NSF PROJECT ABSTRACT

Power-efficient Multi-core Scheduling for Network Applications

High-speed multicore computers are widely available at low cost. They have the potential to process high speed network applications, such as deep packet inspection, and multi-channel real-time streaming of high-definition audio and video. However, they need better scheduling approaches to address the demand of high throughput and low latency of these applications while reducing energy consumption. Current approaches do not consider real-time demands, network traffic patterns and details of the specific multicore architectures. Lack of power-aware core scheduling and thermal management has also led to a great challenge in energy efficiency. Reducing power consumption saves energy cost, reduces the need for building additional cooling, and improves computer hardware reliability.

This project makes fundamental contributions to the design and applications of high-demand network systems. It provides packet level scheduling solutions that improve throughput, reduce latency and energy consumption. The project develops and incorporates network traffic models, architectural details, power consumption models, and thermal models simultaneously for coordinated workload processing and energy consumption. Quality-of-service (QoS) of packet streams/connections is also considered by improving the jitter and out-of-order packet departures. Superiority of the proposed technical approaches is verified through real implementation of network applications and performance and power measurements on commercial multicore architectures. An open-source framework is developed for multicore scheduling, and is made available to the researchers along with optimized multithreaded codes for network applications.