Toward HPC Runtimes that Self-adapt to your Math

Dorian Arnold, University of New Mexico

The continued evolution of high-performance computing (HPC) systems will result in systems comprised of more complex hardware and software components, and compared to today's high-end systems, exascale class systems will be orders of magnitude larger (in terms of node and core counts). To help address the challenge of exploiting massive degrees of parallelism and concurrency that arise in such environments, I propose a self-adaptive framework for extreme scale tools and applications. The proposed framework integrates mechanisms for (1) self-monitoring of dynamic environmental characteristics including resource availability and workload resource demands; (2) self-detection of functional and performance problems; (3) decision processing for evaluating corrective actions; and (4) instantiation of decided actions.

In this talk, I motivate the proposed framework using our research and experiences in the HPC tool domain. Then I describe how this approach may also be useful for scientific applications and their underlying mathematical algorithms -- in an attempt to provoke feedback and discussions for potential application and algorithm-based feasibility and use cases.