(Some) Answers to the Petascale Challenge

Dr. Richard Loft
Computational and Information Systems Laboratory
NCAR

Abstract

Over the past 30 years, the use of supercomputers to obtain numerical solutions of the equations governing weather and climate, along with satellite-based global observations, have led to the realization that the Earth is characterized by a number of complex, highly interactive phenomena that must be understood as a complete system. The problem is complicated by the fact that the sub-components of the Earth system interact on a wide range of spatial and temporal scales that require century long simulations and high resolutions to capture.

The anticipated availability of massively parallel petascale (and, yes, exescale) computers in the next few years offers the climate community a golden opportunity to dramatically advance our understanding of the Earth's climate system and climate change, if they can be harnessed to the task. The fit is not perfect. First, massively parallel systems will impose stringent and unavoidable Amdahl-law requirements on application scalability. Second, the trade-off between resolution and integration rate, both critical factors in climate modeling, are severe. Third, the increasing complexity of petascale systems, e.g. in terms of the numbers of cores on a chip, increases the tension between the system architecture and programmability. Finally, the size and complexity of climate applications make them difficult to port, adapt, and validate on new architectures.

It is becoming increasingly clear that these challenges are fundamental and far-reaching: a business as usual approach extrapolated from existing technologies and algorithmic approaches will likely not be sufficient. This talk will attempt to outline potential sources of algorithm and hardware acceleration necessary on a number of fronts to meet these challenges.