CHALMERS

EUROLAB-4-HPC

Eurolab-4-HPC Roadmap

Paul Carpenter

Barcelona Supercomputing Center

Theo Ungerer

University of Augsburg

Agenda

- EuroLab-4-HPC
- Roadmap Scope, Organisation and Status
- EuroLab-4-HPC Roadmap Topics
- Discussion

Mission: EuroLab-4-HPC will join HPC Systems Research in Europe

EUROLAB-4-HPC

Gap:

- Key HPC stakeholder community (research, suppliers, venture capital) is fragmented and uncoordinated
- Uncoordinated research community

What is needed is

- Join research and other stakeholders around a common long-term research agenda
- Train future technology leaders
- Accelerate innovation

Vision: A European HPC systems virtual Centre of Excellence

EuroLab-4-HPC Project

EUROLAB-4-HPC

- H2020-FETHPC-2014 (same as EXDCI)
- Budget €1.5 M
- Sep 2015 to Aug 2017

CHALMERS

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

HLRS

The University of Edinburgh

Universiteit Gent

INRIA

RWTH AACHEN
RHEINISCH-WESTFÄLISCHE TECHNISCHE HOCHSCHULE AACHEN

École Polytechnique Fédérale de Lausanne

ETH

FORTH Institute of Computer Science

Universität Augsburg University

Technion Israel Institute of Technology

MANCHESTER 1824
The University of Manchester

Objectives

- **join HPC system research groups** around a long-term HPC research agenda by forming an HPC research roadmap and joining forces behind it
- **define an HPC curriculum** in HPC technologies and best-practice education/training methods to foster future European technology leaders
- **accelerate commercial uptake** of HPC technologies
- **build links** between the HPC research community and other stakeholders (suppliers, venture capital, etc.)
- **form a business model** and organization for a future Centre of Excellence on HPC systems

Alignment around a Long-term HPC Systems Research Agenda

- Moore's Law is running out of steam
- Beyond Exascale computing systems
- Disruptive technologies across the stack
- Look ahead 10-15 years

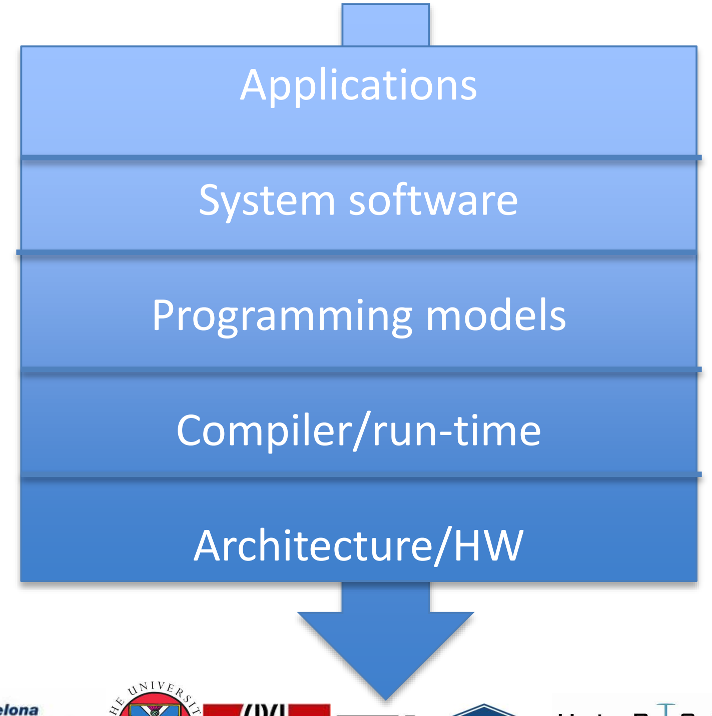

Cross-cutting issues:
performance; energy, dependability, programmability, scalability, etc.

EUROLAB-4-HPC ROADMAP

Roadmap Scope

EUROLAB-4-HPC

- Long-term vision for excellence in European HPC research
- Beyond Exascale targeting 2022—2030
- Include all layers of HPC stack, from applications to hardware
  - Consider adjacent domains: high-performance embedded, data centres, big data
- Close collaborate with HiPEAC Vision and ETP4HPC/EXDCI Strategic Research Agenda
- August 31, 2016: Preliminary roadmap
  - Complete, 40 pages + appendix
  - Will be public after Oct 2016 review
- August 31, 2017: Final roadmap

Roadmap Approach

- Because targeting 2022-2030 will be highly speculative:

1. Select disruptive technologies that may be technologically feasible in the next decade
2. Assess the potential hardware architectures and their characteristics.
3. Assess what that could mean for the different HPV aspects.

