

6th International Conference on

VERY LARGE DATA BASES

1980

REPRINTED FROM

Very Large Data Bases

**SIXTH INTERNATIONAL CONFERENCE
ON VERY LARGE DATA BASES**



IEEE Computer Society



Association for Computing Machinery



Canadian Information Processing Society

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limits of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through the Copyright Clearance Center, P.O. Box 765, Schenectady, NY 12301. Instructors are permitted to photocopy isolated articles for noncommercial classroom use without fee. For other copying, reprint or republication permission, write to Director, Publishing Services, IEEE, 345 E. 47th Street, New York, New York 10017. All rights reserved. Copyright © 1980 by The Institute of Electrical and Electronics Engineers, Inc. Printed in the U.S.A.

Price:
Non-members \$30.00
Members \$22.50

IEEE Catalog Number 80CH1534-7C
Library of Congress Number 79-93333
ACM Order Number 471800
Long Beach Order Number 322

Additional copies available from:

IEEE Computer Society
5855 Naples Plaza, Suite 301
Long Beach, California 90803

IEEE Service Center
445 Hoes Lane
Piscataway, N.J. 08854

ACM Order Department
P.O. Box 64145
Baltimore, MD 21264

Conference Officers

General Conference Chairman
Prof. William W. Armstrong
 Universite de Montreal

Program Committee Co-chairman
Prof. Frederick H. Lochovsky
 University of Toronto

U.S. Conference Chairman
Dr. James B. Rothnie
 Computer Corporation of America

Program Committee Co-chairman
Dr. Robert W. Taylor
 IBM San Jose Research Lab.

Steering Committee Members

Mr. Charles C. Tucker, Chairman
 ACM-SIGBDP

Prof. P. Bruce Berra, Chairman
 ACM-SIGIR

Mr. John K. Lyon, Chairman
 ACM-SIGMOD

Dr. Stuart Madnick, Chairman
 IEEE-TCDBE

Prof. Stanley Y.W. Su
 1979 VLDB U.S. Conference Chairman

Prof. Howard L. Morgan
 1979 VLDB U.S. Program Chairman

Prof. David K. Hsiao
 VLDB Co-founder

Dr. R. Stockton Gaines
 VLDB Co-founder

Dr. Vincent Lum
 Secretary

Program Committee Members

Prof. Catriel Beeri
 The Hebrew U. of Jerusalem, Israel

Prof. Gregor V. Bochmann
 U. de Montreal, Canada

Prof. Giampio Bracchi
 Politecnico di Milano, Italy

Prof. Michael L. Brodie
 U. of Maryland, U.S.A.

Prof. Georges Gardarin
 U. de Paris, France

Dr. Paula Hawthorn
 Britton Lee Inc., U.S.A.

Prof. Beverly K. Kahn
 Boston U., U.S.A.

Prof. Anthony C. Klug
 U. of Wisconsin, U.S.A.

Dr. Ted Linden
 Xerox Corp., U.S.A.

Prof. Dennis McLeod
 U. of Southern California, U.S.A.

Dr. T. William Olle
 T. William Olle Assoc. Ltd., England

Dr. Jesus A. Ramirez
 Sistemas Computacionales Avanzados, S.A., Mexico

Dr. James Rhyne
 IBM San Jose Research Lab., U.S.A.

Mr. Hirotaka Sakai
 Hitachi Institute of Technology, Japan

Prof. Joachim W. Schmidt
 U. Hamburg, West Germany

Dr. Stewart A. Schuster
 Tandem Computers, U.S.A.

Prof. Kenneth C. Sevcik
 U. of Toronto, Canada

Prof. Arne Solvberg
 U. of Trondheim, Norway

Prof. Paul G. Sorenson
 U. of Saskatchewan, Canada

Dr. Don Swartwout
 Bell Laboratories, U.S.A.

Dr. Peter D. Ting
 Bell Laboratories, U.S.A.

Prof. Frank W. Tompa
 U. of Waterloo, Canada

Dr. Wayne A. Walker
 Ontario Hydro, Canada

Ms. Beatrice Yormark
 Interactive Systems Inc., U.S.A.

Western Pacific Coordinator
Prof. Yahiko Kambayashi
 Kyoto U., Japan

European Coordinator
Prof. Hartmut Wedekind
 U. Erlangen-Nuernberg, West Germany

Latin American Coordinator
Prof. Ignacio Mijares
 Inst. Tecnol. y de Estudios de Monterrey, Mexico

Local Arrangements Chairman
Prof. Gerald Ratzler
 McGill U., Canada

Treasurer
Mr. Allan Marjerison
 Touche Ross & Company, Canada

Local Publicity
Mr. Paul-Andre Desjardins
 Ecole Natl. d'Administration Publique, Canada

VLDB Foundation Representative
Dr. Vincent Lum
 IBM San Jose Research Lab., U.S.A.

Registration
Mr. Alan J. Eastley
 Bell Canada, Canada

Travel
Prof. T. Radhakrishnan
 Concordia U., Canada

International Publicity
Prof. Tim Merrett
 McGill U., Canada

FOREWORD

Decision-making in business and government has come to rely more and more on information stored in data bases. Suppliers of hardware and software have been under great pressure to offer solutions to the often extremely complex problems related to storing, accessing, communicating, processing, organizing, and reorganizing this information; and research has indicated some basic principles which have been having an impact on commercial offerings. No one familiar with data base technology as it exists today would, however, claim that we now know most of the answers. On the contrary, alongside developments in traditional application areas, the advent of access to data bases from the home via telephone or cable, perhaps making use of international telecommunication networks, and with other goals than the pursuit of business, has enormously expanded the horizons of the field, and has underlined the importance of research in it. This is particularly true for "very large data bases", although the qualification "very large" is an elusive and time-varying one. What we may think of as "very large" today will probably appear insignificant in size in a decade, and may be storable on a single videodisc about the size of one long-playing record.

