Toward Exascale computing

HPC ADVISORY COUNCIL
NETWORK OF EXPERTISE

After we broke the Petaflop performance barrier (the top 3 systems on the TOP500 supercomputers list demonstrated higher than 1 Petaflop sustained performance), the HPC community is exploring development efforts for breaking the Exaflop barrier. Many organizations and industry collaboration have been established around the world, and multiple media outlets were created to cover the development and collaboration progress.

When and where will the first Exaflop system be built? If we review the performance trends according to the TOP500 supercomputers list performance development graph (Figure 1), we can estimate that the first Exaflop system will be built between 2018 and 2020. According to recent Scientific Discovery through Advanced Computing Conference (SciDAC, July 2010) presentations, the US goal is to have such a system built in the 2018 time frame.

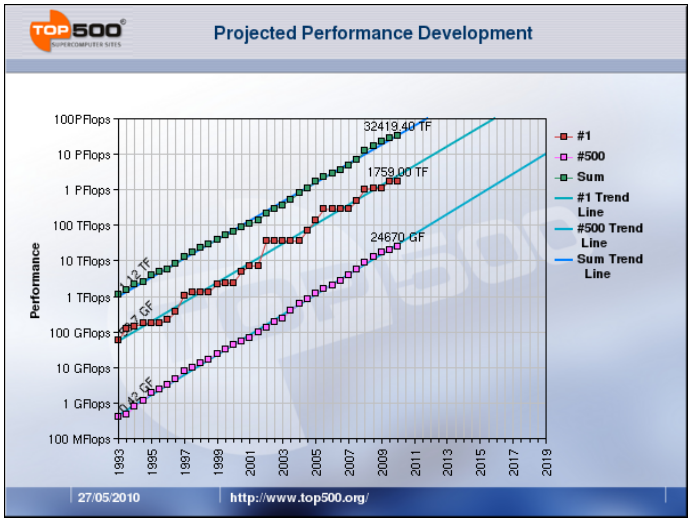

Figure 1: Top500 projected performance development

Estimating where the first Exaflop system will be built is a tougher question. Figure 2 reviews the performance development of the highest ranked systems on the TOP500 from the US, Europe and China, including percentage lines to indicate the performance difference trend between the US's highest ranked system compared to the highest ranked systems in Europe and China.

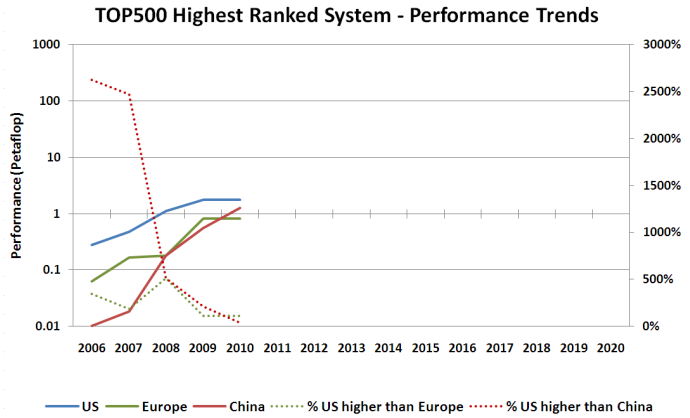

Figure 2 – TOP500 highest ranked system performance trends – US, Europe and China

HPC development in China has dramatically accelerated in the recent years. In 2006, the performance leadership of the US's highest performance system was more than 2500% higher than the best system in China, in the first half of 2010, the gap decreased to less than 40%. Market expectations show that China might host the fastest supercomputer in the world by the end of 2010. While we expect to see 20 Petaflop systems built in the US in 2011-2012 time frame, if this performance trend continues, China might also be hosting the first Exaflop system in the world. Along with China and the US, other countries have a decent chance to host the first Exaflop system – Germany, UK, Japan, Russia, and maybe even India. While we are seeing many cross nation collaboration (Europe, US, etc) we shouldn't eliminate the option of some nations building their own hardware components for constructing the world's fastest supercomputers.

The road to Exascale computing requires collaboration between hardware and software vendors, as well as between vendors and users. The HPC Advisory Council's HPC|Scale group is one example of vendor and user collaboration. Chaired by Richard Graham from Oak Ridge National Laboratory and Gilad Shainer, the HPC Advisory Council's Chairman, the HPC|Scale group includes multiple participators from a large variety of user organizations and vendors. The group's mission is to explore scalability issues, potential optimizations and offloading technologies that will be required in Exascale systems, to minimize jitter and create very effective synchronization capabilities (where these are needed),

and look at overlapping computation with communication to hide the impact caused by the collective nature of the collective operations for example, and other related topics."

What are the key items that we should be looking at? Of course we need to be able to pack as much floating point elements into the smallest space as possible (dense computing, Earth is limited in space…) and we definitely need to decrease power consumption (power efficiency, our energy resources are limited…). Application development is ongoing and we continue to witness more and more applications that can scale fully on the world's largest supercomputers today. The communication elements will become more and more critical in application efficiencies, and offloading networks will be mandatory (can't afford wasting floating point cores on non-compute activities!). The systems will need to be able to dynamically provision among the multiple users and workloads.

Moreover, Exascale computing will produce an exaflood of data that must be stored and shared. We have covered that topic in the From Computational Science to Science Discovery: The Next Computing Landscape paper published on the HPC|Cloud subgroup page.

So what is our forecast? Well… 1-2 Petaflop systems exist today, 20 Petaflop in 2012, 200-500 Petaflop in 2015 and Exaflop in 2018.

Contributing authors:

Gilad Shainer, Brian Sparks, Richard Graham

HPC ADVISORY COUNCIL
NETWORK OF EXPERTISE

350 Oakmead Pkwy, Sunnyvale, CA 94085
Tel: 408-970-3400 • Fax: 408-970-3403
www.hpcadvisorycouncil.com