

Petaflop/s, Seriously

Prof. David E. Keyes

Fu Foundation Professor of Applied Mathematics in the Department of Applied Physics and Applied Mathematics
Columbia University

Sustained floating-point rates on real applications, as tracked by the Gordon Bell Prize, have increased by over five orders of magnitude from 1988, when 1 Gigaflop/s was reported on a structural simulation, to 2006, when 200 Teraflop/s were reported on a molecular dynamics simulation. Various versions of Moore's Law over the same interval provide only two to three orders of magnitude of improvement for an individual processor; the remaining factor comes from concurrency, which is of order 100,000 for the BlueGene/L computer, the platform of choice for the majority of recent Bell Prize finalists. As the semiconductor industry begins to slip relative to its own roadmap for silicon-based logic and memory, concurrency will play an increasing role in attaining the next order of magnitude, to arrive at the long-awaited milestone of 1 Petaflop/s sustained on a practical application, which should occur around 2009. Simulations based on Eulerian formulations of partial differential equations can be among the first applications to take advantage of petascale capabilities, but not the way most are presently being pursued. Only weak scaling can get around the fundamental limitation expressed in Amdahl's Law and only optimal implicit formulations can get around another limitation on scaling that is an immediate consequence of Courant-Friedrichs-Lewy stability theory under weak scaling of a PDE. Many PDE-based applications and other lattice-based applications with petascale roadmaps, such as quantum chromodynamics, will likely be forced to adopt optimal implicit solvers. However, even this narrow path to petascale simulation is made treacherous by the imperative of dynamic adaptivity, which drives us to consider algorithms and queueing policies that are less synchronous than those in common use today. Drawing on the SCaLeS report (2003-04), the latest ITRS roadmap, some back-of-the-envelope estimates, and numerical experiences with PDE-based codes on recently available platforms, we will attempt to project the pathway to Petaflop/s for representative applications.