The Data Engineering Conference ICDE'06

The next International Conference on Data Engineering (ICDE'06) will be held in Atlanta in April, 2006. This conference is the flagship conference of the IEEE Technical Committee on Data Engineering. Atlanta is a great venue, and April is a wonderful time to visit the city, with balmy weather and with magnolias and peach trees in bloom. This year's conference is a very selective conference with high quality papers. Additional information about the conference is available on the conference web site, `http://icde06.cc.gatech.edu/`, including the technical program.

About the Bulletin

I have made a minor change in how the Bulletin is available. Starting with the current issue, the individual papers accessed via the Bulletin web sites will be in PDF, not postscript. My primary reason for doing this is because PDF files are smaller than PS files, and hence download more quickly. I invite your comments on this change. My intent is to convert, over time, the individual papers of past issues into PDF as well. So if you do not like this turn of events, please send me email at `lomet@microsoft.com` telling me why this is a bad idea.

The Current Issue

In the database world, we extract information via very precise query languages. Moving to the world of documents has required our community to master and hopefully enhance in our own way, the technology of the information retrieval community. The world is clearly moving to putting everything online in the hopes that we will learn how to exploit this as an invaluable resource for much of what we do, surely professionally, and perhaps personally as well. We, the database community, should be able to help.

One interesting manifestation of this move toward putting things online is the rapid growth of literature digital libraries, both in professional domains and more generally. This is happening now, as anyone who has consulted DBLP will be aware. But this area also has many challenges. Perhaps we would like to know which papers have been published, to choose an area "at random", in "application recovery". This is a complicated query, much more like an IR query than a SQL query, but it requires more than simply key word search, even when augmented with web link analysis.

It is the desire to extract and exploit information such as the above that makes the current Bulletin issue so important, interesting, and timely. Gultekin Ozsoyoglu has worked in this area himself. So he brings to his editorial duties as a special issue editor, knowledge both of the field and of its research participants. The names of the authors may be less familiar to you than is normally the case with Bulletin authors. But this is an opportunity for readers to very rapidly get the feel for the exciting things that are happening with literature digital libraries and how they might be exploited. I want to thank Tekin for his fine job with the current issue. He has assembled an excellent overview of the current state-of-the-art in this increasingly important area.