu-Young Whang, the TCDE Chair, and organized by the local ICDE organizers. This is a new step in attracting new members to the TCDE and to providing them with information and “enticements” to join.

Our new Working Group on Cloud Data Management held a workshop in conjunction with ICDE. This issue of the Bulletin contains a report on this workshop by the program chairs, Ashraf Aboulnaga and Carlo Curino.

The Current Issue

”New hardware, new problems!” That is one way to view the new hardware scene of multi-core processors, memory hierarchies, flash storage, and increasingly ”remote” disk storage. Another way to view it is ”New Hardware, new opportunities!” There is enormous compute power in modern hardware. And coupled with enormous main memory, query latency can be dramatically reduced. These open the door to order-of-magnitude throughput improvements in database engines.

But order of magnitude gains never come easily. The new hardware landscape requires new database system architectures, and new techniques, for potential gains to be turned into real gains. No longer can we simply count instructions to determine the performance of our algorithms. Even counting memory references will not do the job. Careful attention must be paid to threading models, to reference patterns, to potential blocking behavior, to how to provide durability, particularly in the presence of system crashes.

This brings me to the current issue, which is on main memory database systems. Paul Larson, the issue editor, is a colleague of mine at Microsoft Research in Redmond. He has been intimately involved with the SQL Server Hekaton main memory DBMS effort. So he knows this technical area thoroughly. He also knows about competitive systems and research efforts in this space. The result is that Paul has assembled an issue that brings together both industrial and research contributions that are truly revolutionizing DBMS performance, especially in the OLTP space, where it is frequently possible for the entire database to fit into main memory.

I am very partial to issues such as the one Paul has produced here with the help of authors. Industrial folks and research folks have much to learn from each other. And the current issue is a great example of this. Many of the successful approaches to great performance gains for DBMSs are described in these pages. I want to thank Paul for doing a fine job with the issue and to encourage you all to read and enjoy the papers contained in the issue.

David Lomet
Microsoft Corporation