Database Research: Achievements and Opportunities Into the 21st Century

Avi Silberschatz, Mike Stonebraker, Jeff Ullman
editors

1 Introduction

In February of 1990, a group of database researchers met to examine the prospects for future database research efforts. The resulting report (Silberschatz et al. [1990]) was instrumental in bringing to the public's attention both the significance of prior database research and the number of challenging and important problems that lay ahead. We shall not repeat here the major points made in that report concerning the historical development of relational database systems and transaction management. Rather the reader is encouraged to consult either that report or an on-line document (Gray [1995]), each of which discusses these and other historical achievements of database research. In May of 1995, a second workshop was convened to consider anew the prospects for database research, and this paper reports on our findings.[1] The following points summarize the conclusions of this forum.

- The database research community plays a foundational role in creating the technological infrastructure from which database advancements evolve.
- Next-generation database applications enabled by the explosion of digitized information over the last five years will require the solution to significant new research problems. These problems are grouped in this report into the following broad areas: support for multimedia objects, distribution of information, new database applications, workflow and transaction management, and ease of database management and use.
- A new research mandate for the database community is provided by the technology developments of the recent past — the explosions in hardware capability, hardware capacity, and communication (including the internet or "web" and mobile communication).
- There is a heightened need for governmental and industrial support of basic database research in order to respond to these challenges.

In the remainder of this section, we briefly introduce two major themes that weave through this report. First, the demands placed on database systems are drastically changing, thereby requiring new solutions to fit the altered landscape. Second, the database research community has a long and successful track record creating new solutions and enabling the technology transfer to put these ideas to practical use. Therefore, it is a wise investment of public research monies to ensure that this research community remains healthy and viable.

1.1 The Changing World of Database Management

At its core, a database system is a computerized record-keeping system; it stores and provides access to information. Reduced to its basic components, a database system consists of data, hardware, and software. Although these simple components remain fundamental to database systems, their scope, magnitude, and complexity have expanded mightily.

The last five years have witnessed an unparalleled movement toward data of increasing complexity. The simple business-data-processing information expressed in numbers and character strings, while still important, has been joined by large numbers of multimedia "documents," images, time-series, procedural or "active" data, and myriad other complex data forms. Representative examples of expanded data sources are found in Sections 2.2, 4.4, 4.2, 3, 5.1, and 5.2.

Additionally, low-cost, high-speed hardware components such as multiprocessors based on fast and inexpensive microprocessors, have become widely available. "Off-the-shelf" or commodity disks and

[1] This workshop was supported by the National Science Foundation, Database and Expert Systems Program, under the grant IRI-9521026. All opinions, findings, conclusions and recommendations in any material resulting from this workshop are those of the workshop participants, and do not necessarily reflect the views of the National Science Foundation. The attendees at the workshop were: Phil Bernstein (Microsoft), Ron Brachman (ATT Bell Labs), Mike Carey (IBM), Rick Cattell (SUN), Hector Garcia-Molina (Stanford), Laura Haas (IBM), Dave Lomet (Microsoft), Dave Maier (Oregon Graduate Inst.), Jeff Naughton (U. Wisconsin), Michael Schwartz (U. Colorado), Pat Selinger (IBM), Avi Silberschatz (ATT Bell Labs), Mike Stonebraker (Berkeley), Jeff Ullman (Stanford), Patrick Valduriez (INRIA), Moshe Vardi (Rice), Jennifer Widom (Stanford), Gio Wiederhold (Stanford), Marianne Winslett (U. Illinois), and Maria Zemankova (NSF). Comments on this document were also contributed by Jim Gray (Microsoft).

memories show increased capacity and reduced cost each year. The impact of hardware advances is reviewed in Section 4.1.

Lastly, a new breed of sophisticated DBMSs has begun to appear to manage the demands of this new collection of data and to fully exploit the processing advantages of new hardware. These systems are explored in Section 2.

Not only is the computing infrastructure changing, but also the user community is undergoing a similar revolution. Nearly every human enterprise now includes computerized information processing as an integral component of its operation. From the free-ranging connectivity of the World-Wide-Web (Section 4.4) to the truly astronomical proportions of the Earth Observing System (Section 3.1), the world is on-line and exchanging information. User expansion is illustrated in Sections 3, 4.4, and 5.3.

1.2 The Case for DBMS Research

A primary goal of this report is to provide a sound and reasoned argument that database systems are pivotal to current and developing information-management needs and that, consequently, a commitment to the financial support of database research is a worthwhile investment. A second, major goal of this report is to build on the 1990 report in substantiating the payoff from funding basic database research. The history of database systems research clearly illustrates the connection between basic research, commercial success, and job creation. We observe a consistent pattern of theory building and derivation of working principles, leading to experimental studies and prototype implementations, which evolve, in turn, into commercial products.