The roadmap itself will follow the structure:
"IF technology suitable
THEN foreseeable impact on WG topic could be"

The 5+1 Working Groups

EUROLAB-4-HPC

WG0: Disruptive Technologies (leader: Theo Ungerer, U. Augsburg)

WG1: New technologies and hardware architectures
(leader: Avi Mendelson, Technion, Haifa)

WG2: System software and programming environment
(leader: Paul Carpenter, BSC, Barcelona)

WG3: Vertical challenges: Green ICT, energy and resiliency
(leader Axel Tenschert, HLRS, Stuttgart)

WG4: HPC applications: evolution and requirements
(leader: Paul Carpenter, BSC, Barcelona)

WG5: Convergence of embedded HPC, data centers for big data, and HPC (leader: Babak Falsafi, EPFL, Lausanne)

Current State

- 2016-02: Working groups started work after selecting WG leaders and kickoff telcos
  - 46 roadmapping contributors plus additional technology experts signed up
  - More active ones are welcome!
- 2016-04: Input to FET Proactive Consultation on „Game Changing Technology“
- 2016-07: First Report on Disruptive Technologies available
- 2016-08: Preliminary Roadmap Deliverable submitted
- Final Roadmap under preparation

First Results on Roadmap

- Currently strong efforts are taken in US and Europe towards the Exascale computer
- Further applications emerged besides the high-performance computing (HPC): big data applications and embedded HPC
- Will a convergence of embedded HPC with big data applications and the HPC centers happen?
- HPC will span from on-board computers in car to HPC centers.
  - HPC roadmap scope not restricted to supercomputers
- Computer systems already consume a large part of our natural resources. Green ICT for sustainability will be mandatory.

First Results on Roadmap after the WG Discussions

- Future HPC development will be pushed by new technologies and pulled by new applications, which ones are open?
  - Potentially disruptive new applications, memory hierarchy, hardware accelerators
- However, legacy engineering applications will be continued and have to be adapted to new technologies.
  - Evolutionary scaling for strong scaling applications
  - Not all current applications will scale (weak scaling)
- Roadmap should provide technical background and formulate questions instead of solutions

WG0: Disruptive Technologies (leader: Theo Ungerer, U. Augsburg)

- **Sustaining Technology** (improving HPC HW in ways generally expected)
  - Continuous CMOS Scaling
  - Die Stackaing - 3D-Chip
- **Disruptive Technology in Hardware/VLSI** (innovation that creates a new line of HPC HW superseding existing HPC techniques):
  - Non-volatile Memory (NVM) Technologies
  - Photonics
- **Disruptive technology** (alternative ways of computig)
  - Resistive Computing
  - Neuromorphic Computing
  - Quantum Computing
- **Beyond CMOS**
  - Nanotubes
  - Graphene
  - Diamond

Summary of Potential Long-Term Impacts of Disruptive Technologies for HPC Hardware

- **Processor Logic**
  - if photons, graphene, or nanotube would replace silicon transistors
  - Then (very speculative): much higher clock rates and less heat expected and totally change of current way of computing.
  - Current CMOS technology may continuously scale also in next decade with increasing costs per transistor, power consumption and to less reliability.
  - Die stacking could lead to 3D many-core microprocessors with reduced wire length.
- Memory Hierarchy
- Potential New Hardware Accelerators

Summary of Potential Long-Term Impacts of Disruptive Technologies for HPC Hardware

- Memory Hierarchy
  - 3D stacking will also be used to scale flash memories by 3D flash memories.
  - whole memory hierarchy may change in the upcoming decade with new technologies, in fact memristors, which will deliver non-volatile memory potentially replacing or additional to DRAM.
  - Memristors offer orders of magnitude faster read/write accesses and also much higher endurance than flash.
- Potential New Hardware Accelerators

Summary of Potential Long-Term Impacts of Disruptive Technologies for HPC Hardware

- **Potential New Hardware Accelerators**
  - Resistive Computing could enforce datacentric and reconfigurable computing.
  - Neuromorphic Computing would be efficient in energy and space for artificial neural network applications.
  - Quantum Computing might even solve some problems that couldn't be solved with classical computers with important implications for public-key cryptography, searching, and a number of specialized computing applications.

WG1: New Technologies and Hardware Architectures (leader: Avi Mendelson, Technion, Haifa)

- Principal questions and research challenges for future HPC hardware architectures
  - Impact if power and thermal will not be limiter anymore (due to new materials)
    - Frequency increase vs many-cores
  - … if Dark Silicon not happen
  - … if communication becomes so fast that locality will not matter
  - … if data movement could be eliminated (and so data locality)
  - … if memory and I/O could be unified and efficiently be managed.

WG2: System Software and Programming Environment (leader: Paul Carpenter, BSC)

- Programmer support for complexity
  - Accelerators, deep memory hierarchies, FPGAs, heterogeneous, timing variability
  - Elastic / dynamic execution environments
  - Programming models: tasks, DSLs to separate functionality from implementation
  - Compilers and common runtime systems: load balancing, explicit data transfers, optimized memory layout
- Complex application performance analysis and debugging
  - Will become intractable for humans => machine learning / AI
  - Programmer needs cost model for intuition on performance
  - Distance between the performance problem and application code
- New hardware models of computation
  - e.g. how neuromorphic and quantum computing will affect system software
- Cluster management
  - Resilience, fault prediction, ...

