An overview of technological trends in supercomputing

Fabio Affinito (f.affinito@cineca.it)

SCAI - Cineca
Italy

~~An overview of technological trends in supercomputing~~
New challenges in HPC and what we need to survive

Fabio Affinito (f.affinito@cineca.it)

SCAI - Cineca
Italy

The future

- Number of cores in HPC is continuosly increasing
- Going parallel is a must
- Codes should keep the pace in order to survive
- … but HPC is not everything

A lot of challenges

- Different paradigms of parallelism
  - MPI
  - OpenMP
  - tasking
  - asyncronism
- High-throughput
- New architectures
- GPUs?
- Intel Xeon Phi?

What I've heard in the last 2 days...

- many interesting developments
- some of them tackling the same physical problems
- different kinds of algorithms
- partially sharing the same parts (modules) of code
- documentation? how to do that?
  - for example using Ford...

but...

- it is so difficult to share the work?
- so it is much more harder to face the HPC challenges
  - for example, PW is ready to work on new parallel machines, but is it also PH and other smaller codes?
- we want to make QE more modular, but how?

What is needed?

- a common interest
- a strong commitment
- some good tools
  - for development
  - for testing
  - for the integration

A starting point - IMHO

- writing (good) code is not enough
  - we need to test the code
  - we need to maintain the code
  - we need to document the code (and support the users)
- we need tools (not necessarily new languages or programming models)
- we need good practises

Start from what already exists

- SVN permits the creation of branches
- using branches is sometimes though
  - a branch should be ALWAYS be in sync with the main trunk
- but it permits to incrementally modify the code
- once a branch is well tested and maintained can be then merged with the trunk
- this mechanism permits to avoid disruptive changes, keeping the control on the new developments

So, why not?

- in the past, forks of the trunk often achieved nothing (do you remember PRACE?)
- now we have momentum
  - and this meeting is the a good evidence
- and we have resources
  - MAX CoE is strongly supporting the development of QE

Thanks for your attention

Backup (original) slides

HPC trends

Dennard scaling law

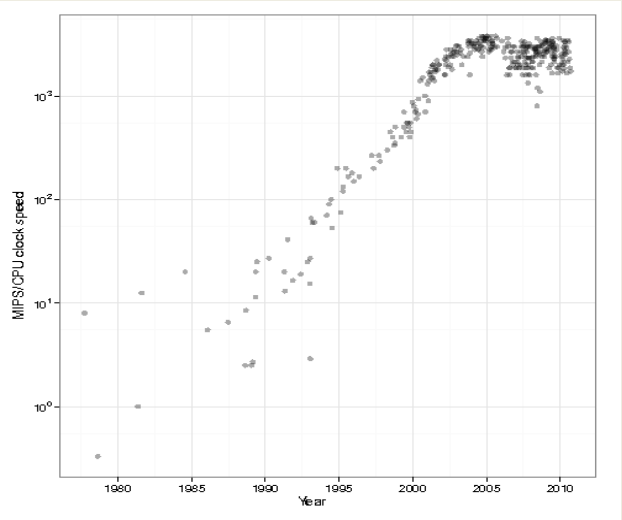

- *The core frequency and performance do not grow following the Moore's law any longer*
- *Increase the number of cores to maintain the architectures evolution on the Moore's law*

Moore 's Law

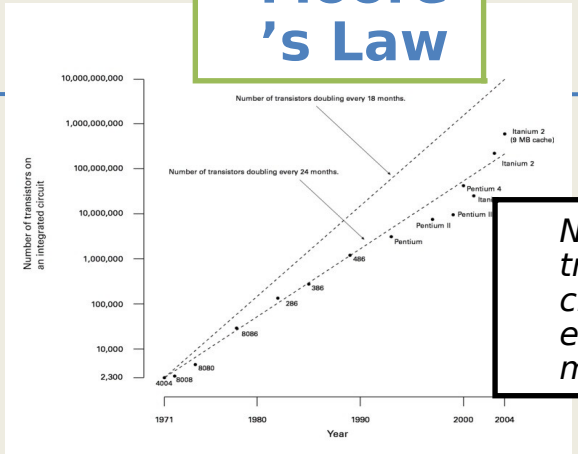

Number of transistors per chip double every 24 month

Amdahl's law

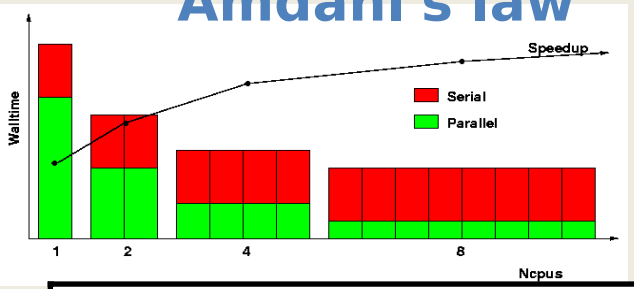

maximum speedup tends to $1 / (1 - P)$
P= parallel fraction

The upper limit for the scalability of parallel applications is determined by the fraction of the overall execution time spent in non-parallel operations.

European Exascale projects

DEEP is based on the concept of a cluster-booster architecture implemented using Intel Xeon processors and Xeon Phi coprocessors. It implements OmpSS as programming paradigm

Mont-Blanc is looking for energy efficient solutions for exascale. The concept is implemented within a cluster of ARMv8 64-bit processors.

The CORAL project

- US ~~exascale~~ *extreme scaling* computing project, involving Argonne (ANL), Oak Ridge (ORNL) and Livermore (LLNL)
- LLNL and ORNL will install a IBM OpenPower machine based on Power9 + Nvidia GPUs (2 or 4 per node), connected with Nvidia Nvlink technology
- ANL will install a machine based on the Intel Xeon Phi processors (Knights Hill) with Intel Omnipath interconnection

Announced roadmaps

Xeon multi-cores:

- 2016 Xeon Broadwell EP
  - 14nm technology (shrink wrt Haswell)
  - up to 22 cores and 44 threads
- 2017 Xeon Skylake EP
  - 14nm technology (tock)

Xeon many-cores:

- Q3 2016 Xeon Phi x200 (KNL)

- 2016: Pascal
  - High bandwith memory stacks
  - NVlink
- 2018: Volta (10nm)
  - 3D memory
  - 10nm shrink

Intel Xeon Phi

Intel Xeon Phi is a multi-core co-processor implementing the MIC (many integrated cores) architecture.

The first generation of Xeon Phi, codenamed Knights Corner (KNC), was released in 2012. The KNC behaved as co-processor, being coupled through a PCIe to the host CPU, similarly to the GPUs.

The main features of the KNC were:

- up to 61 in order x86 cores
- internal ring interconnect
- up to 16GB DDR4 on board
- 512-bit SIMD registers

Knights Landing

Holistic Approach to Real Application Breakthroughs

intel inside XEON PHI