It is in this context of rapid evolution that this Sixth International Conference on Very Large Data Bases is intended to identify and encourage research, development and application of data base systems. Its objectives are the promotion of an understanding of current research, the furthering of exchanges of information and experience gained in the design, construction and use of data bases, and the provision of a forum for the discussion of future research and development. Particular attention has been directed to the data base life-cycle.

The call for papers resulted in over 150 submissions, from which 34 were selected by the program committee. We are grateful for the dedicated work of both the authors and the referees, and we hope that these proceedings will continue to bear rich fruit over the years to come.

We are fortunate to have the sponsorship of the following organizations: the Canadian Information Processing Society (CIPS), the Association for Computing Machinery special interest groups on business data processing (SIGBDP), information retrieval (SIGIR), and management of data (SIGMOD), and the IEEE Computer Society and its technical committee on data base engineering. These groups have thus illustrated their commitment to the data base field both financially and by suggesting members with organizational talent who have contributed to making this VLDB-6 a success.

Many individuals and organizations have contributed to the conference and these proceedings. We acknowledge with gratitude their support, regretting only that space and deadlines have not permitted naming them all. We would particularly like to underline the contributions of Tim Merrett of McGill University as international publicity chairman, seconded locally by P.-A. Desjardins, of Alan Eastley of Bell Canada as registration chairman, and of Gerald Ratzer, also of McGill, as local arrangements chairman. The University of Montreal Centre de Calcul deserves our thanks for use of the SOCUM program for keeping track of registrations, as does Michel Vanier who adapted the program for our particular needs. Allan Marjerison of Touche Ross and Company served diligently in the capacity of treasurer.

Heavy burdens have been carried by the coordinators: Y. Kambayashi, H. Wedekind, and I. Mijares, who solicited participation from other continents. Greatly appreciated efforts have been made by T. Radhakrishnan of Concordia University as travel coordinator, and M. Zolliker of IBM San Jose Research Laboratory. Before her sabbatical, J. Liu served as treasurer.

The increasing cost of international travel has been a problem, but with the help of several public and private organizations, the VLDB-6 travel fund has been able to assist the participants in the program to come to Montreal.

IBM Canada Limited has helped with setting up the conference secretariat, and provided coffee and refreshments during breaks.

Finally, it is a pleasure to acknowledge the fundamental role of D. K. Hsiao and V. Lum in initiating this conference and perpetuating the highly successful VLDB series.

All of the contributors to this conference can feel satisfaction not only for having promoted the ideal of free exchange of scientific knowledge, but also for having contributed to this key area of information processing: data base management, upon which our future society will depend in a most profound way.

W. W. Armstrong, General Conference Chairman

J. B. Rothnie, U. S. Conference Chairman

F. H. Lochovsky and R. W. Taylor,

Program Co-chairmen.

Participant Index

Adiba, M.E.,	Nambiar, K.K.,
Arditi, J.,	Navathe, S.B.,
Badal, D.Z.,	Peeters, E.,
Bayer, R.,	Phillips, D.,
Belford, G.,	Post, J.V.,
Bernstein, P.A.,	Raban, E.,
Blaustein, B.T.,	Reiser, A.,
Brodie, M.E.,	Risch, T.,
Cheng, W.K.,	Rosenkrantz, D.J.,
Clarke, E.M.,	Sadowski, P.J.,
Clemons, E.,	Schmidt, J.W.,
Date, C.J.,	Sedillot, S.,
Demolombe, R.,	Shneiderman, B.,
Deutsch, D.,	Shoshani, A.,
Eggers, S.J.,	Silberschatz, A.,
Elhardt, K.,	Sole, S.,
Epstein, R.,	Spewack, S.H.,
Esaulier, C.,	Spyratos, N.,
Freilich, L.,	Steel, T.B., Jr.,
Gardarin, G.,	Stonebraker, M.,
Goodman, N.,	Studer, R.,
Gotlieb, C.C.,	Swartwout, D.,
Hammer, M.M.,	Tamir, M.,
Hawthorn, P.,	Thomas, G.,
Heller, H.,	Timor, S.,
Herot, C.F.,	Treille, L.,
Honeyman, P.,	Tsichritzis, D.C.,
Hunt, H.B., III,	Vassiliou, Y.,
Jefferson, D.K.,	Walker, A.,
Kedem, Z.,	Walker, W.A.,
Kent, W.,	Walter, B.,
Le Bihan, J.,	Wilson, G.A.,
Le Lann, G.,	Wong, H.K.T.,
Landes, O.E.,	Yormark, B.,
Larson, P.,	Zdonik, S.B., Jr.,
Lindsay, B.G.,	Zukovsky, E.,
Litwin, W.,	
Menasce, D.A.,	
Missinay, T.,	
Mohan, C.,	
Mylopoulos, J.,	

The Conference Organizers wish to thank the following contributors to the Travel Fund:

The Natural Sciences and Engineering Research Council of Canada
 Université de Montréal
 McGill University
 Concordia University

Referees

The Program Committee wishes to thank the following persons who served as referees.