WG3: Vertical Challenges: Green ICT, Energy and Resiliency (leader Axel Tenschert, HLRS, Stuttgart)

- Green ICT
  - Energy efficiency: efficient software stacks, libraries
  - Heat reuse
  - Environmental impact: reducing CO_2, life cycle assessment
- Energy
  - NVM, new materials
  - Improving PUE
- Resiliency
  - Shrinking increases failure rates: 10s atoms per transistor
  - Enormous number of components
  - Compiler level resilience, reliability in runtime, programming models, ABFT

WG4: HPC Applications: Evolution and Requirements (leader: Paul Carpenter, BSC)

- Collect requirements from HPC in science and engineering
  - Including scientific libraries and numerical algorithms
- Working with new PRACE Scientific Case
  - Major effort by EXDCI on HPC science and engineering roadmap

WG4: HPC Applications: Evolution and Requirements (leader: Paul Carpenter, BSC)

- How to cope with complex hardware?
  - Scalability to huge number of nodes (overheads and variability amplified)
  - Conservative in adopting new programming models and DSLs
  - Need to be convinced of long-term support due to long code lifetime
- In situ data visualization, analysis and interactivity
  - Difficulty of moving data between machines
  - Database, and batch & interactive jobs need to coexist
  - Also HPC in the cloud => Elastic resources (doesn't work with MPI)
- What is the „MapReduce“ of scientific computing?
  - No such simple answer
  - Decomposition-driven communication works in some domains
- Load balancing on multiphysics applications
- Bit reproducability is a major problem
- Non-functional properties: want to easily tradeoff accuracy vs cost

WG5: Convergence of embedded HPC, Data centers for Big Data, and HPC (leader: Babak Falsafi, EPFL)

- Big data: two trends changing HPC

- 1. Data-centric science
  - Data analytics complement simulation
  - Science entering 4th paradigm
    - 1: theory, 2: empirical, 3: simulation
  - instrument data (LHC/SKA), simulation data, sensor data, human data

- 2. Warehouse-scale computers
  - Commodity systems and very low prices
  - Increasingly shared concerns: enterprise – HPC
    - Scale out, dynamic resource management
    - High utilization, parallelization (multicore, GPU,FPGA)
    - Low latency interconnect, application resiliency
    - Infrastructure costs, economies of scale

WG5: Convergence of embedded HPC, Data centers for Big Data, and HPC (leader: Babak Falsafi, EPFL)

- Embedded HPC: embedded in physical environment
  - Automotive, aircraft, smart grids, traffic management, civil engineering, intelligent transportation, algorithmic trading
- Performance of HPC from a few years ago
  - But also: time criticality, functional safety, energy efficiency, reliability
- Deep learning important in CPS
- Many-core not efficiently used due to complexity of parallel programming

Main open technical questions

- How will large non-volatile memory impact HPC applications and software?
  - Storage class memory may fill the gap between memory and flash/disc drive
- How to use new technologies in software (NVM, 3D stacking, neuromorphic, accelerators)
- Evolution of MB/core, memory & interconnect BW and latency, …
- What are disruptive ideas in software and applications?
- Any use for special-purpose accelerators like resistive computing (near memory), neuromorphic computing, quantum computing?

Discussion points

- How to cooperate with ETP4HPC and PRACE?
  - Collect application requirements
  - Technical alignment on roadmap topics

BACKUP

Comparison of HPC Roadmaps

	Goal	Timespan	SWOT / political	Scope	Num. pages
HiPEAC vision	Steer European academic research (driven by industry)	"short": 3 years, "mid": 6 years, long: >2020	Y	HPC + embedded	72
ETP4HPC SRA / EXDCI	Strengthening European [industrial] HPC ecosystem	6 years (2014 to 2020)	Y	HPC except applications	92
PRACE Scientific Case	[Academic] need for European HPC infrastructure	8 years (2012 to 2020)	Y	HPC applications	159
EESI (European Exascale Software Initiative)	Development of efficient exascale applications	5 to 10 years	N	Exascale applications	34
BDVA (Big Data Value Association)	Big Data technologies roadmap	2020	-	Big data	45
Rethink Big	Roadmap for European Technologies in Hardware and Networking for Big Data		-	Big data	
ECSEL MASRIA	European leadership in enabling and industrial technologies. Competitive EU ECS industry.	2015 roadmap to about 2025	Y	Electronic components and systems (ECS)	91
Eurolab4HPC	**Academic excellence in HPC**	**8 to 10 years**	**N**	**Whole HPC stack**	**20—25**