Proceedings
of the
Twenty-fifth
International Conference
on
Very Large Data Bases

Edinburgh, Scotland
7-10th September 1999

Editors:

Malcolm Atkinson
Maria E. Orlowska
Patrick Valduriez
Stan Zdonik
Michael Brodie

Ordering Information

Morgan Kaufmann Publishers is the exclusive world-wide distributor for the VLDB proceedings volumes listed below:

	ISBN
1999 Edinburgh, Scotland	1-55860-615-7
1998 New York, USA	1-55860-566-5
1997 Athens, Greece	1-55860-470-7
1996 Mumbai (Bombay), India	1-55860-382-4
1995 Zurich, Switzerland	1-55860-379-4
1994 Santiago, Chile	1-55860-153-8
1993 Dublin, Ireland	1-55860-152-X
1992 Vancouver, Canada	1-55860-151-1
1991 Barcelona, Spain	1-55860-150-3
1990 Brisbane, Australia	1-55860-149-X
1989 Amsterdam, The Netherlands	1-55860-101-5
1988 Los Angeles, USA	0-934613-75-3
1987 Brighton, United Kingdom	0-934613-46-X
1985 Stockholm, Sweden	0-934613-17-6
1984 Singapore	0-934613-16-8
1983 Florence, Italy	0-934613-15-X
1995-1999 5-year set	1-55860-626-2 ($185.00)
1987-1999 13-year set	1-55860-627-0 ($400.00)

Prices are $50.00 per copy for the 1999 volume, $40.00 per copy for all other volumes.

Shipping is free from Morgan Kaufmann within the U.S on prepaid orders. International shipping costs are $7 per volume via DHL/regular mail combination, or $20.00 per volume via international overnight courier. Morgan Kaufmann U.S. accepts credit card payments: the buyer should provide card number, expiration date, and name as it appears on the card for Visa, MasterCard, or American Express credit cards. Morgan Kaufmann also accepts cheque payments in U.S. dollars only; cheques must be drawn on a U.S. bank.

Order from Morgan Kaufmann Publishers
By Mail: Morgan Kaufmann Publishers
Attention: Order Fulfilment Department
6277 Sea Harbor Drive
Orlando, FL 32887 USA

By Phone: 800-745-7323 (from within US \& Canada) and 407-345-3800 (international)
By Fax: 800-874-6418 or 407-345-4060
By Email: orders@mkp.com
By Web: http://www.mkp.com

VLDB 1999 ISBN 1-55860-615-7
ISSN 0730-9317

Conference Organisers

General Conference Chair	Keith Jeffery, CLRC Rutherford Appleton Lab.
Organisation Chair	Jessie Kennedy, Napier University
Programme Chair	Malcolm Atkinson, University of Glasgow
Regional Programme Chairs	
Asia and Australia	Maria Orlowaska, University of Queensland
Europe and Africa	Patrick Valduriez, INRIA
North and South America	Stan Zdonik, Brown University
Industrial Programme Chair	Michael Brodie, GTE Laboratories
Panel & Demonstrations Chair	Ron Morrison, University of St Andrews
Tutorial Programme Chair	Carole Goble, University of Manchester
Organising Committee	
Exhibitions	Lachlan MacKinnon, Heriot-Watt University
Local Arrangements	Albert Burger, Heriot-Watt University
Proceedings Production	Peter Barclay, Napier University
Publicity Chair	Mary Garvey, University of Wolverhampton
Secretary	Jon Kerridge, Napier University
Social Programme	Alex Gray, University of Wales, Cardiff
Treasurer	Mike Jackson, University of Wolverhampton

Programme Committees

Programme Committee Members Asia and Australia

Chair Maria E Orlowska, University of Queensland, Australia
David Abel, CSIRO, Australia
David Cheung, The University of Hong Kong, Hong Kong
Alan Fekete, University of Sydney, Australia
Joseph Fong, City University of Hong Kong, Hong Kong
Yahiko Kambayashi, Kyoto University, Japan
Kamal Karlapalem, Hong Kong University of Science & Technology, Hong Kong
Hiroyuki Kitagawa, University of Tsukuba, Japan
Masaru Kitsuregawa, University of Tokyo, Japan
Dik Lee, Hong Kong University of Science & Technology, Hong Kong
Xuemin Lin, University of New South Wales, Australia
Tok Wang Ling, National University of Singapore, Singapore
Hongjun Lu, National University of Singapore, Singapore
Ramamohanarao (Rao) Kotagiri, The University of Melbourne, Australia
Akifumi Makinouchi, Kyushu University, Japan
Shojiro Nishio, Osaka University, Japan
Beng Chin Ooi, National University of Singapore, Singapore
Dimitris Papadias, Hong Kong University of Science & Technology, Hong Kong
Ron Sacks-Davis, RMIT Multimedia Database Systems, Australia
Rodney Topor, Griffith University, Australia
Kyu-Young Whang, Korea Advanced Institute of Science and Technology, Korea

