Performance Impact of Resource Contention in Multi-Core Systems

High-End Computing

A large percentage of NASA's modeling and simulation work is performed on supercomputing clusters using commodity multi-core processors. While the clusters offer cost-effective peak performance, contention for resources in the memory hierarchy can adversely impact the realized performance. The purpose of our work is to understand how the performance of NASA application codes depends on specific processor architectures.

Resource sharing in commodity multi-core processors, such as the Intel Xeons found in the Pleiades supercomputer at the NASA Advanced Supercomputing facility, can have a significant impact on the performance of production applications. In work done at NASA Ames Research Center, we have used a differential performance analysis methodology to quantify the costs of contention for resources in the memory hierarchy of several multi-core processors used in high-end computers.

In particular, by comparing runs that bind Message Passing Interface (MPI) processes to cores in different patterns, we were able to isolate the effects of resource sharing, and determine that memory bandwidth to a processor socket was the dominant contention factor. This was determined by testing with three benchmarks and four NASA applications on four multi-core platforms.

These results help further our understanding of the specific demands these codes place on their production environments, and each system's ability to deliver performance. This understanding benefits both code optimization efforts and future system procurements.

Robert Hood, Johnny Chang, NASA Ames Research Center
robert.hood@nasa.gov, johnny.chang@nasa.gov

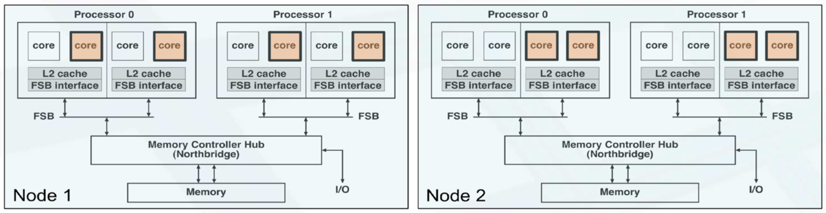

The difference in performance between the two node configurations shown here (where used cores are shown in orange) is due solely to the increased sharing of the L2 cache in Node 2. *Robert Hood, Henry Jin, NASA/Ames*