Testimony before the U.S. - China Economic and Security Review Commission

Date of the hearing:

Thursday, May 10, 2012

Name of the panelist:

Dr. Earl Joseph II

Panelist's title and organization:

HPC Program Vice President, IDC

Title of Hearing:

Hearing on "Assessing China's Efforts to Become an Innovation Society—A Progress Report"

Panel IV: Information and communications technology innovation

"Supercomputers, Technical Servers And Innovation"

Outline:

1. Who Is IDC
2. My Background
3. The Technical Computing Group at IDC
4. Worldwide HPC And Supercomputer Market Trends
5. China Information -- And The US Compared To China Supercomputer Trends
6. Recommendations For Our Country
7. Conclusions And Overall Assessment
8. Charts And Tables

1. Who is IDC

IDC is a market research company that tracks IT markets. We have been tracking IT markets for over 45 years, with offices in over 50 countries. Our research approach is based on strong quantitative numbers tracking and surveys -- details can be found at: www.idc.com

IDC company highlights include:

- IDC is an independent global market intelligence, events, and advisory firm for information technology, telecommunications, and consumer technology markets (ICT).
- More than 1,000 IDC analysts, including in-house statisticians and economists, provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries.
- IDC has been delivering IT intelligence, industry analysis, market data, and strategic guidance since its 1964 founding by Patrick McGovern.
- Our multilingual, multicultural workforce surveys over 360,000 technology users and decision makers annually, delivering unrivaled coverage.
- IDC is a subsidiary of IDG, the world's leading technology media, events and research company.

2. My Background

I have previously worked for a number of computer vendors including Cray, SGI, Concurrent real-time systems, and Sperry-Univac. I have been with IDC for over 13 years, focused on technical computing including supercomputers, buyers, users and vendors. My resume and background information can be found at: www.earljoseph.info

My Ph.D. was in the management of high technology companies with related focuses in quantitative analysis, math, and engineering.

3. The Technical Computing Group at IDC

Some definitions:

- HPC = High Performance Computing = All Types of Technical servers.
- Technical Servers = computer servers sold for use by scientists, engineers, analysts and others for highly computational or data intensive computing.
- Supercomputers = technical servers that cost more than $500,000.

Within the technical computing space, we:

- Attempt to track all HPC servers sold each quarter (around the world), from all vendors. Our data structures go back over 25 years, and we maintain 5 year forecasts in many areas and for most sub-segments.
- Conduct 4 HPC User Forum meetings each year (see: www.hpcuserforum.com)
- Publish over 35 research reports each year.
- Visit all major supercomputer sites & write reports about what they are doing and what their plans are for the future. E.G. We recently visited the large sites across China and are currently writing each one as a case study.
- Assist in collaborations between buyers/users and vendors.
- Assist governments in HPC plans, strategies and directions.
- Assist buyers/users in planning and procurements.
- Conduct an HPC innovation award program every six months.
- Conduct special research studies, e.g.:
    - USG special studies
    - EU 2020 HPC leadership plan
    - EU parallel software scaling plan
    - DOE skills, talent and hiring plan
    - Automotive benchmark analysis
    - Oil/gas HPC usage comparisons
    - NCSA/NSF science underlying applications study
    - Korean HPC strategy implementation plan
    - Council on competitiveness innovation studies
    - Power and cooling studies
    - And many private studies

4. Worldwide HPC And Supercomputer Market Trends

The worldwide technical server market recovered from the 2008 recession very quickly with an impressive 10% growth in 2010 followed by 8.4% growth overall in 2011. IDC is projecting approximately 7.3% yearly growth overall for the next five years.

HPC major trends include:

- All tech severs are around $10 billion a year in sales today (as shown in the first table below), containing 3.3 million processer packages in 112 thousand systems.
- The boarder HPC market, including storage, software and services is around $20 billion today.
- Supercomputer servers represent around $4.4 billion a year today, in 2.8 thousand systems sold each year.
- The major industry/application uses are shown in the second table below, with government labs, universities, bio-life sciences, defense and manufacturing being larger sub-segments.

The supercomputer portion of the HPC has seen the largest growth due to heavy investments by major nations around the world. Supercomputers grew 25% in 2011, with a number of systems costing over $100 million each. The largest supercomputer in the world today, at RIKEN in Japan cost over $500 million.

