

## Guest Editors' message

Gustavo Alonso · David Lomet · Umesh Dayal

Published online: 4 January 2008  
© Springer-Verlag 2007

The 2006 VLDB Conference in Seoul, Korea, was another very successful instance of a VLDB conference and also one that set several new records. The Program Committee received a record number of submissions in all its tracks: 334 papers in the core track, 292 in the information infrastructure track, and 59 in the industry and applications track. To put these data into perspective, VLDB 2006 got 300 more submissions in total than any of the conferences in the 1990s and between 200 and 300 more papers than the previous conferences in the new millennium. Attendance at the conference reached a new all time high with 685 registered attendees.

As PC Track Chairs, we worked hard to ensure that all papers were treated fairly. We wanted the final program to be the strongest possible, covering the exciting new areas as well as the more traditional, reflecting the best research work submitted to the conference. With the high number of submissions, the Program Committee had a huge task in selecting the papers to be accepted to the conference. The reviewing task is a difficult one with tight deadlines and a large number of papers must be thoroughly examined, understood, and written about. We want to thank our Program Committee members for their efforts in making VLDB'06 such a successful conference.

---

G. Alonso (✉)  
Department of Computer Science, ETH Zurich,  
8092 Zurich, Switzerland  
e-mail: alonso@inf.ethz.ch

D. Lomet  
Microsoft Research, One Microsoft Way, Redmond,  
WA 98052, USA  
e-mail: lomet@microsoft.com

U. Dayal  
HP Labs, Hewlett-Packard Company, 1501 Page Mill Road,  
Palo Alto, CA 94304, USA  
e-mail: umeshwar.dayal@hp.com

After the already highly selective process for paper acceptance, we then chose the papers that appear in this special issue of *The VLDB Journal* on “Best of VLDB'06”. To select the papers that would be invited for this special issue, each track followed a similar process: from the top ranked papers, we chose the ones that we thought were most significant, were specially well written, and generated the most excitement and positive reactions from the reviewers. We selected a total of eight papers, four from the core track and four from the information infrastructure track.

The best papers from the core track capture some of the exciting new trends in the evolution of database systems. (1) Gemulla et al. describe a new sampling technique that permits bounded size samples to be incrementally updated. Sampling is playing an increasingly important role in querying and its optimization, where execution cost against samples can be a good estimator of query cost over the base data. (2) XML databases have been the subject of a great many papers at recent database conferences, but none have previously attacked the problem of “XML normalization”. This topic is as important for XML as it has been for relational data, e.g., as a way to eliminate redundancy. This is treated in the paper by Yu and Jagadish. (3) Database systems are increasingly used to store data to satisfy the need to comply with legally imposed retention policies. Such data need to be protected from alteration. How to do this requires new technology if all tampering is to be precluded. This is the subject of the paper by Mitra et al. (4) The paper by Benjelloun et al explores how to incorporate multi-source data with lineage into database systems. How to manage the resulting (and now) uncertain data and how to do query processing involving these data are new challenges addressed here.

The papers in the information infrastructure track cover a wide range of topics that summarize very well the trends in the area. (5) The paper by Jeffery et al. discusses how to deal

with the noisy and unreliable nature of RFID data readers; an important contribution gives the increasingly widespread use of RFID readers in industry and the need to incorporate the data so acquired into the regular IT infrastructure of companies. (6) The paper by Parreira et al. addresses the problem of ranking information in P2P networks and how to make search more efficient in such environments. After many years where the database community was focusing only on the relational database engine, it is encouraging to see how the tried and tested techniques of data management are now been applied and exported to other, perhaps more challenging, areas. (7) The paper by Narayanan et al. also tackles a traditional problem, query optimization, but in a completely new environment and with new constraints,

thereby demonstrating again the potential of many database techniques outside the conventional database engine. (8) The paper by Bernstein et al. explores an increasingly relevant and challenging problem: that of automating and improving mappings between schemas and data collections. In doing so, they bring to the fore one of the topics that has kept researchers busy for many years but where much work remains to be done and one where important, fundamental contributions are still to be made.

The work reported on in this special issue is outstanding, and we can strongly recommend the papers to readers of *The VLDB Journal*. It has been our great pleasure to serve as Track Chairs for VLDB'06 and as guest editors for this special issue on "Best Papers of VLDB'06".