

Evolving the Architecture of a DBMS for Modern Hardware

Paul Larson

Microsoft Research
Redmond, WA 98052-6399
Paul.Larson@microsoft.com

Abstract

The major commercial database systems were designed primarily for OLTP workloads and under the assumption that processors are slow, memory is scarce, and data lives on disk. These assumption are no longer valid: OLAP workloads are now as common as OLTP workloads, multi-core processors are the norm, large memories are affordable, and frequently accessed data lives mostly in the main memory buffer pool. So how can a vendor with a mature DBMS product exploit the opportunities offered by these changes? Rewriting from scratch is not realistic - it is way too expensive and risky. The only realistic option is to gradually evolve the architecture of the system. SQL Server has begun this journey by adding two features: column store indexes to speed up OLAP-type queries, and Hekaton, a new engine optimized for large memories and multicore processors. The talk will outline the design of these features, the main goals and constraints, and discuss the reasoning behind the design choices made.

Biography: Paul (Per-Ake) Larson has conducted research in the database field for over 30 years. He served as a Professor in the Department of Computer Science at the University of Waterloo for 15 years and joined Microsoft Research in 1996 where he is a Principal Researcher. Paul has worked in a variety of areas: file structures, materialized views, query processing, and query optimization among others. During the last few years he has collaborated closely with the SQL Server team on how to evolve the architecture of the core database system.