SuperComputers: Super-Polluters?

In a word: no. Thanks to imperatives for limiting waste heat, maximizing performance, and controlling operating cost, energy efficiency has been a driving force in the evolution of supercomputers. The challenge going forward will be to extend these gains to offset the steeply rising demands for computing services and performance.

In addition, supercomputers also run the models critical to understanding the greatest scientific and environmental challenges of our time – such as global climate change. High-fidelity, kilometer-scale models of weather and climate change require 1000-times more computing performance than the largest supercomputers available today to provide answers to questions that have multi-trillion dollar consequences.

As supercomputing moves into the petaflop era (1015 flops, or FLoating point Operations Per Second), improved energy efficiency and renewable energy sources provide the key to increasing computing power without exacting and undue cost on the environment or curtailing computing power.[1] Indeed, a new generation of best practices have already been identified and implemented in actual facilities.

DRIVING FORCES

As the field of scientific computing matures, the demands for computational resources are growing at a rapid rate. A number of studies from the scientific community, such as the 2004 SCaLeS report[2] and the 2008 DOE E3 report[3], have estimated that by the end of this decade, numerous mission-critical applications will have computational requirements that are at least two orders of magnitude larger than current levels.

Since processor performance has ceased to grow at historical rates, the energy requirements to meet the world's growing appetite for supercomputing have risen sharply. We are entering an era where petaflop high-performance computing (HPC) systems are anticipated to draw prodigious amounts of electrical power. For example, the current 19,320-processor flagship HPC system at NERSC draws less than 2 MW to deliver 100 teraflop (1012 flop) peak performance, while its successors in 2010 are projected to draw as much as 15–20 MW if fully configured. More alarmingly, the DOE E3 report projects an exaflop (1018 flop) HPC system requiring over 130 MW of power. At the average U.S. electricity generation mix, and assuming, for discussion, 8760 hours per year of operation, these levels of power translate to 10,000 and 700,000 tones of carbon-dioxide emissions annually by these individual facilities.4 The corresponding annual energy expenditures, at $0.10/kWh would be $1.7 million and $112.7 million. These values are for the IT equipment only, i.e. excluding cooling.

Operating costs are bound to rise even more quickly than energy use. Oil prices have quadrupled over the past six years, breaching $100 per barrel for the first

[1] Wehner, M, L. Oliker, and J. Shalf. 2008. "Towards Ultra-High Resolution Models of Climate and Weather" International Journal of High Performance Computing Applications (forthcoming).

[2] SCaLeS: "A Science-based Case for Large-Scale Simulation" (http://www.pnl.gov/scales/)

[3] E3: "Simulation and Modeling at the Exascale for Energy, Ecological Sustainability, and Global Security" (http://www.er.doe.gov/ascr/Misc/Energy-ecology-security-initiative.pdf)

[4] US Energy Information Administration, average U.S. emissions, 1998-2000.

time in history.[5] Electricity prices have also skyrocketed. An energy bill of $100-$200 per square foot per year–100-times that of the typical office building—would not be unusual for a contemporary supercomputing facility.

Meanwhile, rising cooling demand is driving the up-front capital costs of cooling equipment towards levels eclipsing those of the IT equipment itself. A recent survey found that 42% of conventional data centers expect to run out of cooling capacity within one to two years.[6] The lifecycle cost of power will exceed the purchase costs of such systems, and, if unchecked, will ultimately limit the practicality of future state-of-the-art HPC platforms. Every dollar spent on paying energy bills is one less dollar available for doing the important work of supercomputing; just another reason to seek efficiency gains.

Greenhouse-gas emissions ("GHGs") are in the fore as climate change looms. GHGs have also taken on new importance after a recent U.S. Supreme Court finding that they should be classified as "pollutants" under the Clean Air Act. A mandatory "cap-and-trade" system for carbon emissions in fact already exists in Europe (with $30 billion in transactions in 2007). Most analysts agree that a similar system will be established in the U.S. irrespective of the outcome of the coming presidential elections. With an eye towards removing uncertainty and managing regulatory risk, through the U.S. Climate Action Partnership (http://www.us-cap.org/) many large industries are in fact demanding mandatory caps. USCAP's members include blue-chip companies like BP America, Conoco Philips, Dow Chemical, Ford Motor Company, GE, GM, Johnson and Johnson, and Xerox corporation.