Some of the reasons for the strong growth include:

- They are used for increasing national innovation, scientific capability, economic growth & job creation.
- It has become a competitive weapon
    - For companies, universities and governments
    - Global competitiveness is driving R&D and better product designs
- Governments view HPC leadership as critical
    - It use to be 1 large supercomputer – now its multiple ones
- There are very critical HPC issues that need to be solved
    - Global warming, alternative energy, safe NE, financial disaster modeling, healthcare, homeland security, …
    - And 3D movies and large scale games are fun
- At the same time, “live” science and “live” engineering costs have escalated
    - And time-to-solution is months and sometimes years faster with simulations

Some interesting quotes and actions from world leaders:

- In 2009, Russian President Dmitry Medvedev warned that without more investment in supercomputer technology: “Russian products will not be competitive or of interest to potential buyers.”
- In June 2010, Rep. Chung Doo-un of South Korea’s Grand National Party: “If Korea is to survive in this increasingly competitive world, it must not neglect nurturing the supercomputer industry, which has emerged as a new growth driver in advanced countries.”

    - The Korean National Assembly then called for the creation of a national five-year plan for advancing HPC
- In February 2012, the European Commission announced that it has adopted a plan to double spending on HPC to €1.2 billion, with much of that money aimed at the installation of additional petascale supercomputer systems

In summary, there is a worldwide battle for supercomputing leadership taking place around the globe -- it's a tool for economic growth and innovation -- large funding is all that is required these days (well mostly funding). There are some major limitations in the ability to effectively use supercomputers today including: Software lags hardware and the gap has been growing at an alarming rate; and the lack of enough talented people to make full use of the supercomputers.

5. China Information -- And The US Compared To China Supercomputer Trends
China has been increasing their investments and installations of all types of technical servers as show in the charts below. Since 2002 they have increased their yearly HPC purchases by an impressive six-fold and their supercomputer purchases by close to 7 times.

The charts below compare the US to China purchases and clearly shows that the US has a sizable lead in total installations, but China's rate of growth is dramatically higher. The US purchases almost five times the amount of supercomputers that China does each year, but is only growing by around 4% a year, while China is growing supercomputer purchases by 22% a year (since 2002). The rate of increase by China could lead to China becoming a major portion of the world market, resulting in vendors customizing future systems to better match requirements from China buyers.

China today spends about $600 million on technical servers and around $375 million on supercomputer systems. In addition, China is building new institutions and new centers to make better use of these supercomputer assets. They plan to have 17 petascale supercomputer centers in the near future. China already has passed up Germany in the number of Top500 supercomputers and will soon pass up all of Europe combined on the Top500 list (see www.top500.org). Counting total supercomputer purchases, China likely exceed Japan and Germany in the next year or so to become the second largest purchasing country in the world (next to the US).

China's use of supercomputers is very different than in most other countries, in that China views supercomputers as a toll and not as a costly asset. So while other centers work very hard to keep the computers fully utilized, China tends to focus on using the computers to support researchers and train new users. By having this extra "head room" it makes it easier for a researcher with a new idea to get substantial time on the computer. It is a more strategic vision than used by other countries.

China desires to obtain leadership in most or perhaps in all areas. China recognizes the importance of HPC for leadership in many areas and is heavily funding HPC to gain a strong standing in science, innovation and for economic growth.

China understands that "To out-compute is to out-compete" (a quote from the US Council on Competiveness). They still have a lot to learn and a lot to learn and need to build out in their HPC ecosystem -- but they recognize this and are addressing it.

China also wants to be known as an innovative leader and not just a copier (or not as just a low cost, cheap supplier). Today, China has the #2 largest supercomputer, it was #1 last year, but Japan built an even large one. China will likely compete strongly to displace Japan as the leader -- and may increase investments even more than currently planned.

China uses their supercomputers on a very broad set of applications as shown bellow, including oil exploration, bio & life sciences, manufacturing, climate, weather, teaching and animation.

Chinese technologies:

- Chinese CPU efforts -- they are developing 2-3 main home grown processors plus a few others. These are based on existing processor technologies like MIPS, ALPHA, etc.
    - From our surveys of buyers in China, most view the Intel x86 as the processor of choice at this time. There are many rumors about China potentially developing a real strong workable processor of its own, but this isn't my area of expertise.
    - This work is innovative in many ways, but evolutionary in others areas in that each processor technology is based on an existing ISA. They have been developing some of these for many years and for over 6 generations, e.g. the Godson/Loongson, MIPS-based processor that Sugon/Dawning is using in a number of products (see http://en.wikipedia.org/wiki/Loongson).
    - They have also developed innovative ways to use these processors with large systems like the TH-1A that uses SPARC-based processors for system control.
    - These processors are seeing success in a number of supporting application areas, e.g. in modems and internet supporting systems.
    - To develop a competitive processor requires both major strength in R&D as well as major investments and skills in building leading edge fabs.
- China has also developed its own very high speed custom interconnect that is used on the TH-1A supercomputer. But it is costly and not broadly used.
    - This is very innovative when combined with the other innovations within the TH-1A design. Only 3 countries have built a large custom peta-scale class supercomputer (US, Japan and China).
    - The interconnect design appears to be similar to other advanced custom interconnects, but its implementation is unique. It is designed to support a very large supercomputer, with higher costs than can be supported in mainstream computers. This is similar to interconnects created and used by US companies like Cray and IBM.
    - It was developed by NUDT (the National University of Defense Technology http://english.nudt.edu.cn/) and was custom designed for the TH-1A

computer. NUDT has built a number of special and innovative supercomputers for many decades in China.

- China also has created an interesting graduate class exercise in developing programs for breaking/checking passwords. The last slide below shows the result of this classroom exercise -- the student's winning program can check 50,000 passwords a second on each node of the computer, 250,000 per second using the GPUs on a node and when run on the whole TH-1A computer 1.8 billion passwords can be checked per second.
    - China recognizes that large scale supercomputers are required for the design, development and pre-testing of advance aviation, ships and military equipment. They proudly display examples of applying supercomputers for high speed fighters, aircraft carriers, satellites, etc.

In summary, China has many weaknesses in HPC -- but they recognize them and are addressing them -- it will take years, but they will likely do it faster than any country has done before.

- China has put into place many undergraduate and graduate programs to teach the use of HPC and how to apply HPC alongside with basic scientific disciplines. China also has a program to bring back to China top researchers and scientists to help train and develop China's eco-system (the goal is to bring back 100 a year).
- China is placing many of their HPC systems into a "Cloud" for smaller organizations and companies to gain access and to learn how to apply HPC in their company.

6. Recommendations For Our Country

Leadership requires major actions, investments and continued actions & innovation over time.

- Smaller 10% - 25% increases do very little.
- Spreading increases across broad groups of existing organizations also do very little to help with leadership. There needs to be a smaller number of targeted areas for leadership.

Today the US is the clear leader in supercomputing, but the level of our leadership has declined and it will require strong investments and leadership to MAINTAIN our current position:

- An increase of $2 billion a year in purchases of supercomputers would be a good start, plus around $2 billion for the infrastructure required to go with these new systems. Then growing it by 10% a year (note -- this must be new dollars, not just shifting of existing funds).
- An increase on the order of $1 billion a year in HPC R&D funded by the USG for developing better systems and building blocks and software.
    - To work well, at least 50% should go to software that will help application jobs performed better, which in most cases means making able to use more processors or scale better.
    - The software R&D must be separate from the hardware R&D, otherwise the hardware will soak up all the funds (as has been the case in the US for decades).
- A good example would be for the US to take the DOE exascale proposal (whatever it finally turns out to be) and funding it at 2x to 3x the plan -- and then asking DOE how

much more they can accomplish and how much faster they can do it with 2x to 3x the funding.
    - Again this needs to be new dedicated funds, not just reshuffling of existing budgets.
- We also need to address the skill sets, training and lack of scientists, engineers and technical experts in our country. We need multiple "Apollo like" missions that are exciting enough to get students interested AND to get experts from around the world to move to the US. We can never graduate enough within the US, so we must also get technical people from around the world to relocate to the US.
    - For example, how many graduating physicists and astrophysicists would move to CERN if given the chance?

7. Conclusions and Overall Assessment

The bottom line is that China has gone from a non-entity in HPC and supercomputers to having one of the largest in the world in only five years. And they are on a trajectory that is dramatically stronger than the US:

- Five years ago they only had a few supercomputers that would make it onto the list of top 500 computers in the world, today they have more than Germany or France, and very soon they will have more than ALL of Europe combined.
- Their largest supercomputer is larger than any known supercomputer in the US (this excludes classified computers).
- They are investing in 17 different petascale centers (petascale = very large supercomputers) around China, many of which are new and are being built from the ground up.

While it is true that they haven't yet fully developed a broader ecosystem to take full advantage of these assets, they are filling out their ecosystem faster than any country has ever done in history. They fully recognize these short comings and are investing heavily in them.