Programme Committees Continued

Programme Committee Members Europe and Africa

Chair Patrick Valduriez, INRIA, France
Serge Abiteboul, INRIA, France
Gustavo Alonso, ETH Zurich, Switzerland
Elena Baralis, Politecnico di Torino, Italy
David Bell, University of Ulster, Northern Ireland
Sonia Berman, Cape Town University, South Africa
Elisa Bertino, Politecnico di Milano, Italy
Catriel Beeri, University of Jerusalem, Israel
Kjell Bratbergsengen, University of Trondheim, Norway
Michael Bohlen, Aalborg University, Denmark
Alex Buchmann, University of Darmstadt, Germany
Richard Connor, University of Glasgow, Scotland
Oscar Diaz, University of the Basque Country, Spain
Jean Ferrie, University of Montpellier, France
Dana Florescu, INRIA, France
Piero Fraternali, Politecnico di Milano, Italy
Theo Haerder, University of Kaiserslautern, Germany
Yannis Ioannidis, University of Athens, Greece
Matthias Jarke, RWTH Aachen, Germany
Genevieve Jomier, University of Paris 9, France
Leonid Kalinichenko, Russian Academy of Sciences, Russia
Nabil Kamel, American University in Cairo, Egypt
Daniel Keim, University of Munich, Germany
Martin Kersten, CWI, the Netherlands
Guido Moerkotte, University of Mannheim, Germany
Giansalvatore Mecca, Universita della Basilicata, Italy
Michele Missikoff, IASI-CNR, Italy
Alain Pirotte, University of Louvain, Belgium
Philippe Pucheral, University of Versailles, France
Oded Shmueli, Technion Israel Institute of Technology, Israel
Stefano Spaccapietra, EPF Lausanne, Switzerland
Anthony Tomasic, INRIA, France
Ozgur Ulusoy, Bilken University, Turkey
Yannis Vassiliou, National Technical University of Athens, Greece

Programme Committees Continued

Programme Committee Members **North and South America**

Chair Stan Zdonik, Brown University, USA
Swarup Acharya, Bell Labs, USA
Rafael Alonso, Sarnoff Labs, USA
Jose Blakely, Microsoft, USA
Anthony Bonner, University of Toronto, Canada
Alex Brodsky, George Mason University, USA
Michael Carey, IBM Almaden Research Center, USA
Mariano Consens, University of Waterloo, Canada
Isabel Cruz, Worcester Polytechnic Institute, USA
Susan Davidson, University of Pennsylvania, USA
Laurent Daynes, Sun Labs, USA
David DeWitt, University of Wisconsin at Madison, USA
Umesh Dayal, HP Labs, USA
Pam Drew, The Boeing Company, USA
Max Egenhofer, University of Maine, USA
Michael Franklin, University of Maryland, USA
Dina Goldin, University of Massachusets at Boston, USA
Nat Goodman, Compaq Corp., USA
Joe Hellerstein, University of California at Berkeley, USA
Rick Hull, Lucent Technology, USA
Tomasz Imielinski, Rutgers University, USA
Roger King, University of Colorado at Boulder, USA
Dennis McLeod, USC, USA
Stuart Madnick, MIT Sloan School, USA
Dave Maier, Oregon Graduate Institute, USA
Alberto Mendelzon, University of Toronto, Canada
Gail Mitchell, GTE Labs, USA
Pat O'Neil, University of Massachusetts at Boston, USA
Frank Olken, Lawrence Berkeley National Lab, USA
M. Tamer Özsu, University of Alberta, Canada
Xiaolei Qian, Securesoft, USA
Ken Ross, Columbia University, USA
Elke Rundensteiner, Worcester Polytechnic Institute, USA
Betty Salzberg, Northeastern University, USA
Len Shapiro, Portland State University, USA
Nandit Soparkar, University of Michigan, USA
Praveen Seshadri, Cornell University, USA
Rick Snodgrass, University of Arizona, USA
Jacob Stein, Sybase, USA
Jeff Ullman, Stanford University, USA
Bennet Vance, IBM Almaden Research Center, USA
Yelana Yesha, NASA, CESDIS and UMBC, USA

