How to Stop Under-Utilization and Love Multicores

(Tutorial)

Danica Porobic
École Polytechnique Fédérale de Lausanne
danica.porobic@epfl.ch

ABSTRACT

Hardware trends oblige software to overcome three major challenges against systems scalability: (1) taking advantage of the implicit/vertical parallelism within a core that is enabled through the aggressive micro-architectural features, (2) exploiting the explicit/horizontal parallelism provided by multicores, and (3) achieving predictively efficient execution despite the variability in communication latencies among cores on multisocket multicores. In this tutorial, we shed light on the above three challenges and survey recent proposals to alleviate them. The first part of the tutorial describes the instruction- and data-level parallelism opportunities in a core coming from the hardware and software side. In addition, it examines the sources of under-utilization in a modern processor and presents insights and hardware/software techniques to better exploit the microarchitectural resources of a processor by improving cache locality at the right level of the memory hierarchy. The second part focuses on the scalability bottlenecks of database applications at the level of multicore and multisocket multicore architectures. It first presents a systematic way of eliminating such bottlenecks in online transaction processing workloads, which is based on minimizing unbounded communication, and shows several techniques that minimize bottlenecks in major components of database management systems. Then, it demonstrates the data and work sharing opportunities for analytical workloads, and reviews advanced scheduling mechanisms that are aware of non-uniform memory accesses and alleviate bandwidth saturation.

Biography

Danica Porobic is a final year PhD student working under the supervision of Professor Anastasia Ailamaki in Data-Intensive Applications and Systems (DIAS) Laboratory at EPFL. Her research focuses on designing scalable transaction processing systems for non-uniform hardware platforms. She has graduated top of her class with MSc and BSc in Informatics from University of Novi Sad and has worked at Oracle Labs and Microsoft SQL Server.