

Evaluation of OLTP deployments in multsocket multicores

Keywords: OLTP, multsocket, multicore, shared-everything, shared-nothing, Islands

Problem: Two are the main deployment strategies for scaling up the performance of online transaction processing (OLTP) systems; the shared-everything and the shared-nothing one. The shared-everything configurations face significant performance challenges since their execution is full of critical sections and, unfortunately, even the smallest of serializations in the execution may have significant impact in scalability [1]. On the other hand, the problems with the shared-nothing approach (where a set of independent database instances operate collectively in order to serve the requests [2]) have to do with load balancing and the need to execute distributed transactions. H-Store takes the shared-nothing approach to the extreme by employing single-threaded shared-nothing instances which have very high single-instance performance because locking, latching, and logging are disabled [3].

The debate between the supporters of the two camps is on-going. Both designs have pros and cons and the decision depends heavily on the workload and the application requirements. The problem gets even more complicated as modern server hardware becomes increasingly heterogeneous. For example, the modern hardware where OLTP systems are called to execute very efficiently on consists of multsocket machines of multicore processors. This study attempts a thorough comparison of the two approaches in modern hardware configurations and proposes a design point called Islands. The Islands design is based on the deployment of (shared-everything) instances at the socket level.

Project: The project consists of three phases. The first phase consists of the implementation of an evaluation framework for shared-everything and flexible shared-nothing configurations on top of a reliable state-of-the-art storage engine. For this project the student will use and extend the Shore-MT storage engine [3] as well as shore-kits, the benchmarking application which is built on top of Shore-MT. The second phase consists of a thorough performance evaluation of the different deployment strategies on multsocket multicore environments in a variety of transactional workloads. Finally, all the findings and conclusions will be summarized in a performance analysis document.

Plan:

1. Extension of shore-kits and Shore-MT in order to reliably run shared-nothing configurations. The student will need to make sure that distributed transactions execute reliably on top of Shore-MT.
2. Implementation of a representative set of transactional workloads for comparing shared-everything and shared-nothing configurations. Except from the traditional benchmarks (such as TAPT, TPC-B, TPC-C and TPC-E), there should be workloads where the user can modify

DIAS: Data-Intensive Applications and Systems Laboratory

School of Computer and Communication Sciences

Ecole Polytechnique Fédérale de Lausanne

Building BC, Station 14

CH-1015 Lausanne

URL: <http://dias.epfl.ch/>

- among others the percentage of distributed transactions, the distribution of requests, the selectivity of queries and the number of partitions (of the shared nothing configurations).
3. Modifications of Shore-MT in order to be able to switch off latching and locking in order to achieve higher single-thread performance, comparable with that of H-Store.
 4. Implementation of the same workloads for VoltDB [5], the commercial version of H-Store which is available publicly.
 5. Thorough performance evaluation in mutlisocket multicore hardware. There are three target machines: a 4-socket machine with 8-core/64-contexts Sun Niagara T2 processors, a 2-socket machine with quad-core AMD Opterons, and a 4-socket machine with six-core Intel Nehalems.
 6. A document that presents the performance study and analyzes the findings.

Supervisor: Prof. Anastasia Ailamaki, anastasia.ailamaki@epfl.ch
Responsible collaborator(s): Ippokratis Pandis (ippokratis.pandis@epfl.ch), Ryan Johnson (ryan.johnson@epfl.ch)
Duration: 4 months