

SIGMOD RECORD

A Quarterly Publication of SIGMOD
the Special Interest Group on
Management of Data

Volume 15, Number 3

September, 1986

CONTENTS	EDITOR'S NOTES	1
	CHAIRMAN'S MESSAGE	2
	ARTICLES	
	Data Base Directions: Information Resource Management— Making It Work, Executive Summary	3
	Elizabeth N. Fong and Alan H. Goldfine	
	A Logic-Programming/Object-Oriented Cocktail	11
	Francois Bancilhon	
	An Event/Trigger Mechanism to Enforce Complex Consistency Constraints in Design Databases	22
	Klaus R. Dittrich, Angelika M. Kotz, and Jutta A. Mülle	
	A Survey of Distributed Deadlock Detection Algorithms	37
	Ahmed K. Elmagarmid	
	The Proteus Bibliography: Representation and Interactive Display in Databases	46
	T. Lougenia Anderson, Earl F. Ecklund, Jr. and David Maier	
	Database Systems Group at Maryland	56
	Database Research Group at Monash University	61
	CALLS FOR PAPERS	63



SIGMOD RECORD

SIGMOD RECORD is a quarterly publication of the Special Interest Group on Management of Data (SIGMOD) of the Association for Computing Machinery (ACM).

We solicit contributions of articles, news, reports, book reviews, and letters. Conference announcements and calls for papers are published if relevant to the interests of the group. All contributions are to be sent to the editor for consideration for publication. Materials are not normally returned. Technical papers submitted to the *RECORD* should be relatively short (5-15 single space pages) and sent to the editor either in camera ready format or

via electronic mail. Papers which are sent in camera ready form should be single spaced on 8.5 by 11 inch paper with at least 1 inch margins.

Opinions expressed in signed articles and letters are those of the writer and do not necessarily express the position of the ACM or SIGMOD. Unless otherwise stated, reports and technical papers in the *RECORD* are unrefereed working papers, and the author(s) should be contacted for reprint authorization. Contact the editor for information on the copyright status of all uncredited material.

SIGMOD Officers

Chairperson

David DeWitt
Computer Sciences Dept.
University of Wisconsin
1210 W. Dayton St.
Madison, WI 53706

Vice-Chair

Paula Hawthorn
Britton-Lee Inc.
Suite 304
1919 Addison St.
Berkeley, CA 94708

Secretary/Treasurer

Fred Maryanski
University of Connecticut
EECS Dept, U-57
260 Glenbrook Road
Storrs, CT 06268

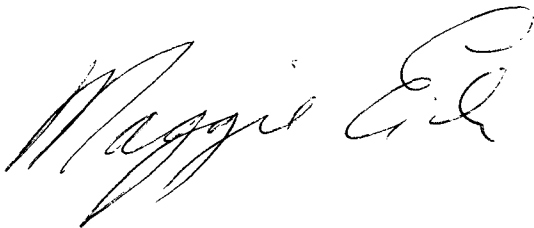
Record Editor

Margaret H. Eich
Department of Computer Science
and Engineering
Southern Methodist University
Dallas, TX 75275
CSNET: eich@smu

Editor's Notes

With this issue of the *SIGMOD RECORD* we should be back on our quarterly schedule for publications: March, June, September, and December. As you can see from the contents of this issue, my plea for contributions to the newsletter has been relatively successful. I have been a little disappointed, however, by the response from students and industry personnel. If you are a student and will soon be looking for a job, any articles printed in the RECORD may help you get that needed exposure. If you work in industry, please don't feel that only research articles will be printed. A large percentage of our membership works in industry and would probably be interested in any recent ideas, new techniques, or experiences that you may have had in the database area. Remember, the need for input is constant. Consider the RECORD as the method to disseminate any relevant database information to a wide audience. I would like to see more technical papers and research overviews. If you have any questions concerning SIGMOD please feel free to send in letters. I will publish all pertinent letters with responses.