In addition, their rate of graduating scientists, engineers and analysts that can apply supercomputers and HPC to help advance their national science standing and improve their economic competitiveness greatly outstrips what the US or Europe is able to accomplish. While we don't have exact figures on these numbers it's not in the 10 times stronger range -- it is much higher.

In summary, the Chinese government, at the national and local levels understands the value of HPC and supercomputers in increasing innovation, advancing science and growth their economy. They are investing heavily in both very large systems and in growing the eco-system required to take advantage of these assets.

8. Charts And Tables:

An Overview of The Worldwide Technical Server And The Supercomputer Market

Revenues	2010	2011
Supercomputers	3,476	4,361
Divisional	1,269	1,246
Departmental	3,279	3,481
Workgroup	1,475	1,213
Total	9,498	10,300
Systems Sold	**2010**	**2011**
Supercomputers	2,560	2,811
Divisional	3,914	3,749
Departmental	19,868	21,199
Workgroup	93,502	84,832
Total	119,844	112,591

WW HPC Systems Revenue by Applications ($M)	
	2011
Bio-Sciences	$1,252
CAE/Manufacturing	$1,095
Chemical Eng	$193
DCC & Distribution	$569
Economics/ Financial	$279
EDA / IT / ISV	$663
Geosciences	$654
Mechanical Design	$63
Defense	$1,005
Government Lab	$2,078
University/ Academic	$1,901
Weather	$454
Other	$95
Total Revenue	**$10,300**
Source IDC, April, 2012	

How the US Compares to China in Technical Servers (Of All Sizes) and Supercomputers