Taken together, these trends will lead to a crisis in HPC in the not-too-distant future, unless vendors and the scientific community work aggressively to develop more power-efficient solutions.

IMPLEMENTING BEST PRACTICES

Many of the efficiency strategies applicable to conventional data centers can be applied to supercomputers, with the exception of some powering-down options for the computers themselves. A list of the most important measures includes:

- Free Cooling - Free cooling is the use of outside air- or water-side cooling via cooling towers only. It allows the facility to turn off compressor systems and save energy. This option works during colder months or cold nights. In most climates, it is effective at least half the time.
- Relaxing environmental conditions - Most centers are overcooling or providing unnecessary humidity control. Taking advantage of the recommended and allowable ranges of environmental conditions setby ASHRAE represents a low cost way to save significant amounts of energy.
- Moving to liquid cooling - Having the HPC community move to direct liquid cooling solutions holds great promise to drastically reduce energy. Liquid cooling solutions that operate with higher temperature chilled water could greatly reduce or eliminate the need for compressor cooling (use of chillers).
- Airflow Management – Cooling demand is affected greatly by the path, temperature, and amount of cooling air delivered to the IT equipment and the separation of hot air removed from it. Best practices include eliminating mixing and recirculation of exhaust, maximizing return-air temperatures by supplying optimally conditioned air directly to the loads, and heat-recovery.
- Air Systems - The air handler fan is typically the second largest energy user in the cooling system. Optimizing designs for use in computing environments, as opposed to using inefficient computer room air conditioners, can save energy.
- Chilled Water Plant Optimization - High-efficiency chillers and variable-speed drives can garner large savings. Optimized sizing and layout can reduce initial costs dramatically.
- IT Equipment Selection - The computers themselves can be made to be more efficient. Optimized power conversions,, reduced power modes, and multi-core processors (see below)–has recently become available, thereby reducing the need for mechanical infrastructure.
- Electrical Infrastructure - Backup power facilities themselves can use a large amount of power, even in standby mode. Careful design, selection of efficient UPS systems, and on-site self-generation can reduce the usage. Shifting to a high voltage AC or DC power distribution system can produce significant savings. (although is less applicable in supercomputers).
- On-Site Power Generation – Savings can be obtained by avoiding the transmission losses and use of the waste heat in absorption or adsorption chillers.

These developments represent a new era in the energy-efficiency movement in HPC. According to a recent report to Congress, compared to the current trends (which include some efficiency gains) energy use from conventional data centers nationally could be reduced by about 60% between 2007 and 2011, valued at over $5 billion and 47 million metric tons of CO2 each year.[7] These savings come without compromising product or

[5] http://www.economist.com/displayStory.cfm?story_id=10436089&fsrc=RSS

[6] Uptime Institute. 2008. "Data Center Capacity and Energy Efficiency Survey."

[7] US EPA. 2007. Report to Congress on Server and Data Center Energy Efficiency: Public Law 109-431. Washington, DC: U.S. Environmental Protection Agency, ENERGY STAR Program. August 2.

data center performance. While supercomputers were not explicitly included in that study, it provides an indication of the opportunities.

EMERGING TECHNOLOGIES & ARCHITECTURES

Future supercomputer facility designs will better integrate the processes of cooling the facility and the IT equipment within, achieving new levels of energy efficiency and first-cost savings. Strategies will include eliminating redundant fans and providing liquid cooling to the chip. For example, Pacific Northwest National Laboratory has identified the potential to save $15 million in capital improvements via improved cooling efficiencies in its own HPC facilities. With one emerging technology—spray cooling—they estimate being able to increase the number of servers by 50% (from 2000 to 3000 per 5000 square feet of floor area), thus deferring otherwise essential new construction.[8]

Even more fundamental innovations will involve rethinking the computing process itself.

Microprocessors designs have recently moved away from exponential scaling of clock frequency toward chip multiprocessors (CMPs) in order to better manage trade-offs among performance, energy efficiency, and reliability as described in the report entitled "The Landscape of Parallel Computing Architecture: A View from Berkeley." [9] Parallelism using arrays of simpler/less powerful processor cores is a more energy efficient way to achieve performance than the traditional approach of using small numbers of complex processors. Industry has already moved forward in this direction with the advent of mainstream products using multicore technology.

