Practitioner Problems in Need of Database Research

Gomer Thomas
Bellcore
444 Hoes Lane, Piscataway, NJ 08854
email: gomer@bellcore.com

The bottlenecks between research and product development are well known. It typically takes a very long time for ideas coming out of research labs to make their way into products, and developers often face practical problems which do not seem to be addressed by currently available research results. The term "technology transfer" is often used to describe the process of overcoming these bottlenecks.

There are two aspects to technology transfer. The most commonly recognized aspect is improving the flow of information about research solutions from research groups to practitioner groups so they can use them in the development of products. Another equally important aspect is improving the flow of information about practitioner problems from practitioner groups to research groups so they can use them as the inspiration for research projects.

At the Seventh International Conference on Data Engineering in Kobe, Japan, April 8-12, 1991, a panel discussion was held on the topic of "Practitioner Problems in Need of Database Research." At this panel the panelists presented a sampling of concrete problems faced by practitioners which appear to hinge on unsolved database research problems. The immediate goal of the panel was to stimulate research on the particular problems described. Broader goals were to demonstrate that practitioner needs form a potentially rich source of inspiration for database research problems and to generate interest in the general question of how such needs can be communicated to the research community on an ongoing basis.

As a further step toward these goals, this issue of SIGMOD Record is publishing the position papers of four of the panelists:

- Kyoji Kawagoe of NEC
- Ravi Krishnamurthy of Hewlett-Packard
- David Reiner of Kendall Square Research
- Antoni Wolski of the Technical Research Centre of Finland

The fifth panelist, C. Mohan of IBM, described some problems related to supporting efficient, robust and flexible storage, maintenance and manipulation of large volumes of data (e.g. >100 gigabytes in a single table) in relational database management systems.

It strikes this author as a very healthy sign that valuable session time at a research conference and valuable pages in a research newsletter such as this one are being devoted to presenting the needs of the practitioner community. As disciplines mature, there is a dangerous tendency for them to become ingrown. It is a tribute to the database research community that this does not appear to be happening in this area.