

# Paradigm Shift to New DBMS Architectures: Research Issues and Market Needs

Panel Organizer:

Sang K. Cha  
Seoul National University & Transact In Memory, Inc.  
[chask@kdb.snu.ac.kr](mailto:chask@kdb.snu.ac.kr)

Participating Panelists:

Anastassia Ailamaki  
CMU  
[natassa@cmu.edu](mailto:natassa@cmu.edu)

Yoshinori Hara  
NEC  
[y-hara@da.jp.nec.com](mailto:y-hara@da.jp.nec.com)

Vishal Sikka  
SAP  
[vishal.sikka@sap.com](mailto:vishal.sikka@sap.com)

## Abstract

*Moore's law has driven CPU power and memory capacity to grow million times since the system R and Ingres projects started thirty years ago. The underlying software technology has also changed substantially. Today, operating systems support POSIX lightweight multithread library and virtually infinite address for efficient utilization of multiprocessor systems with very large memory.*

*Despite these dramatic advances in underlying hardware and software, the initial RDBMS architecture of managing data and indexes as disk-resident block structures remains the same. The heavyweight process architecture is still dominant, incurring costly context switch overhead among multiple processes involved in transaction execution. Typical commercial RDBMS implementations involve several millions of lines of complex code that has been evolving over decades. Because it is extremely risky to overhaul any software of this size, commercial RDBMS implementations are likely to maintain the current, disk-resident heavyweight-process architecture.*

*On the application side, there is growing demand for real-time acquisition and analysis of a large volume of data, especially, continuously arriving stream data. Examples are traditionally found in financial services, telecom, defense and intelligence, logistics, and this list is being expanded to include other domains such as supply chain and retail with the development of RFID technology for ubiquitous tracking of physical objects. These so-called real-time enterprise applications demand*

*orders of magnitude higher performance scalability than existing RDBMS implementations can manage.*

*Recent research addressed the impedance mismatch of these existing RDBMS implementations with modern hardware architecture. For example, latch coupling, a common technique for concurrency control of shared index nodes incurs excessive coherence L2 cache misses on shared-memory multiprocessors. L2 cache misses are expensive, costing the loss of a thousand instruction opportunity as the gap between CPU and memory speeds grows exponentially. Other types of L2 cache misses occur because disk-resident data and index structures are not optimally laid out for CPU processing. With this impedance mismatch, there is a limit in improving the database performance by just growing the hardware capability.*

*To exploit the processing power of modern hardware architecture and to meet the database performance demands of challenging applications, it is necessary to reexamine the underlying premises of the current RDBMS architecture and validate them against the characteristics of modern hardware.*

*The objective of this panel is to review the current state of research toward building new DBMS architectures with key researchers in this area, discuss market needs and requirements with industry leaders experienced in applications, and explore future research directions to promote the research on new DBMS architectures within the database community.*