Welcome Message from the General Chair and the Organizing Committee Chair

Welcome to the 31st Very Large Data Base conference, and welcome to Trondheim and to Norway, the cradle of skiing and object oriented programming. It is 20 years since the VLDB conference last time visited Scandinavia. At that time, our good neighbors in Sweden hosted the conference. Apropos good neighbors: this year, it is 100 years since Norway broke the union between Sweden and Norway. War was threatening, but good diplomacy on both sides allowed for a peaceful dissolution of the union. Today, we like to think of Scandinavia as one of the world's safest and most peaceful places to live.

Preparing for the conference has been the most pleasant experience. This is due to the very good work of all our officers and the high standing of the conference. Special thanks go to the technical program chair Christian S. Jensen. He and his committees have all done a marvelous job.

The large database community in Trondheim has made it easy to handle all the local tasks. Not least have researchers, students, and staff from our own department at the university been engaged in the preparations for the conference.

Norway is a relatively well off country. This leads to fairly high prices. Thanks to our sponsors, we hope that our conference registration fees are at an acceptable level. We are especially pleased to allow students to register for a moderate amount.

Trondheim is one of the oldest cities in Norway. The city had its millennium anniversary in 1997. It is a former capital of Norway. When the Viking Leiv Eriksson set out for Greenland 1005 years ago, he accidentally discovered North America. His voyage started just outside the conference hotel. The historic status of Trondheim is also evidenced by the Nidaros Cathedral. Nidaros is the older name of the city, meaning the city at the mouth of the river Nid. You have the opportunity to experience some of the historic sites. In particular, the program includes a concert in the Nidaros Cathedral on Thursday night, and the reception on Monday evening is held in the old Archbishop's Palace.

Modern Trondheim is the technology capital of Norway. The Norwegian University of Science and Technology (NTNU) and several related research institutions are located in Trondheim. Technology, research, and graduates from NTNU have played a major role in transforming Norway into a modern industrial country. First came the era of hydro electric power, which started a little more than 100 years ago. Then, some 60 years later, started the exploitation of the oil fields in the North Sea. During the same period, ICT activities have flourished in both Trondheim and Norway.

Kjell Bratbergsengen
General Chair

Mads Nygård
Organizing Committee Chair

Foreword by the Program Chairs

Welcome to the *31ˢᵗ International Conference on Very Large Databases*! Over the past three decades, the VLDB conference series has established itself as a prestigious international outlet for research results in database management. The VLDB conferences have become fora for the exchange of the latest concepts and techniques in the development and use of data management technology. This year's program includes a keynote address, a 10-year best paper award presentation, a plenary panel, six tutorials, 30 prototype demonstrations, and presentations of 85 research papers and 18 industrial papers. In addition, a total of 11 workshops are co-located and affiliated with the conference.

This year, the Core Database Technology and Infrastructure for Information Systems tracks attracted 517 submissions, and the Industrial Applications and Experiences track attracted 46 submissions. These figures represent a 12% increase over last year's conference, which itself received 10% more submissions than the previous year. This year also saw a separate program committee for the increasingly popular category of demonstration proposals, of which there were 69. The acceptance rate was 16.4% for each research paper track and was near 40% for each of the other two tracks.

This year's keynote talk is by Bjørn Olstad, who is the CTO of Fast Search & Transfer. The FAST Enterprise Search Platform has been embedded as the information access layer in applications such as Siebel, EMC Storage, and Documentum; and companies like Reed-Elsevier, IBM, Dell, AOL, Factiva, and Reuters use FAST's platform to power information retrieval and analytics solutions. Bjørn is also an adjunct professor at the Norwegian University of Science and Technology, and he has published more than 70 research papers and been granted more than 30 patents.

In his talk, Bjørn will discuss some of the recent advances in search engine technology, including scalable, high-performance indexing of structured and unstructured data, fuzzy matching and ranking schemes, and techniques that involve contextual concept relationships. He will identify emerging trends in search that are brought about by advances such as these. One such trend is the mirroring of database content into a search engine in order to improve query capacity and user experience; another is the use, in applications such as CRM and document management, of search engine technology as the default access pattern for both structured and unstructured data; yet another may lead to a paradigm shift in business intelligence.

The VLDB 10 Year Paper Award is given annually to the author or co-authors of the research paper presented at the VLDB conference ten years earlier that in retrospect is found to have been the most influential paper that year. For year 2005, a committee consisting of Timos Sellis (chair), Christian S. Jensen, Richard Snodgrass, Gerhard Weikum, and Kyu-Young Whang gave this award to David Konopnicki and Oded Shmueli for their VLDB 1995 paper, "W3QS: A Query System for the World-Wide Web."

In year 1995, Mosaic was the browser of choice; and a keynote talk at that year's VLDB conference expressed concern that the Web revolution would take place without the participation of the database research community. The committee felt that this paper was ahead of its time and broke new ground. Many concepts present in this paper are now commonly accepted, including the use of a high-level query language for the web, an API for searches, different graphical representations of the gathered information, a maintenance facility, and query optimization based on indexes maintained by crawlers. At the conference, David and Oded will talk about database-inspired search on the Web.

This year's conference has one panel, titled "Database Publication Practices." The panel is motivated in part by the increases in submission figures at database conferences. For example, the submissions to VLDB have increased from about 300 in 1995 to this year's figure of 517. This has lead to larger PCs and larger reviewing loads. Some feel that this has adversely affected the quality of the reviewing and has increased the randomness

of PC decisions. The panel examines the past and present publication practices in the database research community, taking into account both conference and journal publication; and it explores possible changes to the current practices. The panelists are Phil Bernstein (chair), Surajit Chaudhuri, David DeWitt, Zach Ives, Christian S. Jensen, and Kyu-Young Whang. A plenary session is devoted to this panel due to its broad appeal to the conference participants.

This year's program also includes six tutorials that cover a wide range of topics in data management. The tutorials are the following:

- Approximate Joins: Concepts and Techniques, by Nick Koudas (University of Toronto) and Divesh Srivastava (AT&T Labs)

- Contextual Insight in Search: Enabling Technologies and Applications, by Aleksander Øhrn (Fast Search & Transfer)

- Offline and Stream Algorithms for Efficient Computation of Synopsis Structures, by Sudipto Guha (University of Pennsylvania) and Kyuseok Shim (Seoul National University)

- Personalized Systems: Models and Methods From an IR and DB Perspective, by Yannis Ioannidis and Georgia Koutrika (University of Athens)

- Semantic Overlay Networks, by Karl Aberer and Philippe Cudré-Mauroux (Swiss Federal Institute of Technology)

- XML Full-Text Search: Challenges and Opportunities, by Sihem Amer-Yahia (AT&T Labs) and Jayavel Shanmugasundaram (Cornell University)

The technical program is the result of the combined efforts of a large team. The four program committees had some 171 members, who each typically was responsible for reviewing a dozen papers. Klaus Dittrich coordinated the expanding program of affiliated workshops, Masaru Kitsuregawa was in charge of tutorials, Tore Risch handled panels, Betty Salzberg chaired the program committee of the affiliated Ph.D. workshop, and Øystein Torbjørnsen handled exhibitions. We thank them all for putting their expertise to work on the task of creating an exciting program. Klemens Böhm was in charge of compiling these proceedings – we thank him and his team for doing an expert job. Finally, we extend our thanks to the team of local organizers headed by Kjell Bratbergsengen and Mads Nygård. It has been a pleasure working with them.

Christian S. Jensen   **Laura M. Haas**   **Martin L. Kersten**

Per-Åke Larson   **Beng Chin Ooi**