The first article in this issue is the Executive Summary of a workshop on "Information Resource Management - Making It Work" held at Ft. Lauderdale, Florida on October 21-23, 1985. This workshop was the fourth in a series of workshops sponsored by the National Bureau of Standards on "Data Base Directions". The detailed proceedings of this workshop can be found in the June 1986 (Vol. 9, No. 2) issue of *Database Engineering* or in NBS publication 139 *Data Base Directions: Information Resource Management - Making It Work*. Due to scheduling problems we were unable to publish a joint issue of *SIGMOD RECORD* and *Database Engineering* covering this workshop as had originally been hoped. We, however, look forward to future joint issues with the *Database Engineering* newsletter.

A handwritten signature in cursive script, appearing to read "Maggie Eil". The signature is written in dark ink and is positioned below the main body of text.

Chairman's Message

The wind is blowing 30 MPH and I am dying to leave and go windsurfing but Maggie says I have to write a Chairman's statement first - even though I have nothing to say. So here it is.

As I discussed in the previous regular newsletter, we now seem to have the newsletter back on track. We do continue to need contributions for upcoming issues so please be sure that you send relevant papers to Maggie. Since printing the SIGMOD proceedings as a newsletter is very expensive and our budget is precariously low, we may only be able to afford to print 3 newsletters this year. Next year, as all members start to pay at the new membership rate, we (hopefully) will have our budget problems behind us. We may, however, have to raise dues again.

The recent SIGMOD conference saw a very strong turnout. The program and local arrangements were excellent and I thank Nick Roussopoulos and Carlo Zaniolo for all their efforts in making the conference a success. The 1987 SIGMOD conference is being organized by Mike Stonebraker. The conference will be held May 27-29 in San Francisco and you should be receiving a call for papers in the very near future. The program co-chairmen are Irv Traiger from IBM-San Jose and Umesh Dayal at CCA. The 1987 PODS Conference will be held March 23-25 in San Diego. The general chair is Ashok Chandra of IBM-Yorktown and the program chairman is Moshe Verdi from IBM-San Jose. The 1988 SIGMOD conference is tentatively scheduled to be held in Chicago. Finally, proceedings from the upcoming Object-Oriented Database Systems Workshop which is being held September 23-26 at Asilomar, will be for sale from ACM. Since attendance is limited and the workshop looks to be an exciting one, the proceedings will be especially valuable to those that cannot attend.

As a last issue, there may be some hint of progress on the issue of SIGMOD cosponsoring the VLDB conference. By the time of the next issue of the SIGMOD Record, the issue should be either resolved one way or the other. I will keep you posted.



Data Base Directions Information Resource Management - Making It Work Executive Summary

Elizabeth N. Fong and Alan H. Goldfine
National Bureau of Standards

On October 21-23, 1985, the Institute for Computer Sciences and Technology of the National Bureau of Standards (NBS), in cooperation with the Association for Computing Machinery Special Interest Group on Management of Data (ACM SIGMOD), the IEEE Computer Society Technical Committee on Database Engineering, and the Federal Data Management Users Group (FEDMUG), held the fourth in their series of Data Base Directions workshops. The purpose of this workshop was to assess the nature of current information resource management (IRM) practice and problems, and to explore solutions which have proven workable.

Since the last Data Base Directions workshop five years ago, IRM has been defined, introduced, and applied by many organizations. It is time to evaluate current practice, to identify problem areas, to review what technologies and tools are important and when to apply them to IRM, and to explore the motivations and inhibitors to decentralized and distributed environments.

The workshop was organized into four working panels, which met to discuss:

- o IRM, MIS and the Organization in the 1990s
- o IRM and the Systems Life Cycle
- o Technologies for IRM
- o IRM in a Decentralized and Distributed Environment.

The keynote speaker, Eugene Bloch, Director of Corporate Information Systems and Services for Allied Signal, Inc. spoke from the point of view of a practitioner. As an MIS manager coming from the "real" world of systems developments, operations, budgets, demanding users, and application backlogs, Block claimed that current IRM methods are inadequate. He identified the barriers to making IRM work as:

- o MIS lacks credibility
- o The organizational culture is not ready for IRM

- o Technology is changing too rapidly
- o Organizations believe that it is too costly to replace the large investment in old systems.