Additional Reviewers

Abdulghane A.
Abileah S.
Al-Halimi R.
Albert J.
Avadhanam S.
Baradaram N.
Barish G.
Barrera R.
Barta A.
Bettini C.
Blott S.
Bordia A.
Bouganim L.
Budiarto
Buneman P.
Burkowski F.
Camara G.
Cao Q.
Cart M.
Catania B.
Celis P.
Cetintemel U.
Chandan K.
Chang L.
Chen L.
Cho W-S.
Chrisopoulos A.
Chuan Wu M.
Cicekli N.
Claypool K.
Cosar A.
Crabtree J.
Crestana V.
Dayal V.
Delcambre L.
Desai S.
Ding L.
Dogac A.
Dunkel B.
Finance B.
Flewelling D.
Fonseca F.
Formica A.
Fox N.
Freire J.
Fujikawa K.
Galindo-Legaria C.
Gardarin G.
Garofalakis M.
Gibbons P.
Gorafalakis M.
Goyal R.
Graefe G.
Gravano L.
Guerrini G.
Gupta H.
Haas L.
Hara T.
Hatanaka A.
Heiler S.
Heinz S.
Hinneburg A.
Hornsby K.
Ishii H.
Ishikawa Y.
Jagadish H.
James K.
Jensen P.
Joseph A.
Kalnis P.
Karagoz P.
Katayama K.
Kemme B.
Khan L.
Köller A.
Kolahdouzan M.
Kornacker M.
Kothuri R.
Kuo L.
Lee W-C.
Lee S.Y.
Lei Y.
Li C.
Liebig C.
Liefke H.
Liu H.
Loeser H.
Loukopoulos T.
Lynch C.
Mahnke W.
Mamoulis N.
Mantzourogiannis M.
Marder U.
Martinez J.
Merialdo P.
Mihaila G.
Mishra N.
Miyazaki J.
Monties S
Morishima A.
Muslea I.
Nakano M.
Nayak N.
Nierman A.
Nestorov S.
Ngu A.
Ohmori T.
Okada A.
Oria V.
Orlowski M.
Palpanas T.
Pant G.
Park C-M.
Park Y.C.
Patel B.
Pizzicannella R.
Plazanet C.
Poosala V.
Psaila G.
Rafiei D.
Ramer A.
Rao J.
Rastogi R.
Reddy K.
Reza M.
Ritter N.
Rochat P.
Safar M.
Sahuguet A.
Satoh T.
Seshradi S.
Sellis T.
Shapiro W.
Sharma J.
Shepherd J.
Shi X.
Shimojo S.
Simeon J.
Snider T.
Stefanidis A.
Steiert H-P
Stroe D.
Subramanian S.
Suciu D.
Surjanto B.
Takakura H.
Tamassia R.
Tan K.L.
Tandon A.
Tatbul N.
Teisseire M.
Teorey T.
Theodoridis Y.
Thom J.
Toman D.
Tompa F.
Toroslu I.
Vassalos V.
Voruganti K.
Wallace C.
Wang C.
Wang Y.
Wang K.
Weddell G.
Williams H.
Wong R.
Yang J.
Yeh L.
Yesha Y.
Yokota H.
Yu Y.
Yu J.X.
Zaman K.
Zhang N.
Zhang X.
Zhao W.
Zhao Y.
Zimmerman J.
Zjang X.
Zou C.

Sponsors

Principal Sponsor

Sponsors

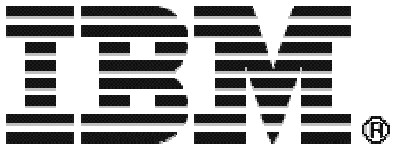

Contributors

Organising and Supporting Bodies

VLDB Endowment

Board of Trustees

President:	John Mylopoulos
Vice-President:	Keith Jeffery
Treasurer:	Stanley Su

Members:

Rakesh Agrawal
Michael L. Brodie
Michael J. Carey
Stefano Ceri
Umeshwar Dayal
Klaus R. Dittrich
Jim Gray
Yannis Ioannidis
Martin L. Kersten
Masaru Kitsugerawa
Maria E. Orlowska
M. Tamer Özsu
Hans-J. Schek
Timoleon K. Sellis
Patrick Valduriez
Gerhard Weikum
Kyu-Young Whang
Jennifer Widom

Preface

The Twenty-fifth International Conference on Very Large Data Bases (VLDB99) was held in Edinburgh, Scotland from the 7th to the 10th September 1999. This book contains all of the material prepared for VLDB99 and as such represents a valuable compendium of current database research, challenges and applications.

