

Panel: Future Directions of Database Research - the VLDB Broadening Strategy, Part 1

Hans-Jörg Schek

ETH Zurich
8092 ETH-Zentrum, Zurich,
Switzerland
Schek@inf.ethz.ch

Abstract

This panel introduces and explains the “future directions” and “broadening” discussions at the VLDB Endowment. Panelists are Michael Brodie, Stefano Ceri, Umesh Dayal, John Mylopoulos, and Hans Schek. We report on the steps that are being implemented at the upcoming VLDB Conferences (Mylopoulos). The main objective of the future direction and broadening discussions is to ensure that database technology remains vital. It should keep its pivotal role as infrastructure for application development for data-intensive, central and distributed applications. In a first part we explain the announcement that has been published through DBWorld (Schek). In a second part we report on statistics taken over the last conferences and on first experiences in the implementation of the broadening strategy this year (Brodie). We continue with steps that are undertaken for the upcoming Rome Conference (Ceri) and we conclude with further perspectives on future directions and broadening (Dayal)

1. Future Directions of Database Research - Changes in the VLDB Conference PC Structure

The following text is taken from [VLDB] after some minor editing. For other interesting observations on the evolution of database research the reader is referred to documents such as [Si+96] and references given there.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.

Proceedings of the 26th International Conference on Very Large Databases, Cairo, Egypt, 2000

1.1 Background

At the New York VLDB Endowment meeting, concerns were expressed that the area of database research may lose the pivotal role it now plays among information system technologies. It was agreed that the Endowment should maintain a watch on trends and future directions in the general area of information management to ensure that database research in general, and the VLDB conferences and journal in particular remain current and relevant. It was decided to set up a small "future directions" working group which maintains an ongoing dialogue with key researchers in the DB community and reports back to the Endowment Board. Working group members were Rakesh Agrawal, Michael Brodie, Michael Carey, Umesh Dayal, Jim Gray, Yannis Ioannidis, John Mylopoulos, Hans Schek, Kyu-Young Whang and Jennifer Widom.

1.2 Situation

The main observation of the working group was the following: While database technology has achieved a high standard in research and development, its future role in a globally distributed information network is less central. Considering the central role of databases in the past in providing a platform for application development, the members strongly felt that actions must be taken. Of course, it is true that databases still do a great job as storage managers for many applications. However, it was felt that the distance to applications and application development has grown substantially during the past decade. Among others, the following examples of application development and application areas were discussed as evidence for the need to act: Managing components, application services, Distributed client/middleware/server computing, Application frameworks, ERPs, XML, e-commerce. To bring such topics into the fold of database conferences, the working group proposed to distinguish between the two main research directions: (1) core database technology and (2) infrastructure for information system development

While (1) is well established and will be further developed worldwide by the international DB research and development community, the infrastructure direction (2) is under-developed and under-represented at conferences. In all examples given above, the database role of providing a storage manager, i.e. core DB technology remains central. However, client application development takes place on other platforms provided by middleware technologies. Or client applications are obtained by customizing pre-fabricated ERP systems. The focus of our community should turn to the investigation of how core technology can become more widespread and usable, by concentrating on the description of new application areas, on the methods and tools for data analysis, design and integration, on the technologies for data deployment in modern architectures (middleware, wireless technology, the WEB), and in general on all the problems and challenges which are due to the need of using very large databases in new contexts.

1.3 Action

In order to evolve into these directions the Endowment has decided in its 1999 Board meeting in Edinburgh to take actions with regard to the selection of tutorials, panels and invited speakers for future VLDB conferences. Most important for the community is the following change in the VLDB Conference PC structure: While in the past there were two PCs for the scientific and industrial tracks of the program, in future VLDB conferences a new track called "IS infrastructure and applications" will be added. Therefore in future Calls-for-Papers for VLDB conferences, three subcommittees will be distinguished, responsible respectively for papers on:

- Core DB technology
- IS infrastructure and applications
- Industrial applications and experience

The new strategy will be fully implemented in the year 2002 but efforts are undertaken to realize it earlier in the upcoming VLDB conferences in Cairo and Rome. The strategy for VLDB2000 and experiences are summarized in [Bro00].

2. Additional remarks and related observations

When relational database systems have been introduced twenty years ago, they have been considered as infrastructure and main platform for development of data-intensive applications. Data independence was considered to be a breakthrough: Programmers were freed from low-level details, e.g., how to access shared data efficiently and correctly, given concurrent access. But by now, the prerequisites for application development have changed dramatically. For instance, communication has become

fairly cheap, and the internet dominates modern information infrastructures. Consequently, the role of database concepts must be re-visited and newly determined.

Carey, Hellerstein and Stonebraker [Regr] observe that all current databases have been designed with the technology of twenty years ago. They state that due to a three-tier architecture, data are at the bottom and application code is away from data in the middle tier. They also state that databases are "bloated" by object-relational features, by stored procedures and triggers, and by warehouse features. They conclude that we should rethink everything.

In a keynote speech Brodie [Bro99] states that states that "the database era nears its end" because DBMSs, in their current form, cannot adequately deal with current heterogeneity and interoperability demands let alone the vast increases in data and transaction volumes of next generation applications. Architectural complexity is another issue in view of the many engine and repository types with ad-hoc solutions for warehousing and mining.

The database group at ETH Zurich [AHST97], [Sch+00], in order to move into new directions, strives for „higher order data independence“ in its hyperdatabase projects. A hyperdatabase is a database over databases and other specialized components. A hyperdatabase applies database technology at a higher level.

2. References

- [AHST97] Alonso, G., Hagen, C., Schek, H.-J., Tresch, M.: *Distributed Processing over Stand-alone Systems and Applications*, In: Proc. of 23rd International Conference on Very Large Data Bases (VLDB'97), Athens, Greece
- [Bro00] Brodie, M.L.: *Future Directions of Database Research – the VLDB Broadening Strategy, Part 2*, this Proceedings
- [Bro99] Brodie, M.L.: *Que Sera, Sera: The Coincidental Confluence of Economics, Business, and Collaborative Computing*, Proc. of the 15th International Conference on Data Engineering, Sydney, Australia, March 1999, pp. 2-3
- [Regr] Carey, M., Hellerstein, Stonebraker, M.: Seminar talk "A sketch of Regres", http://www.cs.berkeley.edu/~gribble/summaries/talks_seminars/regres.html
- [Sch+00] Schek et al. : *Hyperdatabases*, In: Proc. of the WISE2000 Conference, Hongkong, June 2000.
- [Si+96] Silberschatz, A. et al.: *Strategic directions in database systems – breaking out of the box*. ACM Computing Surveys, Vol. 28, No. 4, Dec. 1996, pp. 764-778.
- [VLDB] <http://www.vldb.org/>