A more aggressive approach to mitigating the growing crisis of power consumption in future generations of processing elements is to leverage the enormous resources of low-power embedded processor technology from the consumer electronics industry. The embedded market relies on architectural customization to meet the demanding cost and power efficiency requirements of battery-powered mobile devices such as MP3 players, cell phones, and PDAs. In order to keep up with the demanding pace for semi-customized designs, leading embedded design houses such as Intel, IBM Microelectronics, Altera, and Tensilica have evolved sophisticated toolsets to accelerate the design process through semi-automated synthesis of custom processor designs. One can leverage the tremendous resources of this technology sector to develop a power-efficient HPC system based on application-driven, semi-custom embedded processors. By exploiting a higher degree of parallelism, and some design principles of the low-power embedded computing industry, systems can be built that consume a tiny fraction of the power of existing computing centers without unduly compromising computing performance.[10] For example, proposed specialized Tensilica architecture would achieve a next-generation climate-modeling tool for a task such as 1.5km-resolution climate modeling requiring 2.5 MW (200 petaflop) for the computing infrastructure at an estimated cost of $75 million, as compared a more conventional x86-based system requiring about 180 MW (5 petaflop) at an estimated cost of $1,800 million.[11]

With an eye towards reducing emissions and costs still further—potentially to 100% carbon neutrality—computing facilities are locating closer to clean (and less expensive) sources of hydropower or looking for other ways to generate or procure clean power.

THE VALUE OF BENCHMARKING ENERGY PERFORMANCE

You can't improve what you don't measure. Energy benchmarking is becoming more common in the computing world, and can serve as a valuable indicator of relative performance, in turn pointing the way towards opportunities for improvement.

The Server Metrics[12] extension of the Energy Star program is crafting a rigorous set of metrics and testing methodologies oriented towards web servers. However, unlike the CPU-centric testing regime specified by these procedures, HPC application performance strongly depends on coordinated performance of many servers and the interconnect that binds them together.

The SPEC High-Performance Group is actively pursuing power efficiency benchmarking standards for HPC systems.[13] An important feature of their effort is the emphasis on defining metrics representing useful work per measured watt, which is a superior alternative to current industry trends of advertising peak flops/watt. If consumers buy based on peak flops/watt, efficiency gains are not guaranteed if the architecture is poorly suited for running scientific applications efficiently.

BRINGING IT ALL TOGETHER

Efficient high-performance computing requires not only reliable and efficient design, but also proper facility construction, commissioning, and operation. After all, these facilities are dynamic environments. Workloads, requirements, equipment, and regulations all change. What was efficient when the data center was first built may be far from optimal later.

Maximizing energy efficiency while maintaining computing performance and reliability requires effective systems integration throughout. That is, organizations must implement design-intent documentation, harmonize energy management with core business decision-making, perform benchmarking, and build in-house expertise through training. The climate and our HPC budgets will be better off when these best practices become business-as-usual.

Evan Mills
Staff Scientist, Lawrence Berkeley National Laboratory
William Tschudi
Project Manager, Lawrence Berkeley National Laboratory
John Shalf
Staff Computer Scientist, National Energy Research Supercomputing Center, Lawrence Berkeley National Laboratory
Horst Simon
Associate Laboratory Director, Computing Sciences, Lawrence Berkeley National Laboratory, and Adjunct Professor, Electrical Engineering and Computer Science (EECS) Department, University of California, Berkeley

[8] Hot Topic – Cooling Supercomputers. Pacific Northwest National Laboratory. http://esdc.pnl.gov

[9] Krste Asanović, Ras Bodik, James Demmel, Tony Keaveny, Kurt Keutzer, John D. Kubiatowicz, Edward A. Lee, Nelson Morgan, George Necula, David A. Patterson, Koushik Sen, John Wawrzynek, David Wessel and Katherine A. Yelick. 2008. "The Parallel Computing Laboratory at U.C. Berkeley: A Research Agenda Based on the Berkeley View", Technical Report UCB/EECS-2008-23, EECS Department, University of California at Berkeley, March.

[10] The Economist. "Cool it!" March 4, 2008.

[11] Wehner et al (op cit).

[12] EPA ServerMetrics Workshop. http://www.energystar.gov/serverconference.

[13] Standard Performance Evaluation Corporation, see http://www.spec.org/hpg/