We begin, in Section 2, with examination of several areas where research, often dating back a decade or more, began to influence products in the first half of the 1990s. This section demonstrates the continued payoff to research funding in this area. Then, Section 3 continues with a sampling of the new database applications that will drive future research. We continue in Section 4 with the context in which this required research must take place. Lastly, Section 5 lays out the research agenda which the workshop participants felt strongly was required. Section 6 closes this report with some final thoughts.

2 New Research Achievements

A number of new developments in the database marketplace took place since our 1990 report. As was the case prior to 1990, the great majority of the companies active in these new markets are US-owned companies, and the products these companies provide have their origins in research and prototypes funded primarily by US government agencies. In this section we briefly outline some of the key new developments.

2.1 Object-Oriented and Object-Relational Database Systems

In 1990 there existed a few research prototypes of object-oriented database systems (OODBs). There was considerable debate regarding the nature of such systems and the relationship between OODBs and relational systems. Today there are a variety of commercial OODBs. In aggregate, it is a $75M/year market growing at about 50% per year.

Likewise, in 1990 there were a few research prototypes that combined the good features of relational DBMS (SQL access to simple data types) with the good features of OODBs (modeling of complex data) to create new breeds of "object-relational" database systems (ORDBs) and "deductive-object-oriented" database (DOOD) systems. Today, technology transfer of these ideas and prototypes has created new markets. Moreover, the traditional relational vendors are beginning to move their products in this direction.

2.2 Support for New Data Types

It was recognized in the 1990 report that new applications may require data that is not necessarily composed of records of simple numbers and character strings. Since then, research on a number of other forms of information has found its way into products.

- A decade or more of work on storage and retrieval of spatial data has now appeared in commercial GIS (geographic information systems).
- Long-term exploration for appropriate models of temporal data has now resulted in a number of different proposals for extending query languages to better support temporal data. Among these, the one which currently has the widest support is the TSQL2 proposal, which extends SQL-92.

2.3 Transaction Processing

A core responsibility of a database management system is to support the coordination of many simultaneous users of shared information. The clear definition of transaction management requirements was one of the key contributions of the DBMS research community during the 1970's and early 1980's, as noted in the 1990 report. However, traditional transaction management is not always the appropriate requirement for today's distributed information systems. The study of policies and algorithms for supporting alternatives to atomic transactions dates back to the 1970s. Today, we see that this long research chain is beginning to bear fruit, such as:

- Some commercial products support replicated data, allowing consistent views of the same information to exist at various nodes of a network.
- Some object-oriented systems now support "long transactions," where data is "checked out" and not restored for hours or days.
- Some commercial database systems also support "versions and configurations," the ability to construct a history of related objects ("versions" of that object) and to combine specific versions of different objects into "configurations." These capabilities, as well as "long transactions," are important support for design activities such as software engineering.

3 New Database Applications

To understand our proposed agenda for database research, we illustrate the next-generation applications that drive requirements for new capabilities. In this section, we briefly discuss five such applications.

3.1 EOSDIS

The Earth Observing System (EOS) is a collection of satellites to be launched by NASA starting in 1998; their purpose is to gather information that will support earth scientists concerned with long-term trends regarding our atmosphere, oceans, and land. These satellites are to return to earth about 1/3 of a petabyte (10^{15} bytes) of information per year.

This data is intended to be integrated with existing data and information from other sources such as foreign satellites or nonsatellite observations, and will be stored in the EOSDIS (EOS Data and Information System), a database on a scale not heretofore seen.

EOSDIS is intended to supply the information needs of both scientists and nonscientists. For example, it is imagined that school children will be able to access EOSDIS information to see simulations of world weather patterns, the effect of vulcanism, and so on. Among the many challenges presented by the EOSDIS project are:

- Providing on-line access to petabyte-sized databases and managing tertiary storage effectively.
- Supporting thousands of information consumers with very heavy volume of information requests, including ad-hoc requests and standing orders for daily updates.
- Providing effective mechanisms for browsing and searching for the desired data,

3.2 Electronic Commerce

There are a number of active projects attempting to make information available to support on-line browsing of catalogs and subsequent electronic purchasing of goods. The general goal is to allow companies to supply information about their products to on-line customers (perhaps using an electronic broker as an intermediary). Brokers can aggregate data from several sources, for example by collecting all electronic catalogs on wearing apparel. In turn, a broker can offer "one stop shopping" to the ultimate customer.

Like EOSDIS, electronic commerce involves a very large number of participants interacting over a network. While EOSDIS involves one principal supplier and many consumers, procurement involves many suppliers and many consumers. Further, the participants are mutually suspicious and often have already installed proprietary information systems. Among the challenges of this environment are:

- Heterogeneous information sources must be integrated. For example, something called a "connector" in one catalog may not be a "connector" in a different catalog. Such "schema integration" is a well-known and extremely difficult problem.
- Electronic commerce needs reliable, distributed authentication and funds transfer.

