



## News

[Press Releases](#)  
[Event Announcements](#)  
[CCT Weekly](#)  
[Grants and Funding](#)  
[Student News](#)  
[Archived News](#)

## Power a Major Hurdle on Road to Exascale

*Unless some groundbreaking solutions are forthcoming, exascale computing may remain more fancy than fact*

(Source: [Scientific Computing](#))

When the Defense Advanced Research Projects Agency (DARPA) issued the report "Exascale Computing Study: Technology Challenges in Achieving Exascale Systems" on September 28, 2008, it sent shock waves through the high performance computing (HPC) community. The report flatly stated that current technology trends were "insufficient" to achieve exascale-level systems in the next five to 10 years. The biggest stumbling block? Power.

The document noted that, "The single most difficult and pervasive challenge perceived by the study group dealt with energy, namely finding technologies that allow complete systems to be built that consume low enough total energy per operation so that, when operated at the desired computational rates, exhibit an overall power dissipation (energy per operation times operations per second) that is low enough to satisfy the identified system parameters. This challenge is across the board in terms of energy per computation, energy per data transport, energy per memory access, or energy per secondary storage unit."

That little yellow globe floating serenely at the top of the chart shown in Figure 1 says it all — despite Moore's Law, you can't get there from here. At least not by using the incremental, evolutionary approach favored by most vendors.

### Going for it

That was three years ago. Undeterred, DARPA went ahead and issued a Broad Agency Announcement (BAA) — the agency's equivalent of a request for proposal (RFP) — to develop exascale prototype systems. Contracts were awarded to a number of vendors, laboratories and other institutions including Intel, NVIDIA, Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory, and Sandia National Laboratory.

DARPA is calling for nothing less than revolution in computing. States the BAA, "To meet the relentlessly increasing demands for greater performance and higher energy efficiency, revolutionary new computer systems designs will be essential...UHPC [ubiquitous high performance computing] systems designs that merely pursue evolutionary development will not be considered."

Power is a prime candidate for a major makeover. DARPA wants to develop a single rack capable of delivering in excess of one petaflop on the LINPACK benchmark for a power budget of 57 kilowatts. An unspecified number of these racks should be able to interoperate to address a single application — the building blocks of an exascale system.

Meeting the challenge Bill Dally, Chief Scientist at NVIDIA, is one of the HPC industry luminaries addressing the DARPA challenge. Dally was a member of the 2008 DARPA Exascale Study, so he has no illusions about the enormity of the task.

Dally explains: "The problem is that incremental evolution from today's computer systems won't get to exascale within a reasonable power budget. A conventional HPC system based on CPUs consumes about 5n J/flop — or 5 MW for a one-petaflop system. Scaling this to an exaflops system would take 5 GW. Technology scaling is expected to give us about a factor of four improvement in energy efficiency over this time period, so that will make the evolutionary exascale system take 1.25 GW — still a factor of 25 more than the 50 MW that many regard as a practical upper limit for such a system. Even with GPUs, which at the system level are three to five times more efficient than CPUs, there is still a large power gap to close to realize an exascale system at reasonable power levels."

Dally makes the analogy that improving the energy efficiency of computers by 25x (100x counting the technology scaling) is like improving the mileage of a car from 30 mpg to 750 mpg (3,000 mpg).

"This is a major improvement and not something that is going to come from incremental steps," he says. "This is why we are looking at very different approaches to architecture and programming systems than what is done today. To mitigate the risks associated with some of these technologies, we are focusing our effort where the ratio of gain (in energy efficiency) to risk is highest. We also are pursuing multiple approaches to some problems to give a high probability of coming up with a workable solution. And we are sensitive to the need to provide backward compatibility — even if the underlying hardware looks quite different. Existing customer programs have to run well on this system."

Dally predicts that the exascale system developed in the 2018-2020 time frame will be a heterogeneous computer system with a deep, exposed storage hierarchy. The bulk of the performance will be provided by highly efficient "throughput-optimized" cores — similar to today's GPUs but more efficient and able to handle more general code. A few "latency-optimized" cores — like today's CPUs — will be used to run small portions of the code that are critical-path limited. A deep, software-managed storage hierarchy will be used to supply the majority of data to the bandwidth-hungry execution units from local memories.

The NVIDIA UHPC team is exploring co-designing programming systems and supporting architectures that exploit all available locality in a program to minimize energy due to data movement. This involves a very deep and configurable storage hierarchy and an API that lets the compiler exploit it. They also are developing circuit technologies that minimize the energy needed to move data when it is unavoidable.

"It's going to require a lot of hard work and a number of breakthroughs," Dally concludes. "However, I believe the UHPC goals are achievable. A machine with this level of efficiency doesn't violate any physical laws. We can actually perform the floating point operations with far less power than this. Getting to these power levels is largely a matter of minimizing data movement and eliminating overhead."

### View from LSU

Sandia National Labs, one of the UHPC prime contractors, has asked Louisiana State University's (LSU) Center for Computation & Technology to investigate execution models, runtime system software, memory system architecture and symbolic applications. The LSU research group is lead by Professor Thomas Sterling, an HPC industry leader known for his research on petaflops computing architecture.

Sterling has a somewhat different take on the major obstacles to achieving exascale. He comments, "Many people think power and energy is the number one problem — and it is an enormous problem. But, I consider parallelism to be the biggest hurdle. Even if you have an infinite amount of power, without the requisite parallelism, you will not achieve exascale."

Architecturally, data movement is the primary contributor to power consumption. But by embedding a processor directly adjacent to the memory, Sterling says you can maximize bandwidth and minimize the latency between memory and execution logic. This embedded memory processor (EMP) approach is oriented toward multithreaded HPC.

"Sometimes it's cheaper and more energy efficient to move work to the data, rather than doing a giant gather and moving the data to the work," he adds. "The architecture and the runtime system use parcels or active messages to move an action requirement to another block of remote data and perform the action there. This greatly reduces data movement and saves time and energy as well."

For Sterling, achieving exascale requires a radically new and different software paradigm. Major advances must be made in terms of programming models and tools as well as runtime software — changes that will be just as dramatic as those made by UHPC teams addressing exascale hardware technology.

Will we meet these goals?

Here's what NVIDIA's Bill Dally has to say: "It appears very likely that there will be exascale machines by 2020. The real question is whether they will be developed in the U.S. or whether other countries — those that are investing more heavily in economically critical technologies like computing — will get there first."

**Publish Date:**  
10-14-2011

[Home](#) | [About](#) | [Research](#) | [Programs](#) | [News](#) | [Events](#) | [Resources](#) | [Contact Us](#) | [Log In](#) | [LSU](#) | [Feedback](#) | [Accessibility](#)



Center for Computation & Technology

2003 Digital Media Center • Telephone: +1 225/578-5890 • Fax: +1 225/578-8957

© 2001–2025 Center for Computation & Technology • Official Web Page of Louisiana State University.