The Software of the Future on the Hardware of the Future

Jeff R. Hammond, William R. Scullin and Raymond A. Bair

As the computational science community transitions from petascale to exascale, one of the paradigm shifts expected in software engineering of scientific codes is the need for co-design -- the cooperative interaction of hardware, system software, numerical algorithms, data analysis, etc. -- which presents significant challenges for domain scientists. In this talk, we analyze the cooperative interactions of computational chemistry codes with hardware, operating systems, programming models and communication primitives on Blue Gene/P based upon our experience with NWChem, MPQC, MADNESS, and other packages at the Argonne Leadership Computing Facility. Hybrid programming models are critical in all of these codes in order to achieve asynchronous communication and make optimal use of the limited memory-per-node. Specifically, efficiently scaling NWChem to thousands of nodes required changes to both computational kernels and the communication system, exploiting the unique features of the operating system and the network hardware. Based on our experiences with chemistry codes at the petascale, we discuss possible co-design strategies for these applications on the road to exascale.