These proceedings contain 51 papers that were selected from the 387 submitted papers, three solicited industrial sessions, the ten-year award paper, an extended abstract of one of the two keynote talks, an overview of the panel session and eight selected demonstration descriptions. This provides a total of 77 high-quality papers contributed by 206 authors. It therefore provides excellent insights into:

- recent achievements in database research,
- pressing problems facing the users of databases that will challenge researchers for the coming years, and
- effective techniques for building and using databases.

The conference also included five tutorials:

From Semistructured Data to XML	Dan Suciu, AT&T Labs, USA
A Database Centric Update on CORBA	Sean Baker, IONA Technologies, Ireland
Design and Perception in Information Visualisation	Matthew Chalmers, University of Glasgow, Scotland
Using SQL/J for enterprise database applications: access, procedures and storage	Brian Becker, Oracle Corporation
Metasearch Engines: Solutions and Challenges	Clement Yu, U. of Illinois at Chicago, USA Weiyi Meng, SUNY at Binghamton, USA

A collection of these tutorials is available from the organisers as a separate book.

The papers and tutorials show that there continues to be interest and progress in all of the traditional database topics: optimisation, caching, data model extension, and interaction with programming languages. Sophisticated decision support, based on OLAP, provokes progress in optimisation and approximation. The significance of the large volume of data held as text attracts attention, particularly as XML promises to offer a common notation, so that the real issues of extracting information can be addressed. The interaction with networks is addressed from two points of view: what must databases do to contribute to their management and how can databases be used in conjunction with expanding networks. Considerable practical progress in extending relations towards an object-oriented capability is reported. The integration of databases with information retrieval is demonstrated.

Our ten-year award celebrates the significant contribution of the ARIES algorithm for logging and recovery, and the contribution of Mohan in continuing to develop this widely applicable research topic. The algorithm was first published at VLDB89 in conjunction with its application to nested transactions *ARIES/NT: A Recovery Method Based on Write-Ahead Logging for Nested Transactions*, K. Rothermel and C. Mohan. The programme committee therefore chose to recognise both C. Mohan, IBM Fellow (and currently Visiting Scientist at INRIA) and K. Rothermel, Universität Stuttgart with the ten-year award and to invite Mohan to give the opening talk at the conference in recognition of his sustained contributions.

Sadly, this year we had to include two obituaries. One, for Peter Stocker, recognising many years of contribution to database research and to VLDB conferences. The other, for Cheng Hian Goh, whose potential

for such contributions, evidenced by being an author of two of this year's selected papers, was abruptly terminated by his sudden death.

Under the direction of Malcolm Atkinson, the Programme Committees have done a tremendous job in attracting the largest number of submissions and in maintaining the quality of VLDB as evidenced by these proceedings. The three regional programme committee chairs, Maria E. Orlowska, Patrick Valduriez and Stan Zdonik, and the industrial programme committee chair, Michael Brodie, would like to thank all of their programme committee members who worked so hard and conscientiously to review all of the submitted papers on a very tight schedule. Over one thousand reviews had to be produced in a month. The North and South America region had a particularly arduous task, with 17 papers to review per programme committee member. These programme committee members, in particular, will join us in thanking the 180 colleagues who assisted with the task of reviewing. The processing of this workload was supported by the Puma / Predator system, supplied by Praveen Sesandri, Cornell University. Ela Hunt, Stewart Macneill and Karen Renaud, at the University of Glasgow ran the web site and databases, supported the reviewers and regional chairs and "kept the show on the road". We greatly appreciate their combined efforts. Kathy Humphry helped with many of the tasks that followed paper selection, and Stewart Macneill processed all of the proceedings into its final digital form. Peter Barclay managed the printing, publishing and production process, while Mary Garvey diverted her artistic talents from publicity to prepare the cover.

The conference itself however includes more than is reflected in the proceedings and behind the scene work of several people deserve particular note. Many people in the organising committee under Jessie Kennedy's leadership have committed many hours of careful work preparing for the conference and running it. Mike Jackson, the treasurer believed we could afford the EICC when it didn't even exist and controlled our budget and cash flow when temptations arose. Mary Garvey deserves two thank you mentions as she did the job of at least two people: maintaining the VLDB web site and Publicity. When she agreed to undertake publicity, the web site was another little add-on, however she has done a tremendous job with the web pages, which devoured many of her evenings. Peter Barclay did a thorough job of negotiating with the printers and getting the proceedings delivered on time. A first for VLDB was the demonstrations selected and organised by Ron Morrison. The excellent selection of tutorials was managed by Carole Goble and the exhibition was organised by Lachlan MacKinnon while Albert Burger helped with local arrangements. With the assistance of Jenny Siegel, Jon Kerridge acted secretary with the proviso that he'd metamorphose into Jessie should she not last the pace, although we're not sure how convincing he'd have been! Jon did a stalwart job of technical manager responsible for the student helpers who aided the smooth running of the conference. Alex Gray's expertise and enthusiasm for Scottish culture and dance enhanced the social programme, an important aspect of any international conference.

