Exascale software

technical point of view

Sebastian von Alfthan

CSC – Tieteen tietotekniikan keskus Oy
CSC – IT Center for Science Ltd.

Exascale

- Achievable by 2019
- Major software challenge – paradigm shift away from MPI?
- Heterogeneous computing
- Several project addressing these issues on EU and national level

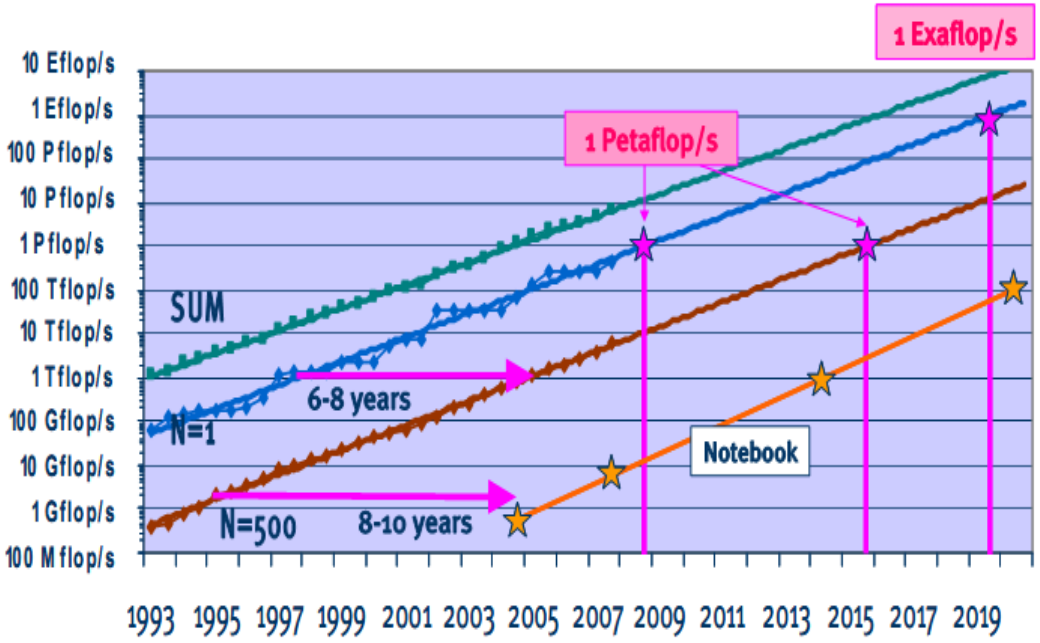

From top500.org:
http://www.top500.org/files/TOP500_Looking_back_HWM.pdf

Challenges

1. Power
   - How to fit inside power budget?
   - Extrapolating current generation processors not good enough; heterogeneous processors
   - Sustaining memory bandwidth requires too much power (from 0.3 to 0.003 bytes/flops ?)
2. Concurrency
   - Billion-way parallelism in applications needed

Challenges

3. Resiliency
   - Resiliency against faults & performance variation
   - Checkpointing difficult, Bandwidth to disks trailing (far) behind
4. Programming
   - Programming models coping with concurrency, heterogeneous hardware

Hardware trends

- Per core speed not increasing
  - Processors with hundreds to thousands cores
  - Single thread performance still important (Amdahls law)
  - Hybrid manycores integrating accelerators appears unavoidable
- Data movement power hungry
  - Difficult providing bandwidth to cores
  - Stacked memory?

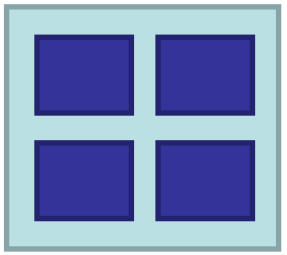

All large core

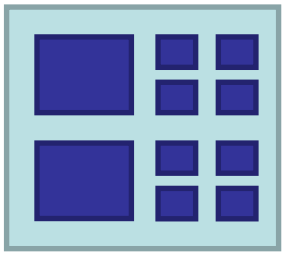

Hybrid manycore

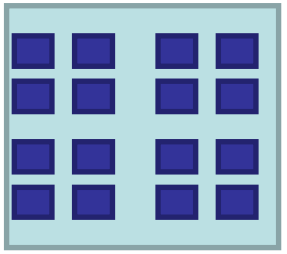

All small core

Hardware trends

- Network
  - Optical networks – low power, high bandwidth
  - Enables high radix routing (Cray)
  - Bytes/Flops will be orders of magnitude lower
- I/O growth rate slow
  - Checkpointing difficult (MTBF vs. write time)
  - Additional memory layer needed?

Software challenges

- Extreme parallelism needed
  - Coping with locality, overheads, latency and contention
  - Low GB/Flops for memory, I/O & Interconnect
- Support for heterogeneous processors
- Resiliency
- Ease of use!

Scaling applications

- Strong scaling
  - Will in practice not be possible…
- Weak scaling
  - Scaling up seems tractable
  - Larger problem or increased resolution
  - Timescales also tend to increase…

Scaling applications

- Coupled multiple image application
  - Stochastic simulations
  - Coupled models
  - Replica simulations
  - Multiphysics

Programming models

- MPI
  - Will not scale – paradigm shift at hand
- Mpi+Threads/accelerators
  - MPI+StarSS (http://www.project-text.eu) / OpenMP
  - MPI cuda / opencl / HMPP / PGI...
- PGAS (upc, caf, x10, chapel)
  - The compiler & runtime can do optimization
- Heterogenity – need to support different compute elements in a portable way
- New models?

Conclusions

- Software platform
  - programming models enabling all levels of parallelism in applications to be expressed
  - “Smart” compilers / run-times /libraries enabling these too run on a heterogeneous platforms with different characteristics
  - Resiliency handled by run-times / system
- Prototypes of hardware platforms

Conclusions

- Applications
  - European scientific communities (& industry) engaged to develop key applications for Exascale
  - Tuning vs. new applications