He claimed that these barriers must be removed, and challenged the workshop participants to develop answers and approaches.

IRM in the 1990s

This panel was charged to determine the economic, political, and technical trends that would shape the IRM functions and organizations during the next decade. In addressing the evolution of the IRM process, the panel used a conceptual model consisting of three levels: the value system, the process structure, and the technical structure.

In identifying the demand for information in an organization, the panel decided on a new approach to IRM called "Information Asset Management (IAM)." An asset was any resource that was not consumed through use, i.e., any resource that was specifically developed for the purpose of being leveraged or reused in the creation of products or services.

A detailed analysis of IAM took place. The assets were structured into five categories: assets required in data acquisition, data storage, data manipulation, data retrieval, and data distribution. Four management functions were defined as: planning, organization, administration, and control. Each of the asset requirements was analyzed with respect to the four management functions.

The panel concluded that enterprise management will have to modify its thinking to deal with information as an enterprise-wide asset, as opposed to a departmental expense. The predictions of IRM in the 1990s include:

- o The nature of applications will change
- o The traditional systems development life cycle is obsolete
- o Users will be a free market
- o Information Services as we know it may go away.

The panel finally identified the primary inhibitors to the evolution of the asset management concept of IRM as:

- o The current structure of information services
- o The current software legacy
- o The current concepts of IRM financing.

IRM and the System Life Cycle

This panel addressed four issues:

- o IRM and the Organization
- o The Management of Change
- o Metadata to Support IRM
- o Methodologies, Tools and Techniques

The working group on IRM and the Organization explored how to help an organization successfully implement IRM. Among the suggested techniques were:

- o Analyzing the "readiness" for IRM
- o Clarifying the organization's objectives
- o Planning carefully, using the Strategic Information System Planning (SISP) method
- o Carefully selecting a pilot project

The working group on Management of Change identified the forces of change as coming from the business environment and technology advances. The identified strategies for dealing with change included:

- o Making change a constant

- o Breaking long range plans into short-term projects
- o Subdividing information architecture
- o Developing flexible methods
- o Reducing the size of changes.

The management of metadata is an important aspect of IRM since metadata describes the information resource and indicates how that data is collected. The scope of metadata discussed by this working group encompassed the total repository of data describing the development and operation of applications. The metadata should support information system planning, database design, data and process creation, maintenance, control and distribution, plus other issues such as security, integrity, reliability and project management. How to effectively represent and manage metadata is still an open question.

The objective of the Methodologies, Tools, and Techniques working group was to determine what tools are required to support IRM throughout an organization's system life-cycle. A list of generic techniques was prepared, and each technique was analyzed for its relevancy to the different stages of the system life-cycle. Future IRM-compliant system life-cycle methodologies will employ tools that cover all stages of the system life-cycle, provide a choice of techniques, and be computer-aided.

Technologies for IRM

This panel was charged with reviewing the current technologies and tools that are important to IRM. The panel first classified the technologies into the following categories:

- o Application development methodologies and supporting tools
- o Information resource dictionary systems (IRDSs)
- o Database management systems (DBMSs)
- o Application generation/development systems

- o Fourth Generation inquiry and report languages for end-users (4GLs)
- o Systems for new application areas, PC/intelligent workstations, local area networks, and database servers
- o Heterogeneous database management.

Each of the technologies was evaluated along the dimensions of: state of the art, benefits and pitfalls in its use, and short-term and long-term outlook. The results of the evaluation are summarized here:

1. Application development methodologies and supporting tools. While application development methodologies have been used for several years, existing methodologies are not sufficiently complete. Most methods center around defining requirements and designing application software and databases. These tools are not widely used because they are not integrated with other tools, and because the cost of training is high. In the short-term, application development methodologies will continue to be uncoordinated and support only specific applications. In the long-term, applications will be created at the conceptual level and automated code generators will generate the application version needed for a particular operating and processing environment.
2. Information resource dictionary systems (IRDSs). The concept of an IRDS has progressed since the last Data Base Directions workshop. The IRDS standard is expected to be accepted by both the private industry and Federal communities. The functionality of the IRDS has expanded to include the support of life-cycle phases, external interfaces, distributed databases, etc. The panel predicted that, in the long term, the IRDS will have enhanced model management to better handle, for example, graphics/image, voice, and non-technical data types; the IRDS will be more integrated with external software and personal computers related to local area networks; the IRDS will support the development of a standard database access and query language; and the standard IRDS will have modules to provide enhanced support to other information processing related technologies.

