Global Trends and Competition

Mark Asch - mark.asch@u-picardie.fr

EXDCI Final Meeting 2017

Barcelona, September 7th 2017

September 7, 2017

The EXDCI project has received funding from the European Unions Horizon 2020 research and innovation programmed under the grant agreement No 671558.

Outline

- Current state of affairs and international trends
  - China
  - USA
  - Japan
  - EU
- Implications for EU and EXDCI(2)
- The way forward…

Recent International Reports

- HYPERION report
- IDC report
- HIPEAC report
- BDEC report

HYPERION REPORT

International State of Affairs I: investment

Exascale Race/Technologies

IDC-Projected Exascale Investment Levels (In Addition to System Purchases)

U.S.

- $1 to $2 billion a year in R&D (including NRE)
- Investments by both governments & vendors
- Plans are to purchase multiple exascale systems

EU

- About 5 billion euros in total
- Investments in multiple exascale and pre-exascale systems
- Investments mostly by country governments with a little from the EU

China

- Over $1billion a year in R&D
- Investments by both governments & vendors
- Plans are to purchases multiple exascale systems each year
- Already investing in 3 pre-exascale systems by 2017/18

Japan

- Planned investment of just over $1billion* (over 5 years) for both the R&D and purchase of 1 exascale system
- To be followed by a number of smaller systems ~$100M to $150M each
- Creating a new processor and a new software environment

** Note that this includes both the system and R&D*

International State of Affairs II: exascale delivery

Exascale Race/Technologies

IDC-Projected Exascale Dates and Suppliers

U.S.

- Sustained ES: 2023
- Peak ES: 2021
- Vendors: U.S.
- Processors: U.S.
- Initiatives: NSCI/ECP
- Cost: $300-500M per system, plus heavy R&D investments

EU

- Sustained ES: 2023-24
- Peak ES: 2021
- Vendors: U.S., Europe
- Processors: U.S., ARM
- Initiatives: PRACE, ETP4HPC
- Cost: $300-$350 per system, plus heavy R&D investments

China

- Sustained ES: 2023
- Peak ES: ~~2020~~ 2019…
- Vendors: Chinese
- Processors: Chinese (plus ~~U.S.~~?)
- 13th 5-Year Plan
- Cost: $350-500M per system, plus heavy R&D

Japan

- Sustained ES: 2023-24
- Peak ES: Not planned
- Vendors: Japanese
- Processors: Japanese
- Cost: $600-850M, this includes both 1 system and the R&D costs…will also do many smaller size systems

International State of Affairs III: feasability?

We Can Build an Exascale System Today?

