



The Lowell Database RESEARCH SELF-ASSESSMENT

Database needs are changing, driven by the Internet and increasing amounts of scientific and sensor data. In this article, the authors propose research into several important new directions for database management systems.

A group of senior database researchers gathers every few years to assess the state of database research and to point out problem areas that deserve additional focus. This article summarizes the discussion and conclusions of the sixth such meeting in Lowell, MA, in May 2003. It follows a number of earlier reports with similar goals, including [1, 2, 5–7].

Continuing this tradition, 25 senior database researchers representing a broad cross section of the field in terms of research interests, affiliations, and geography gathered in Lowell for two days of intensive discussion on where the database field is and where it should be going.

Several important observations came out of this meeting. Information management continues to be a critical component of most complex software systems. We recommend that database researchers increase their focus on the integration of text, data, code, and streams; fusion of information from heterogeneous data sources; reason-

BY SERGE ABITEBOUL,
RAKESH AGRAWAL,
PHIL BERNSTEIN, MIKE CAREY,
STEFANO CERI, BRUCE CROFT,
DAVID DEWITT, MIKE FRANKLIN,
HECTOR GARCIA MOLINA,
DIETER GAWLICK, JIM GRAY,
LAURA HAAS, ALON HALEVY,
JOE HELLERSTEIN, YANNIS IOANNIDIS,
MARTIN KERSTEN,
MICHAEL PAZZANI, MIKE LESK,
DAVID MAIER, JEFF NAUGHTON,
HANS SCHEK, TIMOS SELLIS,
AVI SILBERSCHATZ,
MIKE STONEBRAKER,
RICK SNODGRASS, JEFF ULLMAN,
GERHARD WEIKUM,
JENNIFER WIDOM, AND
STAN ZDONIK

ing about uncertain data; unsupervised data mining for interesting correlations; information privacy; and self-adaptation and repair.

Our community focuses on information storage, organization, management, and access. It is driven by new applications, technology trends, new synergies with related fields, and innovation within the field itself. The nature and sources of information are changing. Everyone is aware that the Internet, the Web, science, and e-commerce are enormous sources of information and information processing demands. Another big source is coming soon: inexpensive

microsensor technology that will enable most things to report their state in real time. This information will support applications whose main purpose is to monitor the reporting objects' state or location. Sensor information processing will raise many of the most interesting database issues in a new environment, with a new set of constraints and opportunities.

The Internet is currently the main driving force

for applications, particularly by enabling cross-enterprise applications. Historically, applications were intra-enterprise and could be specified and optimized entirely within one administrative domain. However, most enterprises are interested in interacting with their suppliers and customers to share information and provide better customer support. Such applications are fundamentally cross-enterprise and require stronger facilities for security and information integration. They generate new issues for the Database Management System (DBMS) community.

A second application area of growing importance is the sciences—notably the physical sciences, biological sciences, medicine, and engineering—generating large and complex datasets that require more advanced database support than current products provide. They also need information integration mechanisms. In addition, they need help managing the pipeline of data products produced by data analysis, storing, and querying ordered data (including time series, image analysis, computational meshes, and geographic information), and integrating with the worldwide data grid.

In addition to these new information management challenges, we face major changes in the traditional DBMS areas, such as data models, access methods, query processing algorithms, concurrency control, recovery, query languages, and user interfaces to DBMSs. These topics have been well studied in the past. However, technology keeps changing the rules. For example, disks and RAM are getting much larger and much cheaper per bit of storage. Access times and bandwidth are improving too, but not as rapidly. These changing ratios require us to reassess storage management and query processing algorithms. In addition, processor caches have exploded in size and number of levels, requiring DBMS algorithms to be cache-aware. These are but two examples of technological change inducing a reassessment of previous algorithms.

Another driver of database research is the maturation of related technologies. For example, over the past decade, data mining technology has become an important component of database systems. Web search engines have made information retrieval a commodity that must be integrated with classical database search techniques. Many areas of artificial intelligence are producing components that could be combined with database techniques; for example, these components allow us to handle speech, natural language, reasoning with uncertainty, and machine learning.

