

Invited Talk

Information Management and System/Storage Technology — Evolution or Revolution?

Speaker Berni Schiefer (IBM Toronto)

Abstract The role and design of systems responsible for the management of information in the enterprise is changing. The kind of information that is being managed is changing as is the way the information is analyzed and made available to users in the enterprise. At the same time system and processor technology is undergoing what some consider a fundamental shift as processor designers grapple with the power/heat issues associated with ever higher frequencies. In this talk we will review the forces that are influencing both the software and the hardware systems.

Biography Berni Schiefer is a DB2 Distinguished Engineer at IBM. He has responsibility for DB2 performance benchmarking and solutions development, including the BCU. He joined the IBM Toronto Lab in 1985 and has worked on SQL/DS and the Starburst experimental relational database at the IBM Almaden Research Lab, prior to working on DB2. His current focus is on introducing advanced technology into DB2 with particular emphasis on processors, performance, XML, Linux, Virtualization and Autonomics.

Panel Discussion

ManyCore-DB: Will you still need me, will you still feed me, when I'm 64?

Heat issues have turned the CPU industry on the path of producing multi-core chips, now doubling the amount of cores per chip in each new manufacturing process (a 2-year cycle). Mainstream CPUs became dual core in 2005 and in 2007 will become quad-core. As downscaling the chip manufacturing process (the driver for Moore's law) is predicted to remain on track, this means that within 10 years, mainstream CPUs will provide huge amounts of cores (64 or more).

This trend is a big gamble by the computer architects on the ability of software architects to efficiently use all these cores, and thus provides a significant challenge to the entire software industry to adopt programming models and software architectures that are amenable to massive parallelism. In many areas of today's software, the amount of parallelism that can be extracted is very limited, and certainly does not scale to 64.

Thus, ubiquitous and mandatory parallelism finally becomes real. This panel is about the massive parallelism challenge of the coming decade to the database community. In fact, if the database community successfully rises to this challenge, the database paradigm might gain significantly in importance in the software industry.

The panel has participants from both computer architects and database architects. Participants are encouraged to be outspoken and where possible controversial in their views, allowing for a lively exchange of arguments (as seen in the previous DaMoN panel).

To structure the discussion, we would like the panelists to outline their views on the following questions:

- Is it indeed true, as sketched in the panel summary, that mainstream computers within 10 years will have ≥ 64 cores?
- How will the rest of the computer architecture scale up with this (e.g. adaptive/configurable caches, or even programmable on-chip memory, massive bandwidth RAM, transactional RAM and the corresponding trends in I/O infrastructure)?
- Are the parallel features in current DBMS's sufficient to exploit these new resources? If not, what areas of database architecture will be affected? Would we even need to revisit database architectures from scratch?
- Is it indeed true, as suggested in the panel summary, that massive parallelism provides a chance for the database community, to further increase its importance in the software industry? As a follow-up, which IT application areas would benefit most from basing their processing infrastructure on parallel data management technology?

Panelists

Goetz Graefe	(Microsoft Research)
Mark Hill	(University of Wisconsin, Madison)
Berni Schiefer	(IBM Toronto)
Todd Walter	(NCR/Teradata)