Connect together 10 Sunway TaihuLight systems

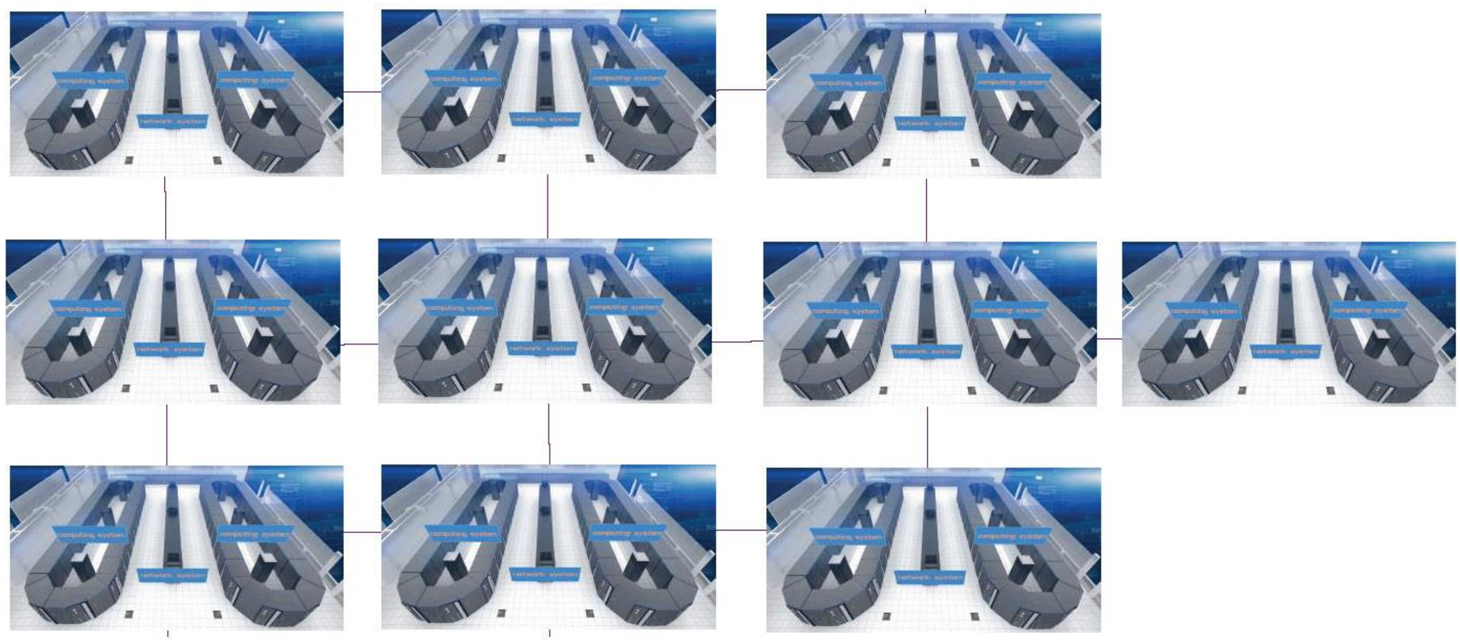

Require 150 MW of power, programming for 100 M threads, and $2.7B price tag

22

Courtesy: J. Dongarra

IDC REPORT

IDC report: RIKEN, December 2016*

This study was carried out for RIKEN by

IDC
Analyze the Future

Special Study

Analysis of the Characteristics and Development Trends of the Next-Generation of Supercomputers in Foreign Countries

Earl C. Joseph, Ph.D.
Steve Conway

Robert Sorensen
Kevin Monroe

- Emerging sectors identified:
  - big data applications,
  - data intensive IoT applications,
  - deep learning and cognitive computing paradigms.

*Released June 2017…

IDC report: conclusions for convergence

- Modest attention being paid to non-traditional HPC software that IDC analysts expect will become increasingly important in the next few years,
  - big data infrastructures built around the Hadoop/Apache Spark (or other alternative) ecosystem and
  - for virtualization schemes such as Docker
- IDC believes that countries that fail to fund development of these future leadership-class supercomputers run a high risk of falling behind other highly developed countries in scientific innovation, with later harmful consequences for their national economies

IDC report: conclusions

- Europe has strong software programs and a few hardware efforts, plus EU funding and support appears to be growing, …, but they have to deal with 28 different countries, and a weak investment community.
- The U.S. has multiple programs, strong funding and many HPC vendors.
- China has had major funding support, has installed many very large systems, and is developing its own core technologies, but has a smaller user base, many different custom systems and currently is experiencing low utilization of its largest computers

HIPEAC REPORT

HIPEAC Vision 2017: IoT and the "edge"

- Energy: to stem the flood of data from the IoT, we must employ intelligent local data processing on remote devices that use minimal energy.
- For compute-intensive tasks, we will still be using the cloud, and that means that connectivity is crucial, yet local processing is also becoming increasingly important so system architects look for accelerators diverting in many cases from traditional Von Neumann architecture.

