

Log in | Sign up



Data Centre Cloud Software Networks Security Policy Business Jobs Hardware Science Bootl
Servers HPC Storage Data Networking Virtualisation Cloud Infrastructure BOFH

DATA CENTRE > HPC

Supercomputer efficiency jumps, but nowhere near exascale needs

Ceepie-geepies blow past BlueGene/Q in Green500 rankings

By Timothy Prickett Morgan, 28th June 2013

4

[The move to Exchange 2013](#)

RELATED STORIES

ISC 2013 HPC geeks ponder 100 petaflopers and quantum supercomputers

ISC 2013 What should replace Linpack for ranking supercomputers?

ISC 2013 Top500: Supercomputing sea change incomplete, unpredictable

It is not precisely the kind of leap that the supercomputer industry needs to reach exascale performance by the end of the decade, but more powerful GPU and x86 coprocessors are enabling more energy-efficient machines, at least according to the latest Green500 rankings.

The Green500 list comes out two or three times a year, usually in the wake of the Top500 supercomputer performance rankings. The latest Top500 list was just [announced](#) at the International Super Computing conference in Germany last week, and ranks supercomputers of all types based on their sustained performance on the Linpack Fortran matrix math benchmark.

The [Green500 list](#), created by Wu-chun Feng and Kirk Cameron of Virginia Tech, starts with the Top500 data, but adds other submissions from the world's HPC labs, and ranks them all on how little juice they can sip while flopping around running Linpack.

IBM's Power-based, all-CPU BlueGene/Q massively parallel supercomputer had been at the top of the Green500 charts for years, but various kinds of hybrid CPU-GPU ar now Xeon-Xeon Phi machines have been breaking into the top of the list with slightl

The Register uses cookies. Some may have been set already. [Read about managing our](#)

cookies

Please click the button to accept our cookies. If you continue to use the site, we'll assume you're happy to accept the cookies anyway.

I'm fine with this [X]

win the Top500 crown

Italian 'Eurora' supercomputer pushes the green envelope



Mi piace  5

Tweet  22

Cray/Appro – but not by nearly as much of an edge.



The interesting thing about the most energy-efficient super this time around, the "Eurora" machine at Cineca in Italy, is not just that it pairs an Intel Xeon CPU with an Nvidia Tesla K20 GPU, but that it uses a special SKU of the Tesla K20X GPU coprocessor from Nvidia and has a single workstation-class, eight-core Xeon E5-2687W processor to drive two K20X GPUs.

The E5-2687W is a bit of a beast, with a [150 watt thermal design point](#), but at 3.1GHz it has plenty of pep. Being designed for a two-socket machine, it has more I/O and memory scalability than the current "Ivy Bridge" and "Haswell" single-socket Xeon E3-1200 chips. What's more, if you want more pep and less greenery, you can fire up the second E5-2687W in the two-socket system.

The pairing of a single fast and hot CPU and two fast and hot discrete GPU coprocessors, ironically enough, lets Eurora deliver 3,209 megaflops per watt of performance. This machine, which fits in a single rack, has only 110 teraflops of sustained performance. This is a perfectly respectable midrange HPC system, and Eurotech would surely be pleased to build a 3.1 petaflops box that fits in nine racks if you wanted to pay for it.

A machine similar to Eurora is installed at Selex ES in Switzerland, based on the same Aurora Tigon servers from Eurotech. It uses the same Xeon workstation processors and the special Tesla K20X GPU coprocessors, delivering 3,180 megaflops per watt running Linpack.

The Aurora Tigon servers (part tiger, part lion and hence a hybrid) have homegrown blade servers with room for two Xeon processors and two embedded K20X GPUs, each with metal plates on them instead of heat sinks so water blocks can be mounted directly on the processing elements to remove their heat. If you can remove the heat efficiently, you can crank the parts faster and get more floating point math done. During the Linpack runs for the Green500 list, only one of the Xeon processors in these machines was activated.

Number three on the Green500 list is the November 2012 champ, the "Beacon" hybrid machine installed at the University of Tennessee. This is the Cray/Appro box comprised of regular Xeon-E5 processors and Xeon Phi 5110P coprocessors, and

The Register uses cookies. Some may have been set already. [Read about managing our cookies.](#)

Please click the button to accept our cookies. If you continue to use the site, we'll assume you're happy to accept the cookies anyway.

I'm fine with this [X]

with AMD FirePro S10000 discrete graphics cards to yield 2,351 megaflops per watt

That is just a smidgen ahead of six BlueGene/Q machines, which are rated at 2,299 megaflops per watt in their various HPC labs in very small configurations. Larger BlueGene/Q machines, such as the "Vulcan" and "Sequoia" machines at the US Department of Energy's Lawrence Livermore National Laboratory, come in at a slight smaller 2,177 megaflops per watt.

The Top500 winner, the Tianhe-2 supercomputer built by the Chinese government for aerospace and physics research, delivered 33.86 petaflops of sustained performance using a mix of Xeon and Xeon Phi computing elements, but it ranked down at number 31 on the Green500 list with 1,902 megaflops per watt.

The energy efficiency of the machines in the Green500 list drops off pretty fast, and the time you are down into the 200s on the list, you are more than an order of magnitude less power efficient than the machines at the top of the list. At the bottom of the list, you are in the range of 40 to 50 megaflops per watt, an embarrassing number that is the result of large machines in the 1 to 2 megawatt range, based on Xeon 5600 or Opteron 6100 processors and slower InfiniBand or Ethernet interconnects.

Companies build supercomputers to last five years or so, and while the energy efficiency they initially get is perfectly reasonable on all-CPU machines, it is clear that where the power envelope is an issue, companies are going to have to use some kind of accelerator and rework their code. And that may not get us to exascale in a 20 to 40 megawatt power budget by 2020.

But it is important to reward progress, and some progress has been made.

If you wanted to build a machine using the CPU and x86 coprocessor technology in the Beacon machine that was at the top of the last November's list, you would need 408 megawatts. But if you scaled up the Eurora machine at Cineca to 1 exaflops, you would need 312 megawatts of juice. Sure, that is a 24 per cent drop, but if you keep this same pace at an annualized rate between now and 2020, you will still need a 42 megawatt nuke plant to power an exascale machine.

The energy efficiency has to go up a lot faster. That probably means integrating interconnects with CPUs and GPUs, welding main memory to the chips, and using very clever optical networking. It's not clear to anyone how we get there – in fact, and more importantly, it's not clear that real-world software would be able to scale across such monstrosities.

But it is still fun to noodle it, ain't it? ®

[The move to Exchange 2013](#)

The Register uses cookies. Some may have been set already. [Read about managing our cookies.](#)

Please click the button to accept our cookies. If you continue to use the site, we'll assume you're happy to accept the cookies anyway.

I'm fine with this [X]