From the outset the VLDB endowment under the leadership of Peter Lockeman and subsequently John Mylopolous were very supportive of the Scottish bid, for which we would like to thank them. Although the VLDB committees undertook a lot of work, the conference wouldn't have been possible without the help of our professional conference organisers, Clansman Monarch, whose professional support was much appreciated. Conferences like VLDB would be much less viable without the support of the industrial sponsors whose contributions were gratefully received.

If you collect these proceedings as a VLDB99 delegate we thank you for coming, welcome you and hope that you will enjoy both Edinburgh and the Conference. If you are reading them after the conference then we hope they will prove beneficial to your research or business.

Jessie Kennedy & Malcolm Atkinson

Obituary

In Memory of Cheng Hian GOH (1965 - 1999)

Cheng Hian Goh passed away on 1 April 1999 at the age of 33. He is survived by his wife Soh Mui Lee and two sons, Emmanuel and Gabriel, to whom we all send our deepest condolences.

Cheng Hian received his PhD in Information Technologies from the Massachusetts Institute of Technology in Feb 1997. Prior to undertaking his PhD studies, Cheng Hian received both a BSc (First Class Honors) and a MSc both in Computer Science from the National University of Singapore where he received numerous awards. He joined the Department of Computer Science, National University of Singapore as an Assistant Professor in November 1996.

He was totally dedicated to his work— teaching as well as research. He enjoyed preparing his lecture notes, and would spend much time in organizing them. He believed in giving his students the best. As a young database researcher, he demonstrated his research capability and made major contributions to the field as testified by his publications in ICDE'97, VLDB'98, ICDE'99, VLDB'99, etc.

He did not stop working/thinking about his research until he was forced to stop by his ill health. Towards his last days, he continued to guide his research students in the hospital. He also completed and co-authored several manuscripts. The submitted versions of his two papers that appear in the proceedings were in fact finalized with his co-authors when he was in hospital.

His last public appearance was at the ER'98 conference (Nov 98) held in Singapore. Despite his weak health, he shared his vision on Information Retrieval in the WWW as a member of a panel on Internet Applications, and as a speaker in the NSF-NSTB sponsored workshop on Databases and the Internet.

Cheng Hian was a very sociable person, and often sought the company of friends. As such, he was loved by those who came in contact with him. He would often go the extra mile to help his friends and colleagues. He was a great person, and had touched the lives of many. We suddenly realized that there are many things that we will never do together again. We'll miss him sorely, and his laughter, and his smile...

Stuart Madnick, Michael Siegel (MIT, USA), Stephane Bressan, Mong Li Lee, Sin Yeung Lee, Tok Wang Ling, Beng Chin Ooi, Kian Lee Tan, and Ke Wang (NUS, Singapore), Hongjun Lu (UST, Hong Kong)

Obituary

In memory of Peter M. Stocker (1927-98)

We regret to record the passing of Peter Stocker, a UK data base pioneer and a trustee of the VLDB Endowment (1984-91). He died suddenly and unexpectedly on 25th November 1998, aged 71, and was mourned by many. He was the founding Professor of Computing at the University of East Anglia and sometime Pro-Vice Chancellor.

Peter trained at Manchester University in the early days of mainframe computers. He took part in the early IFIP Working Conferences on Data Base Management, which were very influential at the dawn of our subject. Later (1982) he organised a very successful international summer school at the University of East Anglia on "Data Bases: Role and Structure".

He is probably best known, first for his pioneering work on self-organising databases, and then on distributed databases and conceptual schema languages. He led an early UK distributed database consortium in the Proteus project. He was an active member of the VLDB Endowment and served on numerous programme committees. In particular, he was Programme Committee co-chair for VLDB'87, when it was first held in the UK.

However, he was also a genuine polymath and his able mind enabled him to make insightful comments in any technical discussion. More than this, Peter was a good friend to many; he had a wonderful impish sense of humour which enabled him to negotiate successfully the most tricky of situations. He was also an inspiring leader and seminar speaker, who had a strong influence on the development of database research in the UK.

He found time to play and follow cricket, to cook for his friends and students and to listen to music. Apart from his children and grandchildren, of whom he was fiercely proud, his main non-academic interest was his garden (and that of his daughter, Anne, in Suffolk). Here he built up a large collection of varied and unusual plants.

Contents

Author Index