

Building a domain-knowledge guided system software environment to achieve high-performance of multi-core processors

Xiaodong Zhang

The Ohio State University, USA

Although multi-core processors have become dominant computing units in basic system platforms from laptops to supercomputers, software development for effectively running various multi-threaded applications on multi-cores has not made much progress, and effective solutions are still limited to high performance applications relying on existing parallel computing technology. In practice, majority multi-threaded applications are highly concurrent programs demanding high throughput, such as concurrent database transactions, massive and independent query requests in Web servers and search engines, and executing many-tasks for scientific applications in a multiprogramming mode. To best utilize the increasingly rich computing and cache resources in multi-core processors (many-cores in the near future) we must address several serious and difficult challenges. First, there are several critical hardware resources for multi-threads to share, such as the last level caches and the memory buses. However, the shared resource management is Largely controlled by hardware. Second, OS scheduler has little knowledge about applications' data demanding and access behavior, making sub-optimal task assignment decisions. Finally, the space allocation for each thread in the shared-cache is demand-based, often causing access conflicts and pollution, significantly degrading overall execution performance.

We have developed a runtime environment connecting multi-threaded applications and system software in a collaborative way, where operating system is guided by application domain knowledge including data access locality and execution behavior to schedule tasks and allocate shared hardware resources for each running thread. We evaluate our environment by concurrent database transactions and multi-threaded scientific computing programs, and show strong performance and Throughput improvement by minimizing cache conflicts and misses in the last level caches of multi-cores. We further develop our system as a general framework to automatically manage multi-threaded applications on multi-core processors.