Aandalen, T.	Haerder, T.	Navathe, S.B.
Aho, A.	Hecht, M.	Newman, P.
Amble, T.	Held, G.	Ng, L.C.
Arisawa, H.	Horgan, B.	Obermanck, R.
Aronson, A.	Hotaka, R.	Oftedal, H.
Badal, D.Z.	Hudyma, R.	Osborn, S.
Batory, D.	Hull, R.	Pelagatti, G.
Becker, M.	Hvasshovd, S.O.	Pirotte, A.
Bell, B.	Ikeda, H.	Reimer, M.
Bernstein, P.A.	Israel, J.	Ridjanovic, D.
Blake, R.P.	Jarke, M.	Ries, D.
Blaser, A.	Jouve, M.	Risnes, O.
Borgida, A.	Kambayashi, Y.	Rissanen, J.
Brown, A.L.	Kameda, T.	Rotem, D.
Casey, R.	Kandzia, P.	Rothnie, J.B.
Ceri, S.	Kerschberg, L.	Santoro, N.
Chen, P.	Kondo, H.	Sarin, S.
Chin, A.Y.	Kongahaug, P.	Schefe, P.
Christodoulakis, S.	Kreps, P.	Schreiber, F.
Codd, E.F.	Kvitsand, S.H.	Sedillot, S.
Cohen, F.	Lacroix, M.	Selinger, P.G.
Cole, W.W.	Ladd, I.	Shu, N.
Cowan, D.	Le Lann, G.	Smith, D.
Cunha, P.	Le Viet, C.	Stiege, H.
de La Beaujardiere, J.	Lehmann, H.	Stonebraker, M.
Driscoll, J.	Lewis, D.	Svanks, M.
Epstein, R.	Lien, Y.E.	Taylor, D.
Fagin, R.	Lindsay, B.G.	Thomas, G.
Fong, E.	Locke, P.W.	Turini, F.
Freilich, L.	Lockemann, P.	Uemura, S.
Galler, B.	Lopez Suarez, A.	Vardi, M.Y.
Garcia-Molina, H.	Lorie, R.	Vassiliou, Y.
Goguen, J.	McCord, R.	Weil, W.
Gomez de Cedron, R.	Mahmoud, S.A.	Weinberger, P.
Gonnet, G.	Maibaum, T.	Weldon, J.L.
Goodman, N.	Mall, M.	Willard, D.
Graham, M.H.	Manola, F.	Wilson, J.
Grant, J.	Melli, L.F.	Wong, E.
Gray, J.	Midtlyng, J.O.	Yang, C.S.
	Mueller, G.	Zaniolo, C.
	Munz, R.	Zdonik, S.B., Jr.
		Zolliker, M.

The Conference Organizers wish to thank the following contributors to the Travel Fund

The Natural Sciences and Engineering Research Council of Canada
 Université de Montréal
 McGill University
 Concordia University

Contents

KEYNOTE SPEECH

- Some Large Questions About Very Large Data Bases* 3
C.C. Gotlieb, U. of Toronto, Canada

SESSION 1A: DATA MODELLING AND DATA ACCESS

- Chairperson: J.W. Schmidt, U. Hamburg, West Germany
- Splitting the Conceptual Schema* 10
W. Kent, IBM General Products Division, U.S.A.
- An Introduction to the Unified Database Language* 15
C.J. Date, IBM Santa Teresa Lab., U.S.A.
- Path Expressions for Complex Queries and Automatic Database Program Conversion* 33
B. Shneiderman and G. Thomas, U. of Maryland, U.S.A.

SESSION 1B: QUERY PROCESSING

- Chairperson: P.D. Ting, Bell Laboratories, U.S.A.
- On Retrieval from a Small Version of a Large Database* 47
A. Walker, Bell Laboratories, U.S.A.
- Estimation of the Number of Tuples Satisfying a Query Expressed in Predicate Calculus Language* 55
R. Demolombe, ONERA-CERT-DERI, France
- Processing Conjunctive Predicates and Queries* 64
D.J. Rosenkrantz and H.B. Hunt, III,
SUNY at Albany, U.S.A.

SESSION 2A: DESIGN AND IMPLEMENTATION OF DISTRIBUTED SYSTEMS

- Chairperson: G.V. Bochmann, U. de Montreal, Canada
- SIRIUS: A French Nationwide Project on Distributed Data Bases* 75
J. Le Bihan, C. Esculier, G. Le Lann, W. Lit,
G. Gardarin, S. Sedillot, L. Treille, INRIA, France
C. Mohan, U. of Texas at Austin, U.S.A.
- Database Snapshots* 86
M.E. Adiba and B.G. Lindsay, IBM San Jose Lab., U.S.A.
- Analysis of Distributed Data Base Processing Strategies* 92
R. Epstein, Britton Lee, Inc., U.S.A.
M. Stonebraker, U. of California, Berkeley, U.S.A.

SESSION 2B: PANEL - DATA ABSTRACTION, DATABASES AND CONCEPTUAL MODELLING

- Chairperson: M.L. Brodie, U. of Maryland, U.S.A.
- Panelists: J.W. Schmidt, U. Hamburg, West Germany
L.A. Rowe, U. of California, Berkeley, U.S.A.
J. Mylopoulos, U. of Toronto, Canada
J.M. Smith, Computer Corp. of America, U.S.A.
D. McLeod, U. of Southern California, U.S.A.

SESSION 3A: SEMANTICS AND INTEGRITY

- Chairperson: D. McLeod, U. of Southern California, U.S.A.
- A Conceptual Model for Semantic Integrity Checking* 111
G.A. Wilson, Computer Corp. of America, U.S.A.
- Fast Maintenance of Semantic Integrity Actions Using Redundant Aggregate Data* 126
P.A. Bernstein, B.T. Blaustein and E.M. Clarke,
Harvard U., U.S.A.
- Knowledge Based Query Processing* 137
M.M. Hammer and S.B. Zdonik, Jr.,
Massachusetts Institute of Technology, U.S.A.

SESSION 3B: DATA BASE LIFE CYCLE

- Chairperson: R.W. Taylor, IBM San Jose Lab., U.S.A.
- A Pragmatic Approach to Database Design* 151
S.H. Spewak, Database Design, Inc., U.S.A.
- The Development and Application of Data Base Design Tools and Methodology* 153
D.K. Jefferson,
D.W. Taylor Naval Ship R. & D. Center, U.S.A.
- Managing Systems That Manage Data* 155
L. Freilich and P.J. Sadowski, Ontario Hydro, Canada
- QUINTILLABIT: Parameters of a Hyper-large Database* 156
J.V. Post, Boeing Aerospace Company, U.S.A.