Participants noted that it is a popular undertaking these days to propose “grand challenges” for various fields of computer science. Each grand challenge is a problem that cannot be solved easily, and is intended as a call to action for a given field, such as The Informa-

tion Utility [2] and Building Systems With Billions of Parts [4]. We all agreed that we could define more grand challenges. In fact, we discussed a few, notably the personal information manager—a database that could store, organize, and provide access to all of a person’s digitally encoded information for a lifetime. In the end, we decided that focusing on a single grand challenge was inappropriate, since information management technology is a critical component in most, if not all, of the proposed CS grand challenges. Moreover, many of those information management challenges are well beyond the state of the art. The existing grand challenges are a full employment act for the database community, so we decided not to add more.

We noted many new applications, technology trends, and synergies with related fields that affect information management. In the aggregate, these issues require a new information management infrastructure. We survey its components here.

Next-Generation Infrastructure

Integration of text, data, code, and streams. The DBMS field has always focused on capturing, organizing, storing, analyzing, and retrieving structured data. Until recently, there was limited interest in extending DBMSs to manage text, temporal, spatial, sound, image, or video data. However, the Web has clearly demonstrated the importance of these more sophisticated data types. The general problem is that, as systems add capabilities, it is difficult to make these additions cleanly. Rather, there is a tendency to do the minimum necessary to have the most important of the desired new features. As a result, these extensions tend to create “second-class citizens,” or objects not usable in all of the contexts where the traditional “first-class citizens” of a DBMS (integers, character strings, among others) may appear. Here are some examples where rethinking the way we handle certain elements could improve the usability of a system.

Object-oriented (OO) and object-relational (OR) DBMSs showed how text and other data types can be added to a DBMS and how to extend the query language with functions that operate on these extended data types. Current database systems have taken their first steps toward supporting queries across text and structured data, but they are still inadequate for integrating structured data retrieval with the probabilistic reasoning characteristic of information retrieval. To do better, we need a fresh approach to both data models and the query language required to access text and data. At the very least, probabilistic reasoning and other techniques for managing uncertainty must become first-class citizens of the DBMS.

A major addition to recent DBMSs is their ability to

add user-defined procedures to the query language. This approach allows the user to add a new data type along with its behavior (methods). Unfortunately, this approach makes procedures second-class citizens. We would like to see code become a first-class citizen of the DBMS, as well.

Triggers and active databases are another source of executable code inside a DBMS. Often, users want to be alerted if a specific database condition becomes satisfied or a certain event occurs. If there are millions of such conditions, it is inefficient or even infeasible to poll the database periodically to see which of the conditions are true. Rather, one wants to specify the monitoring condition and then have the DBMS alert the user asynchronously if the indicated condition becomes true. Commercial vendors have added triggers and alerters to their products, and there has been considerable research on how to make such facilities

scientific discipline is generating huge data volumes, for example, from particle accelerators (physics), telescopes (astronomy), remote sensors (earth sciences), and DNA microarrays (biology). Simulations are also generating massive datasets. Organizing, analyzing, and summarizing these huge scientific datasets is a real DBMS challenge. So is the positioning and transfer of data among various processing and analysis routines distributed through a grid, which requires knowledge of the overall structure of the processing chain and the needs and behavior of each module in it. Meeting these needs will require an integration of data and procedures that allows complex objects and advanced data analysis to be integrated with the DBMS.

It is time to stop grafting new constructs onto the traditional architecture of the past. Instead, we should rethink basic DBMS architecture with an eye toward supporting:

A new form of science is emerging. Each scientific discipline is generating huge data volumes. Organizing, analyzing, and summarizing these huge scientific datasets is a real DBMS challenge.



scalable. However, triggers and alerters have been grafted onto existing DBMS architectures. While it is not feasible to reason completely about code in the general case, it would be useful to have ways for the DBMS to do simple, perhaps only syntactic, reasoning about code objects. For instance, we could hope for the ability to find all code that depends upon a given database object.

