



Enhancing Qthreads for ECP Science and Energy Impact And Sandia ATDM* On-Node Runtime Coordination



Ron Brightwell and Stephen Olivier, PIs

What is Qthreads? Qthreads is a library that provides lightweight multithreading for nimble locality-aware **on-node task parallelism**, exposing this capability through high-level programming models such as the **Kokkos** C++ performance portability library and Cray's **Chapel** language, and via tech transfer to standards like **OpenMP**.

Qthreads at a Glance

Development	2007 - Present
License	Three-clause BSD
Repository	http://github.com/Qthreads

What enhancements are proposed and what is the impact? The project's objective is to develop techniques that improve **network concurrency** for applications that use multithreading coupled with communication, e.g., **MPI+X**. The impact is more efficient use of network and node resources to deliver better **performance** for ECP applications.

What are the first steps to achieving these goals? The initial focus is understanding **use cases** for applications to benefit from better network concurrency, extending earlier work on **task parallel decomposition (TPOD)** [1] within each MPI process of MPI+X programs [2]. Evaluating the performance gap due to serialization of threads' network accesses will guide the development of **requirements** to improve network concurrency for threaded applications.

What is on-node runtime coordination and why is it needed? On-node runtime coordination provides capabilities for the **management of node-level resources**, e.g., cores, threads, and memory. SNL's **ATDM*** software stack uses a component-based architecture, so these resources must be managed to ensure that each component has its own resources to use for the work it needs to do and that performance-killing **oversubscription** of resources is avoided.

* ATDM = Advanced Technology Development and Mitigation, NNSA's Advanced Simulation and Computing program effort within ECP to prepare for future computing systems for stockpile stewardship.

[1] Barrett et al. "Toward an Evolutionary Task Parallel Integrated MPI+X Programming Model." *6th Intl. Workshop on Programming Models and Applications for Multicores and Manycores (PMAM 2015)*.

[2] Stark et al. "Early Experiences Co-Scheduling Work and Communication Tasks for Hybrid MPI+X Applications." *Proc. of 2014 Workshop on Exascale MPI (ExaMPI 2014) at Supercomputing 2014 (SC14)*.