SESSION 4A: HIGH-LEVEL USER INTERFACES

- Chairperson: T. Linden, Xerox Corp., U.S.A.
- OFS: An Integrated Form Management System* 161
D.C. Tschritzis, U. of Toronto, Canada
- A Dialog Interface for Data Base Applications* 167
R. Studer, U. of Stuttgart, West Germany
- Semantics vs. Graphics: To Show or Not to Show* 183
G.A. Wilson and C.F. Herot,
Computer Corp. of America, U.S.A.

SESSION 4B: PANEL - DISTRIBUTED DATABASE SYSTEMS

- Chairperson: G. Gardarin, U. of Paris, France
- Panelists: B.G. Lindsay, IBM San Jose Lab., U.S.A.
R. Munz, Technische U. Berlin, West Germany
J. Le Bihan, INRIA, France
J.B. Rothnie, Computer Corp. of America, U.S.A.
N. Goodman, Harvard U., U.S.A.

SESSION 5A: EFFICIENT DATA ACCESS ALGORITHMS

- Chairperson: F.W. Tompa, U. of Waterloo, Canada *
- Efficient Access of Compressed Data* 205
S.J. Eggers and A. Shoshani,
Lawrence Berkeley Lab., U.S.A.
- Linear Hashing: A New Tool for Files and Tables Addressing* 212
W. Litwin, INRIA, France
- Linear Hashing with Partial Expansions* 224
P. Larson, Abo Akademi, Finland

SESSION 5B: PANEL - DATA PLANNING

- Chairperson: W.A. Walker, Ontario Hydro, Canada
- Panelists: B. MacPherson, Canadian Imperial Bank of
Commerce, Canada
J. Hartigan, General Foods, U.S.A.
L. Wu, IBM Santa Teresa Lab., U.S.A.
L.F. Melli, Ontario Hydro, Canada

SESSION 6A: RELATIONAL DEPENDENCY THEORY

- Chairperson: A.C. Klug, U. of Wisconsin, U.S.A.
- Extension Joins* 239
P. Honeyman, Princeton U., U.S.A.
- What Does Boyce-Codd Normal Form Do?* 245
P.A. Bernstein and N. Goodman, Harvard U., U.S.A.
- Functional Dependencies and Incomplete Information* 260
Y. Vassiliou, U. of Toronto, Canada

SESSION 8B: PANEL - OFFICE AUTOMATION

Chairperson: B. Yormark, Interactive Systems, Inc., U.S.A.

Panelists: E. Scott, Dept. of Transportation,
Washington, U.S.A.

E. Goodman, Exxon Communications and
Computer Science Dept., U.S.A.

B. Kaynor, Davis Polk & Wardwell, U.S.A.

D.C. Tsichritzis, U. of Toronto, Canada

SESSION 7A: CONCURRENCY CONTROL IN DISTRIBUTED SYSTEMS

Chairperson: K.C. Sevcik, U. of Toronto, Canada

Distributed Concurrency Control in Database Systems ... 275
R. Bayer, K. Elhardt, H. Heller and A. Reiser,
Technical U. Munich, West Germany

*Timestamp-based Algorithms for Concurrency Control
in Distributed Database Systems* 285
P.A. Bernstein and N. Goodman, Harvard U., U.S.A.

Update Synchronization in Distributed Databases 301
W.K. Cheng and G. Belford,
U. of Illinois at Urbana-Champaign, U.S.A.

*Non-two Phase Locking Protocols with Shared
and Exclusive Locks* 309
Z. Kedem and A. Silberschatz, U. of Texas at Dallas, U.S.A.

SESSION 7B: FORUM

Chairperson: F.H. Lochovsky, U. of Toronto, Canada

ISO Report on Concepts for Conceptual Schemas 321
T.B. Steel, Jr., A.T.&T. Co., U.S.A.
D.A. Jardine, Queen's U., Canada

*Standardization and the Relational Approach to
Databases: An ANSI Task Group Status Report* 326
M.L. Brodie, U. of Maryland
J.W. Schmidt, U. Hamburg

*EEC Evaluation and Implementation of
Database Systems* 328
E. Peeters, Commission of the European Communities,
Belgium

Telidon and the Human Factors of Videotex Data Bases 330
D. Phillips, Dept. of Communications, Ottawa, Canada

SESSION 8A: APPLICATION DEVELOPMENT

Chairperson: B.K. Kahn, Boston U., U.S.A.

DB1: A DBMS-based Application Generator 335
M. Tamir, T. Missinai, J. Arditi, E. Raban, S. Sole,
S. Timor and E. Zukovsky,
The Weizmann Institute of Science, Israel

*Production Program Generation in a Flexible
Data Dictionary System* 343
T. Risch, Datalogilaboratoriet, Sweden

*An Intuitive Approach to Normalize Network
Structured Data* 350
S.B. Navathe, U. of Florida, U.S.A.

SESSION 8B: PANEL - DATABASE THEORY AND ITS RELATIONSHIP TO PRACTICE

Chairperson: D. Swartwout, Bell Laboratories, U.S.A.

Panelists: A.G. Merten, U. of Michigan, U.S.A.

P.A. Bernstein, Harvard U., U.S.A.

A.V. Aho, Bell Laboratories, U.S.A.

SESSION 9A: DISTRIBUTED SYSTEM RELIABILITY, RECOVERY AND PERFORMANCE

Chairperson: J.A. Ramirez, Sistemas Computacionales
Avanzados, S.A., Mexico

*On the Design of a Reliable Storage Component for
Distributed Database Management Systems* 365
D.A. Menasce and O.E. Landes, Pontificia U. Catolica
de Rio de Janeiro, Brazil

*The Analysis of the Effects of Concurrency Control on
Distributed Data Base System Performance* 376
D.Z. Badal, Computer Corp. of America, U.S.A.

