

Keynote Talk

Database Systems in the Multicore Era

Anastasia Ailamaki

School of Computer and Communication Sciences, EPFL
Lausanne, Switzerland
natassa@epfl.ch

Abstract

Database systems have long optimized for parallel execution; the research community has pursued parallel database machines since the early '80s, and several the key ideas from that era underlie the design and success of commercial database engines today. Computer architectures have shifted drastically during the intervening decades, however, and today the constraints of semiconductor technology combine with Moore's Law to double the number of processors per chip every 18 months. Converting this available raw parallelism into scalable performance is increasingly difficult with conventional servers, for both business intelligence and transaction processing workloads.

This talk analyzes database performance scaling results on future chip multiprocessors and demonstrates that current parallelism methods are insufficient and of bounded utility as the number of processors per chip exponentially increase. Common sense is often contradicted; for instance, the effect of using larger and slower on-chip caches may be detrimental to the absolute database performance. To achieve scalability for database applications on chip multiprocessors, major rethinking of the database storage manager is necessary. First, concurrency needs to be converted into parallelism – a challenging task, even for database systems. Then, parallelism needs to be extracted from seemingly serial operations; extensive research in distributed systems proves to be very useful in this context. At the query processing level, service-oriented architectures provide an excellent framework to exploit available parallelism. I will use the StagedDB/CMP and ShoreMT projects at EPFL as examples to outline the above research directions.

Categories & Subject Descriptors: H.2.4 [Database Management]: Systems—Concurrency, Query Processing, Parallel Databases, Transaction Processing.

General Terms: Performance, Design, Experimentation.

Bio

Anastasia (Natassa) Ailamaki is a Professor of Computer Sciences at the Ecole Polytechnique Federale de Lausanne (EPFL) in Switzerland. Her research interests are in database systems and applications, and in particular (a) in strengthening the interaction between the database software and the underlying hardware and I/O devices, including flash technology, and (b) in automating database design and computational database support for scientific applications. She has received a Finmeccanica endowed chair from the Computer Science Department at Carnegie Mellon (2007), a European Young Investigator Award from the European Science Foundation (2007), an Alfred P. Sloan Research Fellowship (2005), six best-paper awards at top conferences (2001-2006), and an NSF CAREER award (2002). She earned her Ph.D. in Computer Science from the University of Wisconsin-Madison in 2000.