3. Database management systems. DBMSs are a must for IRM. The state of the art of DBMS has matured, and "pseudo-relational" products dominate the micro DBMS market. The group agreed that, at least for the Government sector, standardization, is necessary. The phenomenal increase in the use of DBMSs has had many positive effects on organizations, including improved data shareability, improved data integrity, improved programmer productivity, and improved availability and recoverability of systems. However, there are also negative effects, including unrealistically high expectations that installing a DBMS will magically cure all problems of data management, and the frequent increased resource utilization of DBMSs. In the short-term, will have improved interfaces to the user, to the IRDS, and to applications such as graphics. The long-term will see DBMSs with richer underlying models, support of rule-based systems, and improved utilization of resources.
4. Application generation/development systems. Although application generators are emerging rapidly, (almost every DBMS vendor now offers some kind of application generator), the technical quality of these tools is still immature. These tools are not integrated with existing systems, and give little assistance during analysis phases. However, the code produced by application generators is generally of high quality and very portable. There is no immediate need to standardize the tool, but there is a need to ensure that the code produced by the generators does conform to standards. The use of application generators has not realized its full potential because of cultural barriers and their ineffectiveness in certain application areas such as real time process control systems.
5. Fourth generation languages (4GLs). It is not clear where the boundaries between 4GLs, DBMSs and application generators are. The assessment of the state of the art revealed that the technical quality of 4GLs is uneven and often suffers from severe run-time performance problems. There is a need for core standards so that some primitives can be defined across the board, but the immediate need is for a set of standards/guidelines specifying when and when not to use the 4GLs.
6. New application areas, PC/intelligent workstations, LANs and database servers. A list of requirements for database support for new applications areas was developed. In many cases, the new areas identified

are pushing the frontiers of computer technology. Since these new applications are still evolving, it was felt to be too soon to introduce standards.

7. Heterogeneous database management. The issue here is the storing of data in a heterogeneous environment containing a proliferation of databases and diverse data models and DBMSs. The approach currently taken is either to have locally autonomous databases, or to construct a layer of software to integrate databases by merging the data. No commercial systems are available to deal with the data stored in a heterogeneous environment. Several experimental projects are underway in industry and universities using a multi-layer approach. There is little user experience, and the difficult problems dealing with update are not fully understood. The future calls for research regarding dictionary placement, and distribution of schema information.

IRM in a Decentralized and Distributed Environment

This panel discussed IRM in the context of migrating to a distributed environment.

Using the concept of "spheres of control", the group described several ways in which data can be controlled and shared in a distributed environment. These methods involved local data, interchange data, and shared data with various levels of distributed database management support.

An organization's migration path toward a distributed environment is determined by its starting point. One such starting point is a centralized environment. The critical factor here is the level of sophistication and understanding of IRM. If an organization does not have effective control of its data resources in a centralized environment, it will have that much more difficulty trying to migrate.

The other starting point is a decentralized environment that includes multiple computer sites existing with virtually no communications between them. This means that the different sites will have different levels of sophistication in managing their information resources.

A transition plan to distributed database management and IRM was discussed. The discussions centered around technical and administrative issues, based upon the two starting points. The conclusion reached was that it is more difficult to migrate to a distributed environment if the

starting point is a decentralized environment, because there are many sites trying to maintain control and possibly having different perceptions of where distributed IRM should be going.

Finally, in assessing the state of the art in distributed DBMS systems, the group felt that the distributed systems offered by vendors have specific limitations, but they are clearly steps in the right directions.