USA -- All HPC Servers

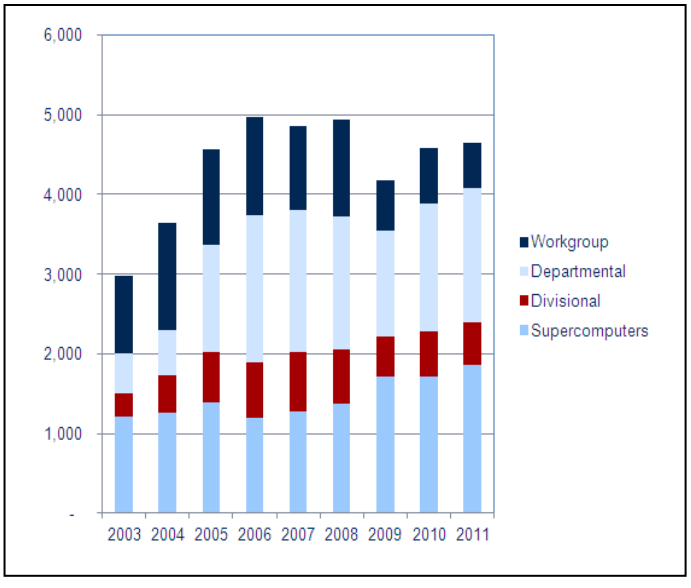

China -- All HPC Servers

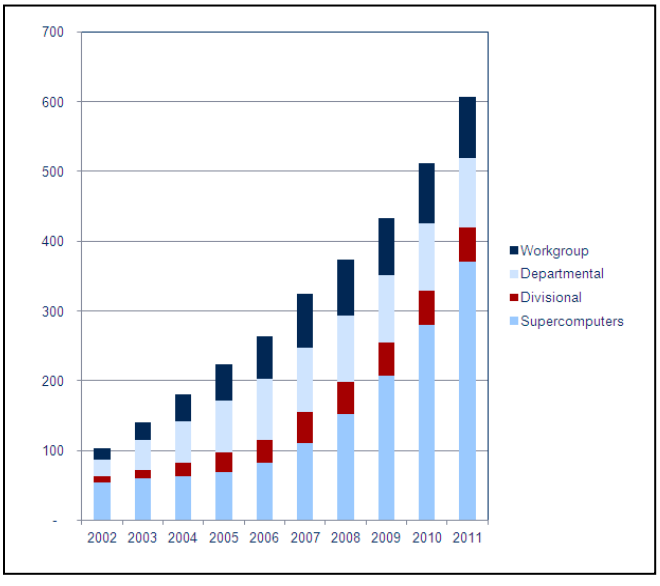

USA -- Only Supercomputers

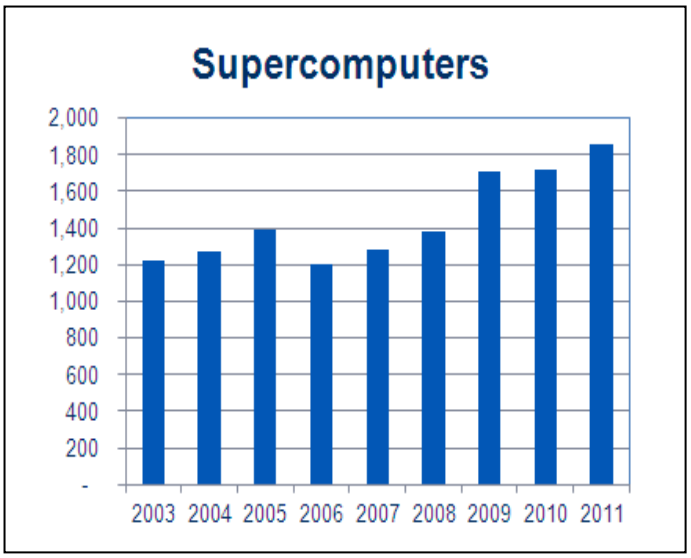

China -- Only Supercomputers

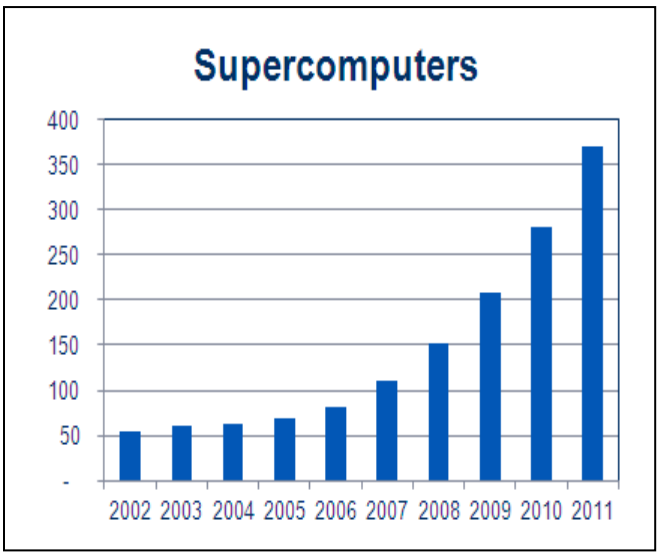

Examples Of How Quickly China Is Increasing Its Budgets And Supercomputer Power

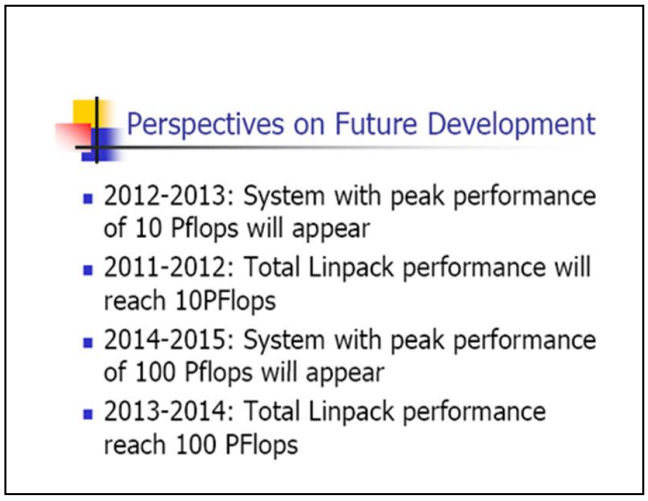

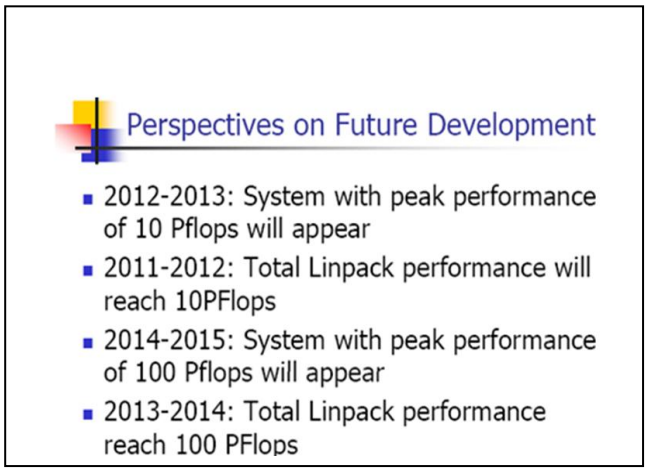

Examples Of How China Is Using Its Supercomputer Power