Bulletin Announcement

I take great pleasure in announcing that all issues of the Data Engineering Bulletin, dating back to 1977, are now available in pdf format via the Bulletin web sites (`http://sites.computer.org/debull/` and `http://research.microsoft.com/research/db/debull/default.htm`). Further, these issues are now also referenced via the DBLP web site as well at `http://www.informatik.uni-trier.de/~ley/db/journals/debu/index.html`. I believe you will find many articles that are of great interest. Some of these articles have continued to be cited many years after their publication but have been hard to acquire until now.

Many people contributed in this effort. I want first to thank the Microsoft Corporation through which the scanning of the issues was accomplished. Next, I must thank the many people who generously contributed issues of the Bulletin from their "archives" so that we can bring to you a complete set of issues. These folks are Phil Bernstein, Umesh Dayal, Stavros Christodoulakis, Mike Franklin, Hank Korth, Guy Lohman, Amihai Motro, Timos Sellis, Gio Wiederhold, and Antoni Wolski. I want to single out for special thanks Sylvia Osborn, who provided many of the very early issues that were particularly hard to find.

Finally, I want to mention that all issues of the Bulletin are now included in and accessible via DBLP. I want to thank Michael Ley for making this happen, and in an amazingly short time so that this information can be included in the announcement.

I would urge you all to visit the Bulletin web site and browse the earlier issues. I think you will be impressed by how interesting and relevant many of the papers continue to be.

The Current Issue

Database provenance is a topic that has more or less escaped attention by our field historically. We were, as a field, pre-occupied with providing the basic data storage and retrieval functionality, ensuring that performance was adequate, and generalizing to deal with issues such as additional data models, distribution, etc. One might characterize this as needing to crawl before one can walk. But our users want us now to walk. They need to know how reliable the data is, who is vouching for it, what are the uncertainties, how precise it is, whether it is under copyright, etc. To deal with this, we need to know where the data has come from and how it was derived. That is, we need to know its provenance.

Dan Suciu together with Peter Buneman have assembled the current issue of the Bulletin on the subject of "Data Provenance". The papers in this issue were carefully selected, some coming from a recent workshop on data provenance principles. I want to thank Dan and Peter for their fine job as editors of this issue, which makes data provenance much more accessible to our community. It serves as a great introduction to an important and suprisingly subtle area.