We expect that several emerging application classes will force data streams to become a first-class part of the DBMS as well. The imminent arrival of low-cost commercial microsensor devices will enable new classes of monitoring DBMS applications. It will become practical to tag every object of importance with a sensor that will report its state in real time. For example, instead of attaching a property tag to items such as laptop computers and projectors, one would attach a sensor. One could then query a monitoring system for the location of a lost or stolen projector. Such monitoring applications will be fed streams of sensor information on objects of interest. These streams will put new demands on DBMSs in the areas of high-performance data input, time-series functionality, maintenance of histories, and efficient queue processing. Presumably, commercial DBMSs will try to support monitoring applications by grafting stream processing onto the traditional structured-data architecture.

Lastly, a new form of science is emerging. Each sci-

- Structured data;
- Text, space, time, image, and multimedia data;
- Procedural data; that is, data types and the methods that encapsulate them;
- Triggers; and
- Data streams and queues

as co-equal first-class components within the DBMS architecture—both its interface and its implementation—rather than as afterthoughts grafted onto a relational core.

The participants suggested it would be better for the research community to start with a clean sheet of paper in many cases. Attempts to add these capabilities to SQL, XML Schema, and XQuery are likely to result in unwieldy systems that lack a coherent core. Because of their dependence on prior standards, we believe strongly that XML Schema and XQuery are already too complex to be the basis for this new architecture. The self-describing record format is a great idea for communication of information, but it is not especially convenient for the DBMS we envision, where procedures, text, and structured data are co-equal participants. A new information architecture cannot be burdened by the political compromises of the past. We believe vendors will pursue the extend-SQL and extend-XML strategies to improve their existing products incrementally. By contrast, the research community should explore a reconceptualiza-

tion of the problem.

A start on such an architecture should be a five-year goal for our community. As a concrete milestone, we look for several substantial prototypes before the next meeting of our ad hoc group.

Information fusion. Enterprises have tackled information integration inside their semantic domains for more than a decade. The typical approach is to extract operational data, transform the resulting data into a common schema, and then load the result into a data warehouse for subsequent querying. Information integration is performed in advance, typically by an extract-transform-load (ETL) tool to build the data warehouse and data marts. This is a feasible approach within an enterprise with a few dozen operational systems under the control of a single corporation.

The Internet completely breaks this extract-transform-load paradigm. There is now a need to perform information integration among different enterprises, often on an ad hoc basis. Few organizations will allow

at Web scale remains elusive. Our community must seriously focus on this issue, or cross-enterprise information integration will remain a pipe dream. The same problem is being investigated in the context of the semantic Web. Collaboration between groups working on these and other related problems, both inside and outside the database community, is important.

There are many other difficult problems to be solved before effective Web-scale information integration becomes a reality. For example, current federated query execution systems send subqueries to every site that might have data relevant to a query, thereby giving a complete answer to every query. At Web scale, this is infeasible and query execution must move to a probabilistic world of evidence accumulation and away from exact answers. As another example, conventional information integration tacitly assumes the information in each database can be freely shared. When information systems span autonomous enterprises, query processing must be done such that each database reveals only the

A semantic heterogeneity solution capable of deployment at Web scale remains elusive. Our community must seriously focus on this issue, or cross-enterprise information integration will remain a pipe dream.



outside entities to extract all the data from their operational systems, so the data must stay at the sources and be accessed at query time. Some commercial products do this today, but with a relatively small, static set of sources within one enterprise.

As mentioned earlier, sensor networks and the new science will be generating huge datasets. These sensors and datasets will be distributed throughout the world, and can come and go dynamically. This breaks the traditional information integration paradigm, since there is no practical way to apply an ETL tool to each such occurrence.

Therefore, one must perform information integration on the fly over perhaps millions of information sources. The DBMS research community has investigated federated data systems for many years. The first of these reports [1] talked extensively about the problem. However, the thorny question of semantic heterogeneity remains. Any two schemas designed by different people will never be identical. They will have different units (one salary is in euros, another is in dollars), different semantic interpretations (one salary is net including a lunch allowance, another is gross), and different names for the same thing (Samuel Clemens is in your database, but Mark Twain is in mine). A semantic heterogeneity solution capable of deployment

minimum information necessary to answer the query, and in conformance with its security policy. A third example is tying information integration to monitoring applications that span multiple data sources. For example, let me know when any of my mileage plans is giving bonuses on hotel stays for chains that have hotels near the sites of conferences or meetings I will be attending.