*Strategies for Handling Transactions in Distributed
Data Base Systems During Recovery* 384
B. Walter, U. of Stuttgart, West Germany

SESSION 9B: DATABASE MACHINES

Chairperson: P. Hawthorn, Britton Lee, Inc., U.S.A.

Panelists: R. Epstein, Britton Lee, Inc., U.S.A.

D.J. DeWitt, U. of Wisconsin, U.S.A.

E. Oliver, Bell Laboratories, U.S.A.

R. Haskin, U. of Illinois, U.S.A.

F. Bancilhon, INRIA, France

SESSION 10A: DATA MODELLING AND SCHEMA DESIGN

Chairperson: P.G. Sorenson, U. of Saskatchewan, Canada

Some Features of the TAXIS Data Model 399
J. Mylopoulos, U. of Toronto, Canada
H.K.T. Wong, Lawrence Berkeley Lab., U.S.A.

Translation Structures of Relational Views 411
N. Spyrtatos, INRIA, France

*Some Analytic Tools for the Design of Relational
Database Systems* 417
K.K. Nambiar, Bharat Electronics, Ltd., India

SESSION 10B: PANEL - DBMS STANDARDS: STATUS AND FUTURE DIRECTIONS

Chairpersons: E. Clemons, U. of Pennsylvania, U.S.A.

D. Deutsch, Natl. Bureau of Standards, U.S.A.

Panelists: F.A. Manola, Computer Corp. of America, U.S.A.

R. Gerritsen, Intl. Database Systems, U.S.A.

J.W. Schmidt, U. Hamburg, West Germany

KEYNOTE SPEECH

University of Toronto

Abstract

There are now in use many data bases, each containing billions of bytes, which have evolved with no reference to any of the topics discussed at length in data base conferences. Examples are the U.S. Census, UTLAS, PRESTO, PLATO, REGULARS, ORBIT and SIDA, to name a few. What can be done to prevent data bases from going the way of the disappearing languages, which most of the well-known models are now? Perhaps data base research should have its primary attention to those systems which are primarily data oriented, rather than those which are primarily operational systems, and which are designed to support a wide range of applications.

There are now in use many data bases, each containing billions of bytes, which have evolved with no reference to any of the topics discussed at length in data base conferences. Examples are the U.S. Census, UTLAS, PRESTO, PLATO, REGULARS, ORBIT and SIDA, to name a few. What can be done to prevent data bases from going the way of the disappearing languages, which most of the well-known models are now? Perhaps data base research should have its primary attention to those systems which are primarily data oriented, rather than those which are primarily operational systems, and which are designed to support a wide range of applications.

In putting these concerns before you I would like to start by referring to the list of Large Data Bases and Data Management Systems, shown in Table 1. One of the systems in this list is well known and widely used by its own particular community, and in this case the entries are representative of a number of similar, perhaps competing systems. If you wish to see some of these systems, called as they are, management systems, the University of Illinois and more recently the University of Toronto, is usually regarded as a data base oriented operational system. SIDA is a data base processing system which allows number oriented data base information. Nevertheless it

can be regarded as a data base system, since it is possible to convert directly from the conventional form of requests to answer queries, perform insertions, deletions, or updates, or provide listings. Sometimes the transaction may initiate a lengthy computation or tabulation job, after all, this can also happen when, in answer to a query, a join is calculated in a relational data base.

The system provide varying services to many users at scattered sites.

In some instances the data base is built around a single entity, as in UTLAS which centres on the bibliographic record, so you might prefer to regard the system as a very large file rather than a data base. But I suggest that once you have a large community of users, with different modes of access, different types of transactions, and different views of the system, the distinction between a very large file and a data base, vanishes. The above implies a rather catholic view of what a data base or data management system is, but perhaps the definition should include a somewhat larger set than that illustrated. The case argument for including SIDA could be used to argue that it is a communication network interconnecting computers and their embedded data bases, and if that is accepted, so too should Telesis and the Sharp's all network. Moreover if we do not insist that the data be alphanumeric, there are vast collections of numerical data, of weather data, and the scientific and engineering information gathered under the auspices of the Committee on Data for Science and Technology. The main point is that these systems, large as they are, in effect

The system provide varying services to many users at scattered sites.

In some instances the data base is built around a single entity, as in UTLAS which centres on the bibliographic record, so you might prefer to regard the system as a very large file rather than a data base. But I suggest that once you have a large community of users, with different modes of access, different types of transactions, and different views of the system, the distinction between a very large file and a data base, vanishes. The above implies a rather catholic view of what a data base or data management system is, but perhaps the definition should include a somewhat larger set than that illustrated. The case argument for including SIDA could be used to argue that it is a communication network interconnecting computers and their embedded data bases, and if that is accepted, so too should Telesis and the Sharp's all network. Moreover if we do not insist that the data be alphanumeric, there are vast collections of numerical data, of weather data, and the scientific and engineering information gathered under the auspices of the Committee on Data for Science and Technology. The main point is that these systems, large as they are, in effect

SOME LARGE QUESTIONS ABOUT VERY LARGE DATA BASES

C. C. Gotlieb

University of Toronto

Abstract

There are now in use many data bases, each containing billions of bytes, which have evolved with no reference to any of the topics discussed at length in data base conferences. Examples are the U.S. Census, UTLAS, PRESTEL, PLATO, MEDLARS, ORBIT and SITA, to name a few. What can be done to prevent data bases from going the way of programming languages, where most of the world continues to use Cobol and Fortran? Perhaps evolutionary models are needed; perhaps data base specialists have to pay more attention to communications. Certainly they have to do more experimental work on operational systems and establish credibility with managers of existing data bases.

