

# User Surveys on Database Research Needs in Finland

Antoni Wolski  
Technical Research Centre of Finland  
Laboratory for Information Processing  
Helsinki, Finland  
email: wolski@tik.vtt.fi

## 1. Introduction

Our laboratory conducts research in software sciences. We are also keenly interested in technology transfer, as our work is often sponsored by industrial companies or industry related research funding organizations.

We often perform needs surveys among potential or actual recipients of our work, in connection with major research programs. During the past 6 years we have carried out two such surveys in the field of database technology. One of them was related to utilization of distributed databases in general, and the other concentrated on the requirements for distributed transactions. Some of the most interesting results are reported in the two sections below.

A few words on the nature of the respondent organizations may be useful. Enterprises of different sizes were involved. However, the most visible group consisted of heavy-weight database users such as banks, insurance companies, wholesale dealers, etc. A typical portrait of their DP-operations was the following:

- nation-wide operations run out of a single mainframe database;
- 3000 to 5000 terminals (including specialized terminals, e.g. ATMs);
- 700,000 transactions per day average load, peaking at nearly two million transactions per day.

## 2. The Distributed Database Needs Survey (1985)

Distributed databases were being highly publicized at that time. However, no workable commercial product implementation existed. The respondents were reluctant to even consider distributing their databases. The arguments we heard were clear and

crisp. They obliterated various "truths" about distributed databases. Here is why the respondents preferred to stay with a centralized solution:

- The decreasing costs and increasing quality of telecommunications made the telecommunications cost saving factor (attributed to distributed databases) insignificant.
- It was believed that the overall system reliability is much better with stringent centralized operations, including stand-by computers, back-up telecommunications links and, above all, manageable and well-drilled recovery procedures.
- In a centralized system all the investments were kept to a minimum; these were hardware (since downsizing was not considered at that time), software, site facilities, and human labour.
- Operations of a distributed and geographically dispersed organization could be successfully based on a single computer site (as was proved by the respondents in the years following the survey).
- Performance of a centralized system was acceptable.
- Commercial offerings of distributed database management systems were highly unsatisfactory at that time.

We believe all of the arguments are still holding. The commercial offerings have improved significantly since then, but they still lack robust global recovery and global management facilities.

On the other hand, our respondents were very interested in integration (or interoperability, as it is said nowadays) of the pre-existing databases which were being established in an uncontrolled way all the time in the organizations. Consequently, we dedicated our efforts to interoperability techniques in the LINDA project during the years 1985 - 1988.

### 3. The Distributed Transaction Management Survey (1989)

We were interested in the needs for distributed transactions in general and in various transaction architectures of interest to the respondents. Some distribution of data resources had happened at the respondents, driven by LAN-based workstation environments proliferating at great speed.

It turned out that the "sacred" values of transaction research were not saluted, e.g. serializability of transactions in general and one-copy serializability in case of management of replicated data.

The need for system-supported serializability was recognized, but only for simple, "canned" transactions. The need to support heterogeneous system environments was also recognized. These requirements were in line with our work on preserving serializability in heterogeneous database environments in the HERMES project during the years 1988 - 1991.

However, in the case of complex (or interactive) transactions serializable execution was not seen as attractive because of the penalty of low concurrency and throughput. This was considered a problem even in a centralized system. In a distributed database the problem became even more severe because of the impact the two-phase commit protocol had on the transaction duration time.

The respondents achieved correctness of interleaved executions in their current high-performance transaction processing applications by making use of the semantics of the transactions in the application code. In most cases this enabled them to execute commutative operations without setting long-term locks. Instead of having to code applications in such a way, they would like to have a general-purpose system for such "semantic" transaction management. Unfortunately, nobody is really eager to deliver anything like that, and there have been only a few research papers on the subject.<sup>[1][2]</sup>

Another area of divergence from the directions of the research community was system support for copies. Let us note that to our knowledge no commercial DBMS supports copies at this time. The researchers who have produced various copy

control methods have concentrated on preserving one-copy serializability. However, one-copy serializability need not be the primary objective of copy support, as illustrated by the following case study of a bank system:

- There is a central (master) database in the system.
- Overlapping snapshots of the central database are located in LAN-based database servers at branches.
- Transactions at branches are run exclusively with the local snapshot data.
- The transactions update the snapshots.
- The transactions are pre-programmed and known in advance.
- The snapshot updates are merged to the central database periodically.

One can see that the problem is that of database snapshot update. It can not be solved if the one-copy serializability is assumed. On the other hand, some semantic correctness criteria can be met if the semantics of the transactions are limited, and if the information thereof is available to the system. We know of only two research papers in this area, and the bank is still waiting for the solution. In the meantime a scheme has been employed whereby the update transactions are directed to the master database.

#### References:

- [1] Kumar, A., and Stonebraker, M., "Semantics Based Transaction Management Techniques for Replicated Data," *Proc. 1988 ACM SIGMOD Conf.* (Chicago, 1 - 3 June, 1988), 117-125.
- [2] Pu, C., and Leff, A., "Replica Control in Distributed Systems: An Asynchronous Approach," *Proc. 1991 ACM SIGMOD Conf.* (Denver, 29 - 31 May, 1991), 377-386.