Sensor data and sensor networks. Sensor networks consist of very large numbers of low-cost devices, each of which is a data source, measuring some quantity—the object's location, or the ambient temperature, for example. These networks provide important data sources and create new data management requirements. For instance, these sensors will generally be self-powered, wireless devices. Such a device draws far more power when communicating than when computing. Thus, when querying the information in the network as a whole, it will often be preferable to distribute as much of the computation as possible to the individual nodes. In effect, the network becomes a new kind of database machine, whose optimal use requires operations to be pushed as close to the data as possible.

Query execution on sensor networks requires a new capacity: the ability to adapt to rapidly changing configurations, such as sensors that die or disconnect from

the network. The query plan needs to change as the network changes, a capability we do not see among database systems today.

Sensors also suggest the need to deal with more complex forms of information integration. A common case is when sensors are not completely calibrated. A value from a sensor must be interpreted in the light of what other sensors are saying. A more complex matter is that the goal of sensor data processing may be to deduce a very high-level fact from very low-level signals. For instance, we may want to combine heat, sound, and vibration sensors to locate a person nearby.

Multimedia queries. Obviously, the amount of multimedia data (pictures, video, audio, and so on) is increasing dramatically. A challenge to our community is to create easy ways to analyze, summarize, search, and view the “electronic shoebox” of a person’s multimedia information. The challenge ranges from Vannevar Bush’s Memex [3] vision to the prosaic task of preparing a multimedia presentation about the children for Aunt Betsy. Any of these goals requires much better facilities for managing multimedia information than are available today.

Reasoning about uncertain data. Traditional DBMSs were applied to business data processing, which typically focused on numbers and character strings. In those application areas, data elements are precise quantities like address, quantity on hand, balance, status, and delivery date. As a result, current DBMSs have no facilities for either approximate data or imprecise queries.

When one leaves business data processing, essentially all data is uncertain or imprecise. Scientific measurements have standard errors. Location data for moving objects involves uncertainty in current position. Sequence, image, and text similarity are approximate metrics. To analyze imprecision, the lineage (provenance) of the data must be tracked, since a scientist needs to know where the data came from (the instruments and their settings) and what cleaning, rescaling, or remodeling, was done to arrive at the data to be interpreted. Clearly, DBMSs need built-in support for data imprecision.

Query processing must move from a deterministic model, where there is an exact answer for every query, to a stochastic one, where the query processor performs evidence accumulation to get a better answer to a query. Users should also be able to ask imprecise queries and have the processing engine include this further source of uncertainty. Of course, with imprecise answers comes a duty for the system to characterize the accuracy offered, so users can understand whether the approximation is good enough for their needs. For instance, information retrieval systems measure precision and

recall to help researchers understand how good an answer is.

Personalization. Several participants noted that query answers should depend on the profile of the person asking. The answer for a domain expert should be different from the answer for a novice. Relevance and relevance feedback should also depend on the person and the context. It should be possible for data at multiple sources organized for one purpose to be offered for other purposes. For example, health information organized for health-care providers should be personalized for individual use (for example, hospital records, prescriptions, drug interaction information, family medical history, immunization records, dental records, and insurance claims). Mass personalization should be feasible in next-generation information systems. We need a framework for including and exploiting appropriate metadata for such personalization.

Participants also noted that personalization and uncertainty produce a need to verify that the information system is producing a correct answer. What happens if the information system is buggy and produces the wrong approximate or personalized answer?

Data mining. Historically, data mining has focused on efficient ways to discover models of existing data sets. These models must expose some useful aspect of the data, while obscuring details not useful for the intended application. Algorithms have been developed by many research communities to perform such operations as classification, clustering, association-rule discovery, and summarization. These techniques are now part of mainstream products from the major DBMS vendors. Today, data mining and business intelligence are heavily used by Fortune 500 end users and by small applications alike. The success of current data mining tools has created a hunger for the next generation of tools. Fortune 500 warehouse users invariably point out they have a single data mining query: “Tell me something interesting.” They are glad to have the current crop of data mining tools, but wish for tools that are better at generating these unexpected pearls of wisdom.