3.3 Health-Care Information Systems

Physicians need to draw upon many different kinds of information in the course of their work. For example, medical records about a single patient may exist at many different hospitals, medical offices, and insurance offices. A history of the patient must be gathered from these. Information about procedures, drugs, diagnostic tools, and other aids to treating patients are in principle available through multiple systems and databases.

In addition, the notes of physicians on a patient, records of diagnostic tests, and insurance and billing information, can be captured electronically and made available for later use. Transforming the health-care industry to take advantage of what is now possible will have a major impact on costs, and possibly on quality and ubiquity of care as well. But there are problems to be solved before we can reach that point, such as:

- Integration of heterogeneous forms of legacy information.
- Access control to preserve the confidentiality of medical records.
- Interfaces to information that are appropriate for use by all health-care professionals.

3.4 Digital Publishing

The publishing industry, like health-care, will undergo profound changes in business practices over the next few years. It is becoming possible to store books and articles electronically and deliver them over high-speed networks to consumers. Further, the notion of a publishable document is broadening to include audio and video, graphic images, lectures, annotations, and other elements that convey information. The total amount of information available to be put on line dwarfs the petabyte-sized EOSDIS database, and the amount that is likely to become available in the near future is an order of magnitude greater.

A corollary of the changes to come in publishing is that the education industry draws much closer to publishing. Instead of live lectures to small groups of students, educational products that are part text, part video-lecture, and part interactive training can serve much of the needs of thousands of students. These prospects motivate the following research directions:

- Management and delivery of extremely large bodies of data at very high rates. Typical data consists of very large objects — in the megabyte to gigabyte range — and may require delivery with real-time constraints.
- Protection of intellectual property, including cost-effective collection of small payments and inhibitions against reselling of information.
- Organization of and access to overwhelming amounts of information.

3.5 Collaborative Design

Large projects such as aircraft are now being designed and built by consortia of independent companies. Information regarding such projects often has a multidecade lifetime, supporting both maintenance and subsequent modifications. Designs may be simulated on a computer — for performance, feasibility of assembly, and correctness — before becoming physical reality. Designs evolve both before and after first manufacture, causing a proliferation of current, tentative, and historical configurations of related information. Different design disciplines tend to use distinct design tools with different underlying models and notations. Moreover, designs often last longer than the tools that produced them, so different components of one design may end up having been developed by different versions of the same tool. We are thus led to consider problems such as:

- As with several of the other paradigm problems suggested above, we must deal with integration of varied sources of information, including legacy sources.
- Collaborative design requires support for new forms of concurrency control and sharing mechanisms in the underlying databases.
- Interacting collections of processes, such as design and simulation, require "workflow" management in which long-term transactions interact in sound ways.
- Support for versions of the design for a single component and support for configurations combining versions of many components is essential; see Section 5.3.3.

4 Trends That Affect Database Research

Before presenting our agenda for database research, we present some observations about trends that impact this research. In this section, we discuss in turn hardware trends, the directions of the DBMS vendors, research and the business climate, and the World-Wide-Web.

4.1 Technology Trends

For fifty years there has been an exponential improvement in a number of parameters that measure our ability to compute. Each of the following have improved by a factor of 10 or more every ten years.

1. The number of machine instructions executable in a second.
2. The cost of a typical processor.
3. The amount of secondary storage per unit cost.
4. The amount of main memory per unit cost.

These improvements in price/performance of critical components make it possible, each few years, to provide solutions to new classes of problems and to create new kinds of products and services that were previously beyond our reach.

We expect these trends to continue unabated into the next millenium. Moreover, in the last few years, these trends have been joined on their upward spiral by two new parameters that heretofore had improved, but not as explosively as the others, namely:

- The number of bits transmitted per unit cost.
- The number of bits transmitted in a second.

These new trends result in an environment where it is possible to deal with terabytes of data and complex queries in a cost effective manner.

4.2 Database Architectural Trends

While less profound than the exponential growth of computing and communication, there have been a number of important changes in the way databases are structured and used over the past five years. We note the following:

- While relational systems were well on their way to supplanting earlier approaches in 1990, we now observe that the relational approach is today ubiquitous. Relational databases are today used in applications ranging from those running on top of very large parallel architectures (e.g., NCR 3600) to those running on home computers.
- Client-server architectures have migrated from file systems to database systems over the last five years. We expect it will become progressively more common for database servers to be accessed remotely over networks.
- The traditional, record-based data that populated relational databases of five years ago has been joined by various kinds of "multimedia" data. This trend is fueling the success of ORDB and is causing relational vendors to dramatically enhance their engines to deal with a richer data model. In short, systems that are merely relational DBMSs will be the legacy systems of the next decade.