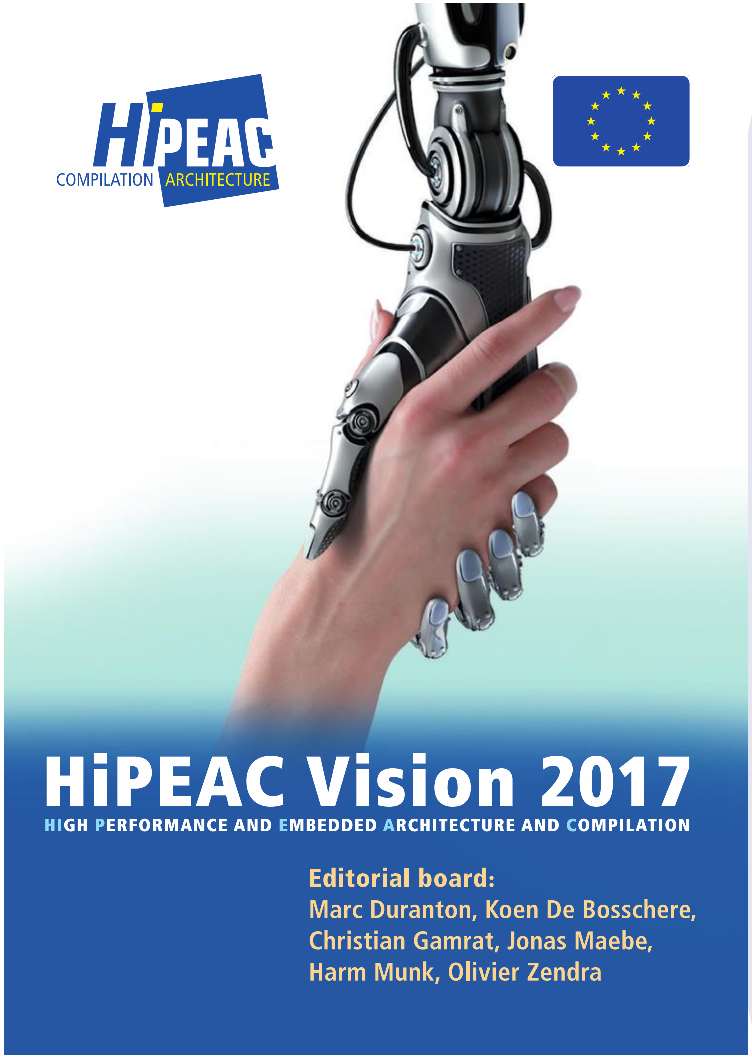

HIPEAC Vision 2017: IoT and the "edge"

- In IoT devices, communication inherent part; compromise between storage, communication and computation
- New disruptive technologies, like non-volatile storage, change the data storage and distribution landscape. This enables leaving the current cloud computing for the fog computing approach where most computation could also be performed locally, as edge computing.
- The universe of IoT systems will create ZetaBytes of data, but most of it will be 'dark data': data that is written once and never read again. The current 'gold rush' consists of trying to extract meaningful information from this data using various data analytics techniques.
  - Dedicated hardware or reconfigurable systems could lead to further improvement.
  - Secondly, moving to a hierarchical approach, in which first analyses are performed at or near the capture location, can save energy. eg. streaming analytics, where data is analysed on the fly (or "in the loop").

HIPEAC Vision 2017: IoT and the “edge”