A challenge for data mining research is to develop algorithms and structures for sifting through the database looking for such pearls, which can run in the background using spare system resources. Another important challenge is to integrate data mining with querying, optimization, and other database facilities such as triggers. We hope data mining will move in this direction and beyond algorithms for elementary operations. There was also a feeling that CS and IT curricula should include more about the proper use of data-mining tools.

Self-adaptation. One consequence of the widespread use of DBMS technology is a shortage of com-

petent database administrators. Modern DBMSs are much more complex than their counterparts of 20 years ago. Today, a database administrator must understand disk partitioning, parallel query execution, thread pools, and user-defined data types. These concepts were not present in the systems of yesteryear. Indeed, current DBMSs are thought to be too difficult to use. To address this shortcoming, the major vendors have all embarked on projects to simplify database administration.

Such projects have at least two components. First, current DBMSs have a large collection of tuning knobs that allow an expert to extract the best performance from an operational system. Often, such tuning is done by the vendor's engineers, at great expense to the customer. Most systems engineers doing this tuning do not, in fact, have a deep understanding of the meaning of the knobs. Instead, they have seen many system configurations and workloads and carry a crib sheet of tuning parameters that had optimized other systems. When presented with a new environment, they reach for the crib sheet closest to the configuration at hand and use those settings.

It should be possible to perform tuning using a combination of a rule-based system and a database of knob settings and configuration data. There has been substantial progress in this direction, mostly by the vendors, in such areas as dynamic resource allocation, the selection of physical structures and, to some extent, the selection of materialized views (redundant data maintained by the DBMS to speed up the execution of certain queries). In our opinion, the ultimate goal is "no knobs," where all tuning decisions are made automatically by the system, perhaps guided by default policies, such as the relative importance of response time and throughput, or by personal profiles that summarize user needs. Thus, more sophisticated models of user behaviors and workloads are a prerequisite for progress in this area. We believe real no-knobs operation is possible and recommend this goal as a major focus of the research community.

Many new applications that use DBMSs are going to require unattended operation. In addition to no-knobs tuning, the DBMS must be able to recognize internal malfunctions and malfunctions of communicating components, identify data corruption, detect application failures, and do something about them. Such capabilities require making the DBMS more self-aware and providing it with explicit models of the information system in which it participates.

Privacy. The widespread adoption of the Internet has created a surge in information available about individuals. Moreover, there is plenty of computing power available to perform correlations among databases.

This convergence has made it possible to discover an unprecedented amount of information about individuals. Mundane data, such as every address at which someone has ever lived, is readily obtained. It appears easy to discover every person who has ever lived at a given address, making it easy to discover all the neighbors any given person has ever had. It is claimed that the names of persons who rode on the same airplane flight can also be obtained. Furthermore, identity theft is a troubling problem, since it is not difficult to obtain all the information needed to apply for a credit card in someone else's name.

Data-oriented security research was prevalent in the 1980s but has died down since. We see a need to revitalize this subfield, but with a distinctly different orientation. Today, we need to address the concerns, policies, and mechanisms to support multiple individual options and controls on information held by third parties. This collection of issues is likely to be quite different from those of the pre-Web security model of a single organization protecting its data from unauthorized access. While laws will continue to have an important role in addressing information privacy and related security issues, we can change the set of options and reach better points in the space of privacy-security trade-offs by advancing what is technically realizable.

Since much of this information correlation is performed by DBMSs, our community can work on security systems that include a component dealing with the prospective use to which the data will be put. Access decisions should be based not only on who is requesting the data but on to what use it will be put. Moreover, declarative systems that specify the purpose of the data request are something our community should be good at, since we have dealt with data-oriented declarative specifications in other contexts.

Trustworthy systems. Privacy is only one aspect of the broader issue of trustworthy systems that safely store data, protect it from unauthorized disclosure, protect it from loss, and make it always available to authorized users. There is increasing interest in digital rights management—protecting intellectual property rights and allowing private conversations. Furthermore, ensuring the correctness of query results and data-intensive computations is a growing concern, especially for embedded systems such as health care smart cards and other medical applications. Logical inferencing technology, such as theorem proving and model checking, may be helpful in validating correctness. The information management community should play a central role in addressing these needs and enhancing DBMSs with mechanisms to support these capabilities.