4.3 Research and the Business Climate

The economics of research and development has changed significantly in the past five years. In general, corporations that traditionally supported substantial basic research have had to cut back, as profit margins on many product and service lines have shrunk. Their research has been reoriented toward short-term projects intended not as prototypes but as frontier work intended directly for market. Moreover, government research agencies have in the recent past felt a similar need to reorient from long-term or basic research to short-term, closely managed projects.

However, there are indications that database research is regarded positively by government and industry. For example, Toole and Young [1995] gives a prominent place to information systems in future strategic research, unlike some earlier documents of this type. The growth of information-intensive industries and the importance of increasing the efficiency of business practices have each put database technologies at or near the top of corporate concerns.

The trend in industrial database research, while mixed, has many positive elements. Companies that have had to restructure their research enterprise have frequently maintained or increased their database component. While many of the companies that are most closely identified with database systems as products have remained substantially outside the research community, some other companies that five years ago did not regard database systems as central to their mission have moved aggressively to build a database research component.

4.4 The Information Superhighway Just Rolled Through Your Living Room

No futuristic discussion can be called complete without a discussion of the World-Wide Web (WWW or Web for short). While people were debating the nature of the "Information Superhighway" or "National Information Infrastructure," the Web — an informal collection of connected "documents" based on HTML (Hypertext Markup Language) — started growing at an astronomical rate. Recently, the number of Web bits carried by the Internet has grown 15–20% per month, or a factor of 10 growth per year.

As a result, the number of active users of the Internet has passed 10% of the US population by some estimates. Following similar technology-based trends such as VCRs or audio CDs, we expect that internet access will reach the great majority of the population just a few years after breaking the 10% barrier. As a consequence, we expect over the next few years several orders of magnitude of growth in the availability and use of information over the Internet. We expect that the provision and use of such information will become a concern of every individual.

Databases and database technology will play a critical role in this information explosion. Already Webmasters (administrators of World-Wide-Web sites) are realizing that they are database administrators, with a different title. Many large Web sites are turning to DBMS technology to keep track of the ever increasing numbers of stored objects. In addition, innovative sites are already prototyping traditional DBMS applications such as electronic catalogs using the Web as an infrastructure.

5 New Research Directions

Against this backdrop, our workshop participants suggested the following research agenda as the most important way to enable the new generation of applications mentioned in Section 3. Indicated research directions are grouped into five major categories:

- Problems associated with putting multimedia objects into DBMSs.
- Problems involving new paradigms for distribution of information.
- New uses of databases.
- New transaction models.
- Problems involving ease of use and management of databases.

5.1 Support for Multimedia Objects

The explosive growth of the Web, as well as the challenges of EOSDIS, electronic commerce, and digital publishing, creates a collection of challenges which must be addressed for future database systems. Below are the principal areas for research involving multimedia data.

5.1.1 Tertiary Storage

Indisputably, multimedia data is very large, and its very bulk presents us with new challenges. As mentioned in Section 3, new applications like EOSDIS or electronic libraries involve data volumes in the petabyte range. Despite the exponential growth in disk sizes, these volumes are unlikely to be maintainable solely on magnetic or magneto-optical disk in the foreseeable future. We are thus presented with problems of managing a new level in the storage hierarchy, called *tertiary storage*. This third level uses storage devices that are orders of magnitude slower than "secondary storage" (disks), yet also of vastly greater capacity. Tertiary storage devices include compact-disk juke boxes or tape silos, and typically use a mechanical arm to physically move the desired tape cassette or CD to a reader.

In a sense, access to tertiary storage is by buffering selected data items on secondary storage, just as access to secondary storage is by buffering selected data into main memory from disk. However, the raw numbers and the proportions make the optimization of the tertiary-to-secondary transition rather different from the secondary-to-main transition. As just one example, today much tertiary information is found on tape cassettes. Not only are tape cassettes slower to access than blocks on a disk by three orders of magnitude (seconds vs. milliseconds), but finding data not at the beginning of a tape cassette can increase the access time by 1–2 orders of magnitude beyond that, while access to the middle of a disk block does not increase access time substantially. Thus, the question of where in a tape cassette information is placed is significant today, while the position of data within a disk block or track is relatively unimportant.

5.1.2 New Data Types

Each form of multimedia information (data type) requires its own collection of first-class concepts (operations and functions), along with a high performance implementation involving appropriate data structures and access methods. As a simple example of the challenge, a recent experiment benchmarking current object-oriented systems noted a vast difference in the performance of such systems on a large text object. When asked to find the last character of a megabyte-long text string, some systems retrieved the whole string and then applied the "find last" operation, while others were able to retrieve only the last byte or some short tail of the string. We need to think carefully about:

- The operations available for each type of multimedia data, and the resulting implementation tradeoffs.
- The integration of data involving several of these new types.

