



## News

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

## Preoccupied with Exascale

(Source: [HPCwire](#))

Every HPC event you attend this year will almost certainly devote much attention to the drive toward exascale. The HPCC conference in Newport, Rhode Island, this week was no exception. Besides the topic of the "missing middle," which I covered in my [previous post](#), exascale computing was probably the biggest single focus at HPCC this year. That makes sense, given that the supercomputer crowd is always leaning forward, and exascale is obviously the next big milestone.

Or is it? After hearing so much about exascale over the last couple of years, I'm starting to wonder about the rationale of devoting so much effort to what is essentially an arbitrary milestone based on the nomenclature of our decimal numbering system. Why not think about the challenges of 100-petaflop, or even 10-petaflop systems?

For that matter, why not devote more resources to figure out how to make today's single-petaflop and multi-teraflop systems fundamentally better? Currently, there are only a handful of applications that can use a petaflop of computing. And only a small number of sites can even install a petaflop machine, given their cost (100-plus million dollars) and energy expense (several million dollars per year). In 10 years, exaflop machines will be equally rare and underutilized.

Getting applications to use cutting-edge supercomputers to the fullest extent has always been particularly difficult. Our track record of preparing software -- application-level or system-level -- for systems 10 years into the future is rather poor. I'm not sure what more we can expect, though. The hardware characteristics of systems not yet born are, by definition, difficult to anticipate.

To mitigate that problem, the HPC digerati are turning to "co-design" (i.e., developing hardware in conjunction with software) for exascale designs. It sounds like a wonderful idea, but I'd be hard-pressed to think of success stories using this approach. There is a reason hardware comes first: it's the basic foundation upon which the higher abstractions of software are created. To some extent, co-design seems like trying to teach the baby while it's still in the womb.

At HPCC, four of the 18 sessions focused almost exclusively on exascale, and many of the others at least touched on the topic. The one that particularly caught my attention, though, was the UHPC panel that discussed the work under development for DARPA's Ubiquitous High Performance Computing program. The panel had the principals of each of the four UHPC projects (Angstrom, Runnemedede, X-Caliber, and Echelon) talk about their respective approaches and provide an update on their work.

The scope of this article doesn't allow me to elaborate on the specifics of each UHPC effort here (but watch this space for additional coverage in the future). In this context, my main interest is pointing out that UHPC is -- as panel moderator Thomas Sterling pointed out -- not an exascale program, per se. The DARPA RFP that defined this effort focused on "extreme computing" and developing power-efficient hardware, software stacks, operating systems, and programming environments that can scale down as well as up.

One of the goals of UHPC is to produce an architecture that delivers one petaflop in cabinet, with a max power draw of 57 KW. It is these cabinet-sized system that are likely to be widespread in the US DoD (and elsewhere) by the end of the decade. By contrast, exascale systems will be rare and initially serve as special-purpose machines, much as the petascale systems of today are.

Building better software and hardware for today's level of supercomputing is a laudable goal. There's plenty of backfilling to do in this regard, and that's why I think the journey to exascale will be more important than its destination.

That's the good news. The bad news is that there is concern that UHPC funding may be axed. At HPCC, rumors were floating about that money to support this effort will not be forthcoming. This was brought up at the panel session, and although all the participants seemed aware of the funding uncertainties, no one knew how this might play out.

In fact, the US government's budgetary angst was a topic of discussion throughout the HPCC conference, and there was plenty of pessimism to go around. The general consensus was that given the political climate, government-funded HPC might be on the brink of its own recession. The InterSect360 forecast delivered at the conference predicted government HPC would grow modestly this year, but that forecast could turn south quickly if federal and local budgets start slicing off science and technology programs.

Today's political climate will be especially problematic for exascale work. The community has never done a great job of explaining the societal payback for high performance computing that would generate urgency for those in the government. It's difficult enough to distill the value of HPC into sound bites, but because exascale HPC is especially hard to explain to non-science types, that work will be particularly hard to sell. The multi-year time frame for exascale is another big disadvantage, given the rather short-term outlook of most politicians in this time of tight money.

With that in mind, HPC may indeed be entering a period of limited public support. If so, the community may have to refocus its priorities, as unpalatable as that seems. Exascale will inevitably happen. Moore's Law, heterogeneous architectures and optical interconnects will see to that. But we may end up drifting into exascale rather than driving it.

**Publish Date:**  
04-01-2011