New user interfaces. It is a long-standing lament that the database community does too little about user

interfaces. The desktop now has the horsepower to run sophisticated visualization systems. However, for a given type of information coming from a DBMS, it is not clear how best to render it visually. A small number of slick visualization systems oriented toward information presentation were proposed during the 1980s, notably QBE and VisiCalc. There have not been comparable advances in the last 15 years, and there is a substantial need for better ideas in this area.

Thirty years of research on query languages can be summarized by “we have moved from SQL to XQuery.” At best, we have moved from one declarative language to a second declarative language with roughly the same level of expressiveness. It has been well documented that end users will not learn SQL; rather, SQL is a notation for professional programmers. We see in other communities a number of ideas that could affect database-oriented research on interfaces. The information retrieval community has used keyword-based querying for decades. Browsing has

obsolete devices (such as special storage drives), or because the application needed to interpret the information no longer works. Preventing this information loss requires mechanisms for migration, to copy information from deteriorating or obsolete media, and for emulation, to capture methods that can interpret information that is stored for long periods.

Metadata also plays an important role. For example, the capture of scientific data requires the lineage of the data be captured along with any programs required to gain access to the data. Metadata describing context may also be essential. For example, tables on social service spending in Germany over time might not explicitly show that the 1983 tables are in deutschmarks while the 2003 tables are in euros, or that the 1983 table only included states in West Germany. Without context, the data is worthless, even if it is accessible. In general, society would benefit immensely by our information management research community architecting such a store, whose content remains accessible in a use-

Thirty years of research on query languages can be summarized by “we have moved from SQL to XQuery.” At best, we have moved from one declarative language to a second declarative language with roughly the same level of expressiveness.



become increasingly popular in a number of areas.

Perhaps most interesting are the research opportunities suggested by the term “semantic Web.” While it may be unclear what the concept truly entails, much of the recent work has centered on ontologies. An ontology characterizes a field or domain of discourse by identifying concepts and relationships between them, usually in a formal language. This work may support information integration, since a fundamental problem in that area is the inability to combine databases that at a deep level are talking about the same thing, but do so using different terminology. Work on ontologies may likewise enable users of databases and other resources to use speech or natural language to query in their own terminology. The database community should be looking for opportunities to exploit these developments in future database management systems.

100-year storage. As an ever-increasing fraction of the world’s information is stored digitally, there is a need for indefinite electronic storage of information. However, even archived information is disappearing, because it was captured on a medium that is deteriorating (such as photographic film or magnetic tape), because it was captured on a medium that requires

ful form indefinitely. Such a store would, to the extent possible, automate the process of migrating content between formats or maintaining the hardware and software that each document needs to be used. It would also manage the metadata along with the stored documents.

Query optimization. Many of the participants saw query optimization as an important element of some of the endeavors discussed here. The general principle is that when we deal with very large volumes of data, we tend to manipulate that data in a regular way. This regularity allows very high level languages like SQL or XQuery to be used successfully in the database world, but almost nowhere else. Very high-level languages demand a competent optimizer. Participants suggested we need further work on optimization of information integrators for semistructured query languages like XQuery, for stream processors, for sensor networks, and possibly for other domains as well.

Many uses of SQL systems involve sequences of relatively simple queries, embedded in a host language program, all cooperating to accomplish one task. We see the need to consider interquery optimization involving large numbers of queries, even in the very traditional, pure relational setting.

Next Steps

Many of the suggested research directions are very long-term goals. However, the Lowell meeting participants had several suggestions for steps that should be readily accomplishable by the time the next meeting is held in five or so years. Several goals are mentioned here, such as the proposal to reconsider DBMS architecture to handle new data types, approximate reasoning, and treating procedures and data as co-equal. Here, we note some short-term goals.

Information integration research would be well served by generating a test bed and collection of integration tasks. Doing so would allow easy access to a test platform for anybody with an integration idea. The test bed would allow a controlled way to compare solutions, including the opportunity for researchers to gain bragging rights if their technique is currently the best at one of these integration tasks. It would also help generate interest in this research area. Several researchers noted TREC (trec.nist.gov/data.html) served the same purpose in the information retrieval community.