5.1.3 Quality of Service

Delivering multimedia data to many users presents several new research problems. In general, when data is large, access and delivery can easily be a bottleneck. However, large data objects are frequently accessed in a very predictable fashion. For instance, a video server delivering movies to many households at once can assume that each request will remain in force and must be delivered at the standard rate until a "stop" button is pushed. Thus, there are opportunities to optimize access based on predicted use, and those predictions will normally be quite accurate.

Moreover, some multimedia information comes with severe delivery constraints. For example, video generally must be delivered at a fixed rate, and there is a penalty for late delivery, since the movie will appear to blink out or undergo a time distortion. Yet the audio associated with a movie is under a more severe constraint. An appropriate algorithm can interpolate video frames fairly well, but it is not really possible to interpolate missing audio. To make matters more complex, certain kinds of video/audio can support different kinds of delays or delivery failures. For instance, a video of a course lecture often can tolerate frame rates below one per second, since the important visual information may be only what is being written on the board or pointed to on a slide.

These examples suggest the need for research into "quality of service." That is, associated with each form of multimedia data is a collection of options and issues regarding:

1. How best do we assure timely and realistic presentation of the data in its desired form?
2. When the system is unable to meet all expectations for service, how do we gracefully degrade service? Can we interpolate or extrapolate some of the data? Can we reject new service requests or cancel old ones?

5.1.4 Multiresolution Queries

Traditionally, database queries have dealt only with precise notions: "what is the destination of flight 123?" or "what is the balance of account 45678?" Many new applications are well served by queries that use more vague concepts, such as the best available value found among an indeterminate set of loosely integrated resources. For example, if we would like to see a satellite image for a certain time and region, we must ask a resource like EOSDIS for the "best" match, perhaps an image that includes the region and that was taken as close to the desired time as possible. There is thus a need for new query languages or extensions to old ones that incorporate the degree of precision and modes of relaxing the requirements of a query as first-class notions.

Similarly, there are experimental systems that retrieve images from a database of images according to imprecisely defined characteristics, such as shape, color, or texture. Systems of this type offer the potential for retrieval by content in the world of images, video, and other media, similar to our current ability to retrieve character or numeric data by value. However, there is substantial research yet to be done in this direction.

5.1.5 User Interface Support

While SQL or higher-level forms built on top of SQL queries has served for access to conventional, record-based data, multimedia data often require new user interfaces that must be supported by a database management system. For example, geographic or spatial data (maps) are often best queried through a display of the space, upon which the user can outline regions that may be hard or impossible to describe as an SQL query. Likewise, we mentioned in Section 5.1.4 the potential for querying image databases through an interface that allows descriptions of colors, shapes, and other characteristics. Similar problems of facilitating user queries comes up for each form of multimedia data.

Multimedia information also invites us to develop new means for browsing, searching, and/or visualizing the content of multimedia databases. For example, a course may be stored as many hours of video. We might like to know if the course is worth taking, or we may wish to find within that course ten minutes or an hour's worth of material on some subtopic. It is necessary to provide suitable access, and many approaches are possible, such as sample frames, text-based indexes, or the capability to search for segments having specified characteristics.

As another example, satellite images are often far too large to transmit quickly to the user. However, appropriate samples can offer the important characteristics, allowing the user to retrieve and scan many images quickly to find a desired characteristic such as a volcanic plume. One can extrapolate from these examples the opportunities that are available for allowing users to *visualize* the content of massive data objects quickly and efficiently. Providing this capability is an important challenge for the next generation of database systems.

5.2 Distribution of Information

As previously mentioned, the World-Wide-Web is a a large, autonomous distributed systems, whose nodes are increasingly becoming database systems. Likewise, digital publishing entails a distributed system in which there is a very low level of trust between the client and a server. Although distributed database systems have been extensively investigated by the database research community, and the results of their endeavors moved into commercial products, the new environment facilitated by the Web requires rethinking many of the concepts in current distributed database technology. This section indicates the required research agenda.

5.2.1 Degree of Autonomy

The databases and other information sources that are connected through a network frequently are owned by different participants. Distributed health-care systems (Section 3.3), distributed design systems (Section 3.5), and the Web all exhibit this characteristic. Autonomous participants in a distributed system present many special problems to distributed database systems.

The distributed system has to cope with the possibility that some participants may refuse to accept a connection. Different participants may use systems with different capabilities. For instance, a design system may best support consistency of the design through distributed constraints or active rules whose triggering conditions involve more than one autonomous participant. How do we monitor the necessary conditions if participants do not have, or are unwilling to use, active rules in their own systems?

5.2.2 Accounting and Billing

In locally autonomous systems, the server may insist on remuneration in return for performing a service for a client. Previous distributed database systems, in contrast, have assumed that data was a private resource of one corporation or entity, and consequently have ignored this thorny topic.

