


<http://cra.org/ccc/theimpactofnitrd>

## More and Moore: Growing computing Performance for Scientific Discovery

Katherine Yelick

University of California, Berkeley & Lawrence Berkeley National Laboratory

The past two decades of growth in computing performance have enabled breakthroughs across many disciplines, from biology to astronomy, and from engineering to fundamental science. Within the energy space, computation is used to develop alternative energy sources, to design energy efficient devices, and to understand and mitigate the impacts of energy choices. Demand for computational capability across science grows unabated, with the growth coming from three general areas:



- **Scale:** Computational models have become more detailed and sophisticated, but future simulations will combine even more physical phenomenon interacting in more complex ways. For example, combustion simulations of low emissions devices currently use petascale simulations, but capturing the chemical reactions to simulate alternative fuels will require more detailed exascale simulations.
- **Volume:** Science and engineering disciplines use simulations to screen through a large set of possible designs to narrow down the most promising solutions. This can make companies more competitive by reducing the time from concept to product, e.g., one project looks through tens of thousands of simulations from potential battery materials, while another examines the dynamics of tens of thousands of proteins to understand Alzheimer's and other diseases. Orders of magnitude more computing are needed to handle realistic environments and a broader set of applications.
- **Data:** Simulations run on petascale systems produce petabytes of scientific data, but those volumes are matched or exceeded by the data coming from other experimental and observational devices such as sequencers, detectors, and telescopes. These massive data sets will drive the need for increased computing capability, as the scientific process requires automated filtering, analysis, visualization and information retrieval.

Two decades ago, the field of high performance computing was facing a major crisis, as the vector supercomputers of the past were no longer cost effective. Microprocessor performance increases were tracking Moore's Law, making massively parallel systems the cost effective choice for science. But this hardware revolution came with significant challenges for algorithm and software developers. Funded in large part by Networking and Information Technology Research and Development (NITRD) Program investments, researchers developed new software tools, libraries, algorithms, and applications and adapted to that change, while also increasing the capability and complexity of the problems being simulated.



Today, the field again faces a major revolution in computer architecture. The clock speed benefits of Moore's Law have ended, and future system designs will instead be constrained by power density and total system power, resulting in radically different architectures. The challenges associated with continuing the growth in computing performance will require broad research activities across computer science, including the development of new algorithms, programming models, system software and computer architecture. While these problems are most evident at the high end, they limit the growth in computing performance across scales, from department-scale clusters to major computational centers.