There was considerable discussion about how to construct such a test bed. While there are many design issues to enable controlled experiments, the most pressing problem is obtaining appropriate datasets. One potentially practical proposal was for 10–20 CS departments to make public some of their nonproprietary classroom scheduling or other similar data. Any member of a CS department who can make such a dataset available is urged to contact Mike Stonebraker (stonebraker@csail.mit.edu), who will coordinate the construction of this test bed. Perhaps a more sophisticated test bed with much larger datasets can be proposed.

There was considerable support for the position that many of the traditional research focuses were solved problems (for example, ACID properties). Other people countered that sea changes in related technology might require revisiting these problems, and that there is always the potential of finding simpler mechanisms that can be more easily implemented or more powerful mechanisms that can be more broadly applied. There was also a heated discussion concerning whether stream processing systems would require a new DBMS engine or whether current ones could adapt to the new requirements successfully.

Our discussion also centered on the level at which information integration should occur. Although many thought the DBMS was the best place to do it, others thought it would be more natural and widespread at the application (Web services) level. There was also considerable discussion about whether Web services would make any progress on addressing semantic heterogeneity issues. Some thought that de facto standards would emerge for such services, while others countered

that the electronic component community has been trying hard for years to standardize on a set of services (Rosetta Net) with limited success.

We close this article with two observations on which there was general agreement. First, the database research community should avoid drawing too narrow a box around what we do. We must explore opportunities for combining database and related technologies that can improve the usage of information, such as information visualization technologies, which has often been left to the domain of other research communities. To broaden the set of technologies database researchers apply, they need to expand their breadth of competencies. Consider the plasterers' union, which decided many years ago, when wallboard was being introduced, that it was not their competency. As plaster was replaced by wallboard, the union lost out. This fate could befall the DBMS community if it does not respond to the new challenges of integrating related technologies with information management.

Second, it was noted that the average age of participants at these meetings has been increasing. On the other hand, there are more young database researchers than ever before, as evidenced by the large number of junior faculty in databases. We recommend the next meeting invite a broader mix of age groups within our community. **C**

REFERENCES

1. Bernstein, P.A., Dayal, U., DeWitt, D.J., Gawlick, D., Gray, J., Jarke, M., Lindsay, B.G., Lockemann, P.C., Maier, D., Neuhold, E.J., Reuter, A., Rowe, L.A., Schek, H.-J., Schmidt, J.W., Schrefl, M., and Stonebraker, M. Future directions in DBMS research—The Laguna Beach participants. *SIGMOD Rec.* 18, 1 (1989), 17–26.
2. Bernstein, P.A., Brodie, M.L., Ceri, S., DeWitt, D.J., Franklin, M.J., Garcia-Molina, H., Gray, J., Held, G., Hellerstein, J.M., Jagadish, H.V., Lesk, M., Maier, D., Naughton, J.F., Pirahesh, H., Stonebraker, M., and Ullman, J.D. The Asilomar report on database research. *SIGMOD Rec.* 27, 4 (1998), 74–80.
3. Bush, V. As we may think. *Atlantic Monthly* (July 1945), 101–108.
4. CRA Conference on “Grand Research Challenges” in Computer Science and Engineering (Warrenton, VA, June 23–26, 2002); www.cra.org/Activities/grand.challenges/.
5. Silberschatz, A., Stonebraker, M., and Ullman, J.D. Database systems: Achievements and opportunities. *Commun. ACM* 34, 10 (Oct. 1991) 110–120.
6. Silberschatz, A., Stonebraker, M., and Ullman, J.D. Database research: Achievements and opportunities into the 21st century. *SIGMOD Rec.* 25, 1 (1996), 52–63.
7. Silberschatz, A., Zdonik, B., et al. Strategic directions in database systems: Breaking out of the box. *ACM Comput. Surv.* 28, 4 (Dec. 1996), 764–778.

Correspondence regarding this article should be addressed to Jeff Ullman (ullman@gmail.com).