Although this is the sixth international conference on very large data bases, it is the very first I have attended personally, and you might therefore wonder why I was so presumptuous as to accept the invitation to be a keynote speaker. There are three points I would make in my defence. First, my own researches, and those of my students have progressed from data structures, through file structures, to data bases, and more recently to very large data bases, so I have a personal stake in the field. Second, I have examined, rather carefully, the publications of the previous conferences, as well as those of some other data base conferences, so I do have an idea as to what the principal topics of earlier conferences have been. Finally, and most important, I have some serious concerns about data base research, and would like to make suggestions on certain directions such research might take to ensure that data base studies become more useful and gain wider acceptance.

In putting these concerns before you I would like to start by referring to the list of Large Data Bases and Data Management Systems, shown in Table 1. Each of the systems in this list is well known and widely used by its own particular community, and in many cases the entries are representatives of a number of similar, perhaps competing systems. It may surprise you to see some of these systems labelled as data management systems. PLATO, developed at the University of Illinois and marketed by Control Data Ltd., is usually regarded as a computer-assisted instructional system. SITA is a message processing system which allows member airlines to exchange information. Nevertheless it

can be regarded as a data base system because many of the carriers have their reservation computers connected directly into the network, and it is possible to present queries to and conduct transactions on these computers. The systems listed in Table 1, have the following properties to justify considering them as data bases or data management systems:

- Except for the U.S. Census, they are all on-line.
- Embedded in them are very large collections of alphanumeric files - at least 10^{10} bytes and often much more.
- They can all be viewed as collections of data along with transaction processing programs. Often the transactions take the conventional form of requests to answer queries, perform insertions, deletions, or updates, or provide listings. Sometimes the transaction may initiate a lengthy computation or tabulation but, after all, this can also happen when, to answer a query, a join is calculated in a relational data base.
- The systems provide varying services to many users at scattered sites.

In some instances the data base is built around a single entity, as is UTLAS which centres on the bibliographic record, so you might prefer to regard the system as a very large file rather than a data base. But I suggest that once you have a large community of users, with different modes of access, different types of transactions, and different views of the system, the distinction between a very large file and a data base, vanishes. The above implies a rather catholic view of what a data base or data management system is, but perhaps the definition should include an even larger set than that illustrated. The same argument for including SITA could be used to classify Tymnet as a data management system because it, too, is a communication network interconnecting computers and their embedded data bases. Once Tymnet is accepted, so too should Telenet and I. P. Sharp's APL network. Moreover if we do not insist that the data be alphanumeric, then there are vast collections of numerical data, e.g., weather data, and the scientific and engineering information gathered under the auspices of CODATA, the Committee on Data for Science and Technology. The main point is that there exists a large number of what are, in effect,

Table 1. Some Large Data Bases and Data Management Systems

System	Description	Similar Systems	Reference
U.S. Census 1970 1980	Data on individuals, grouped in blocks, census tracts 1970 basic files - 1700 reels of magnetic tape (800 b.p.i.) - approx. 2×10^9 characters 1980 approx. 10^{10} characters Plus numerous derived files, processing programs	HUD - EHAP Housing and Urban Development Experimental Housing Allowance Program	1
UTLAS University of Toronto Library Automation System	Catalogue records of library books 3,250,000 unique records (approx. 1000 bytes) plus duplicates, derived records Growth rate is 110,00 records/month	OCLC Ohio College Library Centre 5×10^6 unique records	2
PRESTEL	Two-way communication system between data base at broadcasting station and home television in the U.K. Initially (Sept. 1979) 100,000 pages		3
PLATO Programmed Logic Automatic Teaching Operations	Computer assisted instructional system Developed at U. of Illinois; marketed by CDC Over 700 lessons on automechanics, chemistry, mathematics, music, physics, political science, psychology, etc.		4
MEDLARS MEDLINE Medical Literature Analysis and Retrieval Systems	Abstracts of journal articles in medical field 2.7×10^6 citations (Aug. 1979) Increasing 220,000 per year	Chemical Abstract Service (CAS) N.Y. Times Information Bank, etc.	5
ORBIT	System Development Corporation On or off-line searching of a collection (about 30) of data bases consisting of journal abstracts, e.g., MEDLARS Operates on Tymshare, Telenet	DIALOG Lockheed Information Systems	6
SITA Société Internationale de Télécommunications Aéronautiques	Message switching network connecting 175 airline carriers Includes many on-line reservation computers	SWIFT, BANK WIRE (banking networks)	7

operational data management systems based on very large collections of specialised data.

There are some additional observations to be made about the systems described in Table 1.

Communications play an integral role in them. In fact from an engineer's point of view, it plays the central role, providing on-line access to the user. For the communication engineer the key problems relate to transaction statistics, data transmission rates, network routing, and response times as compared with those of integrity, consistency, deadlock, and data access paths, of concern to the data base designer.

Although each system is relatively well defined and narrow in its area of application, in total they display a wide range of applications, organizations and methods of providing service. In some the data is completely centralized (UTLAS, ORBIT, PRESTEL, and PLATO); in others it is highly distributed (SITA). Sometimes the transaction is a read-only (ORBIT); usually a range of operations is possible. In some, access to the data is through a single, well defined key (e.g., SITA through airline, date, flight number); in others there may be any one of a set of keys (e.g., in UTLAS where Library of Congress Catalogue No., International Standard Book No. [ISBN] and precise title are all keys, each with its own index); in still others (e.g., ORBIT) search is by a boolean construct on attributes.

All of the systems have evolved over a considerable period of time. None is less than five years old. Design as it progressed, usually focussed on communication aspects, or on record format. Problems of data organization, record linkage, data access paths and data integrity were solved, or rather avoided, by adopting very primitive solutions, e.g., single key indices, rigid query format, strict hierarchical structure, centralized control with lockout, and so on. As a consequence, the systems have evolved with little or no influence from the ideas, theories, and models that have been the subject of data base conferences.