In an "information for sale" environment, innovative strategies must be implementable, such as "metering," whereby a client pays a small amount for each access to remote data. Collecting such fees efficiently is a subject for research. Obviously it doesn't make sense to "do a dime's worth of work to collect a nickel."

A second interesting question concerns the interaction between query processing strategies and billing rates. For example, suppose we want to explore the literature on dinosaurs. The local museum may provide a free service suitable for this query, but it may not be as complete as a commercial bibliographic service. We would like our query engine to be aware of the difference in costs and pose the question first to the free source. The query engine could then modify the query for the expensive source, saying in effect "tell me all references about dinosaurs except for the following 2000 references that I already know about." On the other hand, the commercial service is unlikely to want to answer that query, because it involves a lot of work with little or no result and (depending on the charging algorithm) little resulting fee. Thus, there are opportunities for research that interconnects charging policies, service-purchase policies, algorithms for querying "cheaply," and algorithms for managing billing information

5.2.3 Security and Privacy

Distributed systems involving many autonomous participants need to support security of information. Sometimes, privacy is the motivation. For example, health-care systems must provide patient information instantaneously to physicians who need it. But the same system should not permit unauthorized access to the data. Other times, it is the commercial value of the data that must be protected. Examples are distributed design (Section 3.5) or electronic publishing (Section 3.4). Some important research goals are:

1. Develop extremely flexible authentication and authorization systems, supporting access according to a variety of "roles" played by individuals. For instance, an individual in the role of physician to a particular patient should have different access privileges than when merely in the role of physician or in the role of citizen.

2. Find mechanisms that will support the sale of information to large numbers of users whose identity is not known to the seller.

5.2.4 Replication and Reconciliation

A fundamental problem of operating a distributed database concerns what happens when the network disconnects the nodes into two or more groups. If a physician gets on an airplane with a copy of his patients records he should be allowed to make annotations (updates) in his data base, even though he is disconnected from his network connection while on the airplane.

Obviously, the connected components of a distributed database system should continue to operate independently as best they can. Queries and data modifications that can be performed by a single component are executed, while those that require more than one component must fail.

For efficiency reasons, data is often duplicated at more than one site. When the network is connected, it is possible to keep the copies consistent. However, when the network disconnects, the copies may become different. On reconnection, some mechanism must *reconcile* the copies, forming one copy that represents all changes.

In the traditional view of distributed databases, disconnection is very infrequent, and the recombination and reconciliation process can be quite complex and time-consuming. In the new information environment, as the above example illustrated, we foresee situations where disconnection is a very frequent event. Thus, it is imperative that we discover high-speed protocols and algorithms for reconciliation.

In addition, the increasing dependence on information systems means that in many applications 100% availability is required, so called 7 x 24 operation. For example, an automated call-routing system cannot shut down and halt or cancel calls. Some of the problems of increasing system availability involve improvements to the reliability of hardware and software. However, in the database domain, we need to enhance availability by research into replication architectures with identical data copies, that can allow the system to proceed correctly in the face of partial failures.

5.2.5 Data Integration and Conversion

Extensive information systems, such as the health-care support systems suggested in Section 3.3, involve the interconnection of information resources having a wide variety of formats and models. It is generally believed that some integrating notation and model must stand in the middle among all these sources. Each source is *wrapped* by a component that translates between the viewpoint of the source and the shared, global viewpoint. Higher-level products may then be built from these wrapped sources.

Many questions are raised by the problem of converting and integrating disparate data sources:

1. What should the integrating model be?

2. What tools are necessary to make the use of arbitrary information sources in integrated systems as easy as using stand-alone databases?

3. How do we extend the idea of a data dictionary to support correct use of terminology among heterogeneous, integrated information sources? This problem is essentially the "ontology problem" of AI.

4. A possible approach to combining disparate information sources is through the use of *mediators*, which can perform customized integration, perhaps with additional filtering or processing. Their role is similar to that envisioned for "agents" in the AI community. The best utilization of this approach is an important research topic.

5.2.6 Information Retrieval and Discovery

It seems evident that increasingly more information will be available on the "web," a collection of informally connected resources on the Internet. The informality and distributed control characteristic of the web contrasts markedly with the structure and control found in today's distributed database systems. This new environment points up the importance for tools to integrate heterogeneous information that was foreseen in the Silberschatz et al. [1990] report.

However, the nature of the web presents further problems, or extreme examples of the earlier problems for heterogeneous data. For instance, we must cope with:

- Data whose schema is unclear, changes without notice, or whose structure is irregular.

- Data whose precise definition is unclear and/or whose reliability is unclear.

- While database technology has been very effective at creating indexes and other support for searches of well structured information, it is essential that these techniques be extended and adapted to the unstructured world of the web.

