

Enhanced Disk-Based Databases towards Improved Hybrid in-Memory Systems

Authors : Samuel Kaspi, Sitalakshmi Venkatraman

Abstract : In-memory database systems are becoming popular due to the availability and affordability of sufficiently large RAM and processors in modern high-end servers with the capacity to manage large in-memory database transactions. While fast and reliable in-memory systems are still being developed to overcome cache misses, CPU/IO bottlenecks and distributed transaction costs, disk-based data stores still serve as the primary persistence. In addition, with the recent growth in multi-tenancy cloud applications and associated security concerns, many organisations consider the trade-offs and continue to require fast and reliable transaction processing of disk-based database systems as an available choice. For these organizations, the only way of increasing throughput is by improving the performance of disk-based concurrency control. This warrants a hybrid database system with the ability to selectively apply an enhanced disk-based data management within the context of in-memory systems that would help improve overall throughput. The general view is that in-memory systems substantially outperform disk-based systems. We question this assumption and examine how a modified variation of access invariance that we call enhanced memory access, (EMA) can be used to allow very high levels of concurrency in the pre-fetching of data in disk-based systems. We demonstrate how this prefetching in disk-based systems can yield close to in-memory performance, which paves the way for improved hybrid database systems. This paper proposes a novel EMA technique and presents a comparative study between disk-based EMA systems and in-memory systems running on hardware configurations of equivalent power in terms of the number of processors and their speeds. The results of the experiments conducted clearly substantiate that when used in conjunction with all concurrency control mechanisms, EMA can increase the throughput of disk-based systems to levels quite close to those achieved by in-memory system. The promising results of this work show that enhanced disk-based systems facilitate in improving hybrid data management within the broader context of in-memory systems.

Keywords : in-memory database, disk-based system, hybrid database, concurrency control

Conference Title : ICECECE 2015 : International Conference on Electrical, Computer, Electronics and Communication Engineering

Conference Location : Singapore, Singapore

Conference Dates : January 08-09, 2015