Examination of the session titles and the proceedings of the five previous VLDB conferences shows that there has been a relatively strong emphasis on such topics as data base design, conceptual models (semantic, relational, ANSI-SPARC), language interfaces, functional dependencies, data base machines, and problems of consistency and integrity. It is true that there have also been sessions on physical design, implementation, performance evaluation, and experimental systems, but these have very seldom been concerned with existing data bases, such as those listed in Table 1. The focus of VLDB conferences, and of other data base conferences such as the one sponsored by ACM-SIGBDP, SIGIR and SIGMOD in December 1979 on the "Entity Relation Approach to System Analysis and Design", has been on theoretical models, proposals, and experimental systems, rather than on large, operational data bases.

Why is this a concern? Well, we are all familiar with what has happened in programming languages. Beginning with Algol 60, there has been a continuous stream of research, design and proposals with respect to programming languages. A great deal of it has been of excellent quality, and we have a much better understanding today of the features a language should have to make it easy-to-use, consistent and versatile, than we did when the first languages were constructed. There have been definite benefits from this theory - certainly we know how to implement a compiler on a microcomputer, efficiently and quickly. But it must also be admitted that all this work on programming languages has had relatively little impact on the vast majority of computer users who have not had formal training in computer science. Cobol and Fortran, languages which originated in the fifties, continue to be the principal language of use by those in business and science. It is true that APL has gained some currency among scientists and engineers but this is an old language too. And there is some possibility that Pascal and ADA will be adopted widely, but is also true that Basic, a very primitive language indeed, is being seriously proposed as a standard for microprocessors. It is very unlikely that the newer languages will displace Cobol and Fortran; there is simply too large a literature, and too many users for conversion to take place. At best the new languages will find a place besides the older ones which will continue to thrive.

The lesson about data bases to be drawn from this is that there is a real danger that the research carried out in universities and industrial laboratories, as described in VLDB conferences, will have very little effect on the large, operational data bases used by the world at large, and on which most money is spent. One can point out that the data bases exemplified by those in Table 1 were established before there was any theory at all about models and design, and that we can only expect to influence the design of future systems. But it is likely that the existing large data bases will continue to grow, and that practises centering around them will be representative of how data bases are built and used. Moreover it is not at all certain that current ideas on data base theory and design are finding their way into the newer data bases presently evolving. Right now many corporate data bases are being assembled with the aid of commercially available DBMS systems such as IMS and System 2000. Some of these are quite large, occupying hundreds and even thousands of megabytes of storage. But studies by a student of mine, Geovane Magalhaes, carried out on an assembly of corporate data bases at the Hydro Electric Power Commission of Ontario show that the assemblage is not being influenced by current theory. In fact, if we go closely into the DBMS systems, we are likely to find that some standard, very well known data structuring techniques, such as the use of hash tables or inverted lists for searching may not be implemented, or if they are, it will be done in the crudest way possible. I hasten to add that this is not because those in charge of developing and using the systems at Ontario Hydro are unaware of current ideas. Indeed many of them participate

actively in data base conferences. The data they are working with are files originating with different operating divisions of the organization. These historical files are a given, as are the operations to be carried out on them, the computer operating environment, and the commercial tools for data base management. Under these circumstances it is very difficult, if not impossible, to apply conceptual designs, relational models, or integrity constraints when assembling the overall database. There is ample evidence that this experience at the Hydro is not atypical of what is happening elsewhere. In general, operational priorities prevent experimentation, and historical legacies still make it difficult to apply new theoretical approaches.

So the questions referred to in the title are these: How can we ensure that the research being done on data base theory, design, and implementation will affect the major data bases being used by the world at large? What can be done to prevent data bases from going the way of programming languages, where so much work has remained outside the mainstream of computer application?

It is not easy to come up with proposals that will obviously have the desired effects, but I would like to make some general suggestions that might be useful.

We could accept the view so much is already invested in the current, large operating data bases, that most of them will never be replaced by completely new systems, designed from scratch. Instead we can concentrate on methods to modify them incrementally, eventually bring them into a form such that they may be viewed as implementations of the conceptual data bases with the properties we want. As an example, access to pages of Prestel, the British two-way information system based on television receivers in the home, is purely hierarchical. At each stage the user selects his next page from a menu of displayed alternatives. Critics have argued that the Prestel data structure is too restrictive; unless a user knows the exact number of the page he wants*, he has to traverse the tree up and down too often, and since charges are based on the number of pages displayed, the costs are too high. One can certainly think of gradually adding cross links to the Prestel data items in a way which would make the data base look like a network rather than a tree, and this could be done without rebuilding the whole data base at once. It is not certain that such an enhancement would make a decisive difference to users. There are other limitations of Prestel which also reflect the fact that the design was standardized years ago, e.g., those relating to colour, and picture resolution. The limitations due to choice of hardware may be the more important, but I would think that those due to data structure and access path could be corrected without resorting to a massive reconversion. The

*Prestel does provide users with a directory of page contents.

suggestion then, is that more attention be paid to examining how operational data bases such as those listed in Table 1, can be gradually converted to modern, more powerful ones. Going back to the case of programming languages, it is worth noting that both for Fortran and Cobol, all through the years, there have been committees who have made recommendations to bring the languages up-to-date, and this has undoubtedly been a factor in the languages' survival.

I have already noted that much of the design work for on-line data bases is focussed on the communication aspects. It is reasonably certain that communications will continue to be central in data bases, especially if the present trends to local processors connected by the ethernet or other type of network, and to distributed systems, continue. The VLDB programs have, of course, had sessions on distributed processing, but it is true to say that most of the people working on data base design, theory, and implementation, or at least those coming from our university computer science programs, do not know enough about the hardware and software of communications. This will have to change. In my view the separation between data bases and communications is not tenable in the long run, and only those who are well versed in both disciplines will be able to make effective contributions. For example, a discussion of security and integrity which omits any reference to communication protocols is likely to overlook some of the principal problems.

