Opportunities for Integrating Tasking and Communication Layers

Dylan T. Stark* and Brian W. Barrett
Sandia National Laboratories
*dstark@sandia.gov

Abstract

This talk will discuss efforts toward system software composition for current and future HPC architectures in the context of requirements and opportunities for a new Chapel compiler target. HPC processor architectures over the last decade evolved from monolithic processing cores to an ecosystem comprised of multi- and many-core chips and accelerators, along with asymmetric memory systems, and potential heterogeneity in ISA and/or performance characteristics. Likewise, HPC interconnects have adapted to optimize communication patterns involving smaller, more frequent, messages over bursty bulk data transfers. Chapel is well-positioned to deal with the wealth and uncertainty of architectural options with a global-view tasking language, compiler infrastructure, and modular runtime back end supporting third-party libraries for performance-critical runtime components. However, the current modular back end comprises independent components for parallelism, communication, and memory management that may limit overall performance.

We contend that there are opportunities to improve performance and scalability of Chapel by designing towards a more unified runtime system. Over the last year, we have been developing a unified runtime system called the Scalable Parallel Runtime (SPR) to study the integration of task parallelism and messaging. The SPR integrates the Qthreads light-weight threading library and the Portals 4 network API to provide a tightly-coupled solution for high-performance large-scale distributed-memory tasking applications. Qthreads is already a Chapel tasking layer target, and a new communication layer has been prototyped to use Portals 4. This provides us with the opportunity to evaluate the requirements on the Chapel compiler infrastructure for supporting a unified runtime system interface, such as the one presented by the SPR. Our evaluation includes analysis of a set of Chapel benchmarks to characterize application execution patterns and call pressure through the runtime interface to give greater insight into runtime system composition, Chapel compiler and the runtime interaction, and compiler and runtime system co-design for current and future HPC systems.

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.