The IoT “jungle” – beware of divergence and dispersion…

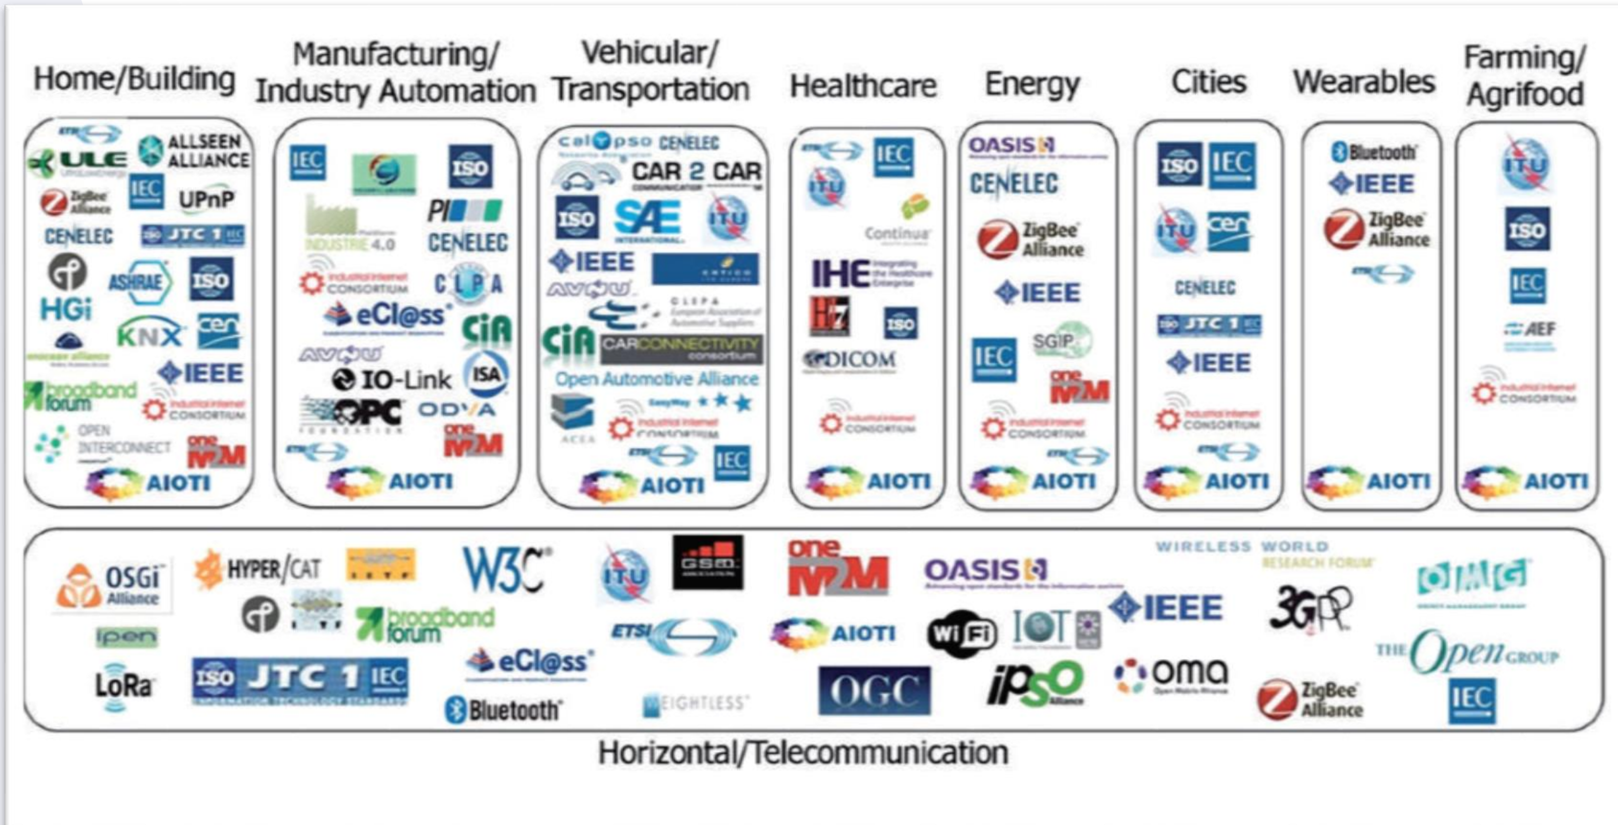

Urgent need for standards and “shaping strategy”.

This is the end, my friend…

[The Doors, 1967]

The end of the centralized computing model?

- The rise of the edge is imminent! (see previous slides)
- Fundamental asymmetry: we can cache downstream (centre-to-edge) but **not** upstream (edge-to-centre) – at least not on low-cost, low-energy devices.
- More and more large scientific instruments are out on the edge, eg. SKA, LIGO, LSST, weather/climate sensors, satellites, seismic arrays, etc.
- What does this imply for future roadmapping in Europe?

So, where do we stand in Europe?

- In the (pure) exascale race, we're 2-3 years behind… (IMHO)
- We need to concentrate on domains where we have (consensual) strengths and move forward on these:
  - federated service platforms for convergence,
  - algorithm and software development for convergence,
  - analytical methods and tools for convergence,
  - applications and workflows using convergence.
- The ETP4HPC-SRA, BDVA, PRACE, EDI are perfect instruments for this.
- EXDCI(2) will play a central role in coordinating this effort, and the future BDEC initiatives...