5.2.7 Data Quality

There has always been a problem of validating the contents of a database in the face of unreliable data input mechanisms. However, the new applications often involve the combination of information found among many resources, and these resources may be of differing reliability. Thus, we are faced with the need to invent methods for evaluating the reliability of such information. Further, we must be able to query the reliability or the *lineage* (description of the origin or origins) of our data. Ideally, the reliability and/or lineage of data will be first-class concepts in new query languages.

5.3 New Uses for Database Systems

Traditionally database systems were used to support business data processing applications, and much DBMS research was focused in this direction. However, three important new uses for the technology have recently emerged. Each poses a new environment for which DBMS technology must be adapted. These three uses have been termed "data mining," data warehousing, and data repositories in the marketplace, and we discuss each in turn.

5.3.1 Data Mining

There is today a great deal of enthusiasm for *data mining*, the extraction of information from large bodies of data often accumulated for other purposes. For example, airlines are getting much better at keeping planes full by "mining" their reservations data for patterns of travel. There is a famous story of a convenience store that discovered there was an unusual correlation in the afternoon between people who bought diapers and people who bought beer; hence, they moved beer and diapers close together and placed potato chips between them, thus increasing sales of all three items.

Data-mining queries pose some unusual problems:

1. They tend to involve aggregations of huge amounts of data.

2. They tend to be ad-hoc, issued by decision makers who are searching for an unexpected relationship.

3. In applications such as trading of commercial instruments, there is a need for extremely fast response, and the figure of merit is total elapsed time, including the writing, debugging, and execution of the query.

4. Often, the user cannot formulate a precise query and his real question is "Find me something interesting".

Thus data mining applications are well served by research into:

- Optimization techniques for complex queries, such as those involving aggregation and grouping.

- Techniques for supporting "multidimensional" queries where the data is organized into a "data cube" consisting of a quantity of interest (e.g., sales) broken down into "dimensions" such as day, item, color, store, supplier.

- Optimization techniques involving tertiary storage (see Section 5.1.1).

- Very high-level query languages and interfaces that support nonexpert users making ad-hoc queries.

5.3.2 Data Warehouses

A *data warehouse* is a copy of data from one or more databases. There are many uses for and approaches to "warehousing." For example, a store may keep a copy of its cash-register transactions in a warehouse for data mining (see Section 5.3.1). A warehouse may store information from (parts of) many databases for emergency purposes. For example, earthquake planners maintain copies of civil infrastructure information (roads, bridges, pipelines, etc.) in a warehouse, because at the time the information would be needed for disaster response, it is unlikely that the required information could be obtained by remote access to databases maintained by the cities and towns closest to the quake site. Yet another example is the use of a warehouse as an "instantiated view" of integrated information. We discussed in Section 5.2.5 the opportunity to build mediators, which are in effect views of many information sources integrated to form a coherent whole. Warehouses can replace mediators by physically storing the integrated information; in contrast, the mediator makes the information available by forwarding queries, much as a view does.

Some research problems associated with creating and using warehouses resemble those involved with information integration in general, while others are special. For example, we need:

1. Tools to create *data pumps*, modules that sit above an information source and obtain from that source those updates that affect the stored information at the warehouse, translating them into the global model and schema of the warehouse.

2. Methods for "data scrubbing": making the data consistent, identifying values that although different represent the same object (e.g., "Sally

Jones" and "S. A. Jones"), and identifying implausible values.

3. Facilities to create and maintain a metadictionary, informing users how data was obtained.

5.3.3 Repositories

There has evolved a class of applications called *repositories*, that can be characterized as storing and managing both data and *metadata* — the information about the structure of the data. Databases in support of computer-aided design, including CASE (software engineering) systems, are examples, as are many document management systems. These systems are distinguished by having to cope with frequently changing metadata, a characteristic of all design environments.

- Repositories must maintain an evolving set of representations of the same or similar information. For example, a program module may be represented as source code, object code, intermediate code, flow diagram, use–definition chains, and documentation. Relationships among these representations must be maintained by the repository, so changes to one representation of an object can be propagated to related representations of the same object.

- Repositories must support notions of *versions* (snapshots of an element evolving over time) and *configurations* (versioned collections of versions.) For example, different releases of a software system would normally be constructed as a configuration composed of a particular version of each source file.

- Repositories must support evolution of the structure of information and its metadata cleanly, e.g., without recompiling when new properties of data or new relationships are added.

The research goal is the creation of "repository management systems" analogous to today's DBMSs.

5.4 Workflow and Transaction Management

As databases become commonplace and are being used in ways not initially anticipated by the business community, the traditional transaction model has come under question. A transaction may now span multiple "independent" databases, and may no longer complete its execution in a very short duration of time.