In my opinion the most urgent work that has to be in the data bases is in the area of experimentation. There have, of course, been many experimental versions of different models, and these have been described at previous VLDB conferences. These have also been sessions on implementation, practice, and application development, and more such sessions are scheduled later in this conference. But now that there are many large data bases in actual existence, each with hundreds or even thousands of users on-line every day, there is the possibility of learning much more about workload profiles, storage references, and resource usage than has hitherto been known. Obviously the frequency distribution of transactions according to type depends on how queries are formulated in the user-language. But it is difficult to design a user-language without some a priori ideas on what kinds of queries are important. Similarly the logical or conceptual design of a data base must be easily interpreted in terms of the actual data items users reference frequently, and the relationship between these items. In other words the real test of abstract, conceptual models is whether it is possible to efficiently and naturally map the operations in them into things users want to do. To apply this test we must know what users really do with data bases. Only then will the models be useful for design, implementation and operation, and serve as effective working tools.

To my knowledge there is only one instance described in the open literature of experimental data gathering from an operational data base while in service. The methodology for this experiment

is described by Tuel and Rodriguez-Rosell⁸, the main findings are given in Ragaz and Rodriguez-Rosell⁹ and in Rodriguez-Rosell¹⁰, and some of the conclusions that can be drawn from the data are reported by Smith¹¹. On reading about the complexity of the experiments carried out by Rodriguez-Rosell and his colleagues it is easy to see why more such work has not been done. After capturing the data, in itself a non-trivial task, a large IBM computer was dedicated for an appreciable time to reconstructing the data base and its operating system, and to analyzing the results obtained. When Geovane Magalhaes and I attempted to do similar work on other systems we found the same difficulties. Although there is a very large IMS data base in our university on academic records, financial systems, and personnel, we found it impossible to conduct experiments in this environment. We might have been able to meet the concerns of those in charge of the system about access to confidential records (there are some legal issues here), and gathering data in a way which would guarantee no interference with operations. But the equipment resources were simply not available, because taking the observations needed a monitoring computer of appreciable power. Such resources were available at the Hydro, where we were able to carry out the kind of experiments we wanted to, and we are indebted to Dr. Walker and his colleagues for giving us the opportunity to do so.

I think that the position with regard to experimental work on operational data bases is very similar to the position of experimental work on operating systems ten years ago. At that time the importance of collecting hard data on real systems was recognized. It was difficult to collect such data, however. Initially complete computers were dedicated as monitors. Later hardware and software monitors were developed, and more recently methods of collecting the data are built into the computer hardware or into the operating system itself. Before this could be done it was necessary to learn what to observe. Certainly we need better, i.e., less expensive tools for taking measurements on data bases. Many systems do provide memory for recording statistical counts, but I do not think that any system includes a simple method of capturing the storing of memory references encountered in processing a query, or "photographing" the indices, directories and data dictionaries at the start of a run. Such interpreter-like capabilities would be useful.

It is likely that experimental observations on data bases will be harder and more expensive to get than those on operating systems. Just copying a data base can cost thousands of dollars. But the history of experimental observations in performance evaluation should be an encouragement to persevere with experimental work on data bases. There the observations soon yielded some immediate results because it was possible to observe device contention, resource imbalance and take corrective actions. In the longer term the existence of measurements has contributed to the development of

the theory of queueing networks of processors, a theory which has already made a significant difference in our understanding of how to manage a complex computer system, and one which will undoubtedly become part of the main framework of computer science. It may well be that as experimental data on data bases become available, we will be able to narrow down the range of conceptual models we are now confronted with, and concentrate on one or two which tie in well with implementations and observations, so that they are really useful. In the short term observations should also allow us to tune and improve existing systems, with attendant improvement in performance. This may give those of us doing research work on data bases greater credibility with those responsible for operating them, something I am not sure we always have.

In summary, I believe that it is time for a shift of emphasis in the direction of research on very large data bases. The shift would involve working more closely with existing large data management systems, paying more attention to communication aspects, and developing tools and techniques for collecting and analysing experimental data on operational systems. It does not imply downgrading of the value of conceptual models and theoretical research. But the opportunities are now here to look for ways to tie theory in with practise, and until this is achieved Data Bases, as a discipline, will not mature.

References

1. "Census Summary Tapes: What's New for '80", Data User News, Vol. 15, no. 3, March 1980. Department of Commerce, Bureau of the Census, Washington, D.C.
2. Matthews, J.R., "The Four Online Bibliographic Utilities: A Comparison", Library Technology Reports, American Library Assn., Nov-Dec, 1979.
3. The PRESTEL Users' Guide and Directory, Vol. 2, no. 3, July 1979.
4. PLATO Courses: 1980 Catalogue, Control Data Education Co., 1980.
5. National Library of Medicine, 8600 Rockville Pike, Bethesda, Maryland 20014, U.S.A.
6. System Development Corporation, 2500 Colorado Ave., Santa Monica, CA 90406, U.S.A.
7. Schwartz, M. Data Communication Network Design and Analysis, Ch. 2, Prentice-Hall, 1977.
8. Tuel, W., Jr., and J. Rodriguez-Rosell, A Methodology for Evaluation of Data Base Systems, IBM Research Report RJ 1668, San Jose, CA, October 1975.
9. Ragaz, N. and J. Rodriguez-Rosell, Empirical Studies of Storage Management in a Data-Base System, IBM Research Report RJ 1834, San Jose, CA, October 1976.
10. Rodriguez-Rosell, J., "Empirical Data Reference Behaviour in Data Base Systems", The Computer Journal, Nov. 1976, pp. 9-13.
11. Smith, A.J., "Sequentiality and Prefetching in Database Systems", ACM Transactions on Database Systems, Vol. 3, no. 3, Sept. 1978, pp. 223-47.