BDEC REPORT

BDEC = Big Data & Exascale Computing

- Objective: elaborate international roadmap and guidelines for achieving ***convergence*** between HPC and HDA
- BDEC holds 1-2 meetings per year since 2013 (US, Japan, EU)
  - Frankfurt meeting (@ISC'16), BoF (SC'16) organized by EXDCI
  - Latest meeting (#5) was hosted in Wuxi (for the 1st time!) by China in March 2017.
  - Next meeting (early 2018) to be held in USA

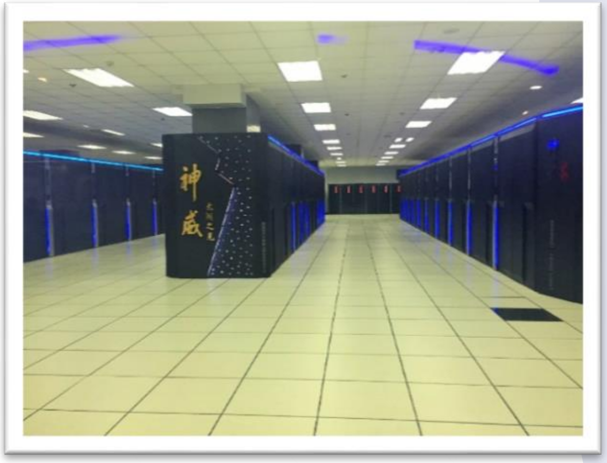

BDEC report: Pathways to Convergence

1. Split in 2 paradigms
2. App, workflow convergence
3. Centre convergence
4. Edge convergence
5. Conclusions, recommendations

The BDEC "Pathways to Convergence" Report

BDEC Committee

June 16, 2017

Contents

1 Introduction

The story of this report, and of the series of Big Data and Exascale Computing (BDEC) workshops that it summarizes, is at least in part the story of the high performance computing (HPC) community's response to two major inflection points in the growth of scientific computing in the 21st century. The first marks the disruptive changes that flowed from the end of Dennard Scaling, c. 2004 Dennard et al. [15], which gave rise to the era of multicore, manycore, and accelerator-based computing, as well as a variety of other complex and closely related problems of energy optimization and software complexity. The second, occurring nearly simultaneously, was the relatively rapid emergence (c. 2012) of "Big Data" and large scale data analysis as voracious new consumers of computing power and communication bandwidth for a wide range of critical scientific and engineering domains. Because the BDEC community has its roots in traditional HPC, the marked difference in the ways in which this community has struggled to absorb and adapt to these two watershed transitions, with varying degrees of success, provides essential context that informs a reading of this report.

To those who witnessed previous transitions from vector and shared multiprocessor computing, the response to the end of Dennard scaling was comparatively familiar and straightforward. Indeed, the dramatic effects of the end of Dennard scaling on processor and system designs were very much on the mind of the group of HPC leaders who gathered at the annual Supercomputing conference in 2009 to form what would become the International Exascale Software Project (IESP) [16]. Along with the European Exascale Software Initiative (EESI)[1] and a parallel effort in Japan, these grass roots exascale efforts were the progenitors of the BDEC.

Although the challenges that confronted the IESP"s vision of exascale computing were unquestionably formidable—orders of magnitude more parallelism, unprecedented constraints on energy consumption, heterogeneity in multiple dimensions, and resilience to faults occurring at far higher frequencies—they all fit comfortably within the problem space and the ethos that defined traditional HPC. Accordingly, participants

[1] http://www.eesi-project.eu

The way forward…

From “roadmaps” to “shaping strategies”

- Q: How do you draw a roadmap when the terrain is moving under your feet?
- A: You don’t – you adopt a “shaping strategy” (Deloitte)
  - Shaping View
  - Shaping Platforms
  - Shaping Acts and Assets

Thank you!

- Please consult www.exascale.org for reports and presentations from all the BDEC meetings, where you can download an early draft of the document "Pathways to Convergence".
- "Pathways to Convergence" report will be officially presented at Supercomputing 2017 in Denver USA.
- You can contact us at:
  - mark.asch@u-picardie.fr
  - tmoore@icl.utk.edu
- **Questions and/or comments?**