5.4.1 Workflow Management

Often, business processes involve both machine-based steps, where a database or other information resource is used, and human-based steps, where a person is required to intervene. For example, a travel report is filed by an employee, entered into the computer system by their secretary, automatically formated into the required reimbursement form, and routed to a clerk who electronically approves or rejects the form. If approved, the form is passed to a bookkeeping system that enters the expense and generates a check. Also, the presence of multimedia data heightens the need for workflow management closely coupled with a DBMS. The task of digitizing a paper document is a sequence of steps with intervening human activity. Such steps include scanning, optical character recognition, error checking, and registration of the finished document.

As these examples suggests, such processes require special kinds of data management that support sequences of related events. Some events may have long delays associated with them, e.g., the clerk is on vacation and their backup is at lunch. Events may involve branching and even backtracking, e.g., the form is rejected but needs to be corrected for subsequent approval. Like the repositories mentioned in Section 5.3.3, these *workflow* systems require their own kind of "workflow management system" that supports their special requirements. Workflow processes also require tools for their design, creation, and management. Some of the technology involves transaction models and systems to be discussed subsequently in Section 5.4.2.

5.4.2 Alternative Transaction Models

Transactions are processing units with capabilities important to traditional database systems: atomicity (all-or-nothing gets done), serializability (no unexpected interactions among transactions), and recoverability (if a transaction is completed, then its result is permanent, even if the system subsequently crashes.) These capabilities remain important in many of the new applications, but the accepted ways of implementing transactions are often unworkable. For example, certain applications involve very long steps; examples are CAD systems where parts of designs are checked out by a designer for hours or days, and workflow systems, where individual human-based steps can take hours or days.

Some alternative transaction models have been proposed, including nested transactions that break long transactions into smaller steps and "sagas" that involve the ability to undo the effect of steps that get blocked by later steps. The next research steps involve enhancing these models and building effective support for a wider class of transactions, including CAD and workflow applications as well as others that will come up as we integrate larger numbers of autonomous information sources.

5.5 Ease of Use

Because the role of information in society is expanding so rapidly, the role of the database in society is expanding equally. The small number of large systems a few years ago have been joined by large numbers of smaller installations (as well as additional large ones). However, the complexity of deploying and using such systems has not kept pace with the spread of the systems.

The next generation of database management systems needs improved interfaces, not only for the end-user, but for the application programmer and system administrator. A goal for user support is database systems with the ease of use of spreadsheets, which are today often used as rudimentary database systems.

Installation and upgrading of database management systems is today a daunting task, much more so than for most other kinds of software. This fact is no surprise because of the size and complexity of DBMSs and the variety of physical configurations upon which they must run. However, we recommend a research program whose goal is to create intelligent tools to assist system managers in installing and configuring their systems. In effect, operational systems are presided over by a staff of "wizards" who keep them operating efficiently. Needed is an "electronic wizard," who can lower the manpower cost of operating large DBMS applications.

A related problem is the intelligent support for performance enhancement. One of the most common calls to database-company help lines is a complaint that such and such a query runs too slowly. There are theories for index selection and database schema design that often can solve such problems. What is needed is an electronic "physical database design tool".

6 Conclusions

The technological environment of the United States, and indeed, of the world, is changing rapidly, and with these changes have come expansions to what we view as the natural domain for database technology. Because the information demands of the changing world are stressing the limits of database technology, the database research community must move quickly and aggressively to address the new challenge areas. The opportunities and demands for research in database and information systems today — ranging from highly theoretical investigations of new models and algorithmic foundations, to the creation of ambitious new system prototypes — have never been greater. Yet the funding level for database research remains substantially below that of other areas of comparable importance, as is made clear by data in Toole and Young [1995].

Therefore, the workshop participants recommend a renewed call to action to those government agencies charged with assessing and promoting research in the database target areas, and to those commercial enterprises that benefit from database research. In conclusion, we reiterate two important recommendations of the 1990 report (Silberschatz et al. [1990]), which remain timely.

- The NSF together with the other agencies represented in the Federal Coordinating Council on Science, Engineering, and Technology that fund basic research should, with input from both academia and industry, develop a strategy to ensure that basic database research is funded at a level commensurate with its importance to scientific research and national economic well-being.

- U.S. industrial firms with a substantial stake in database technology should vigorously support existing programs and the development of new programs that provide funding for basic university research in the database area.

We anticipate a decade of exciting achievements in both academic and industrial database research circles, and we look forward to a dynamic and proactive response from government and commercial entities in supporting and furthering this vital and expanding field.

7 References

Gray [1995]: `http://www.cs.washington.edu/homes/lazowska/cra/database.html`.

Silberschatz at al. [1990]: "Database systems: achievements and opportunities," *SIGMOD Record* **19**:4, pp. 6–22. Also in *CACM* **34**:10 (Oct., 1991), pp. 110-120.

Toole and Young [1995]: `http://www.hpcc.gov/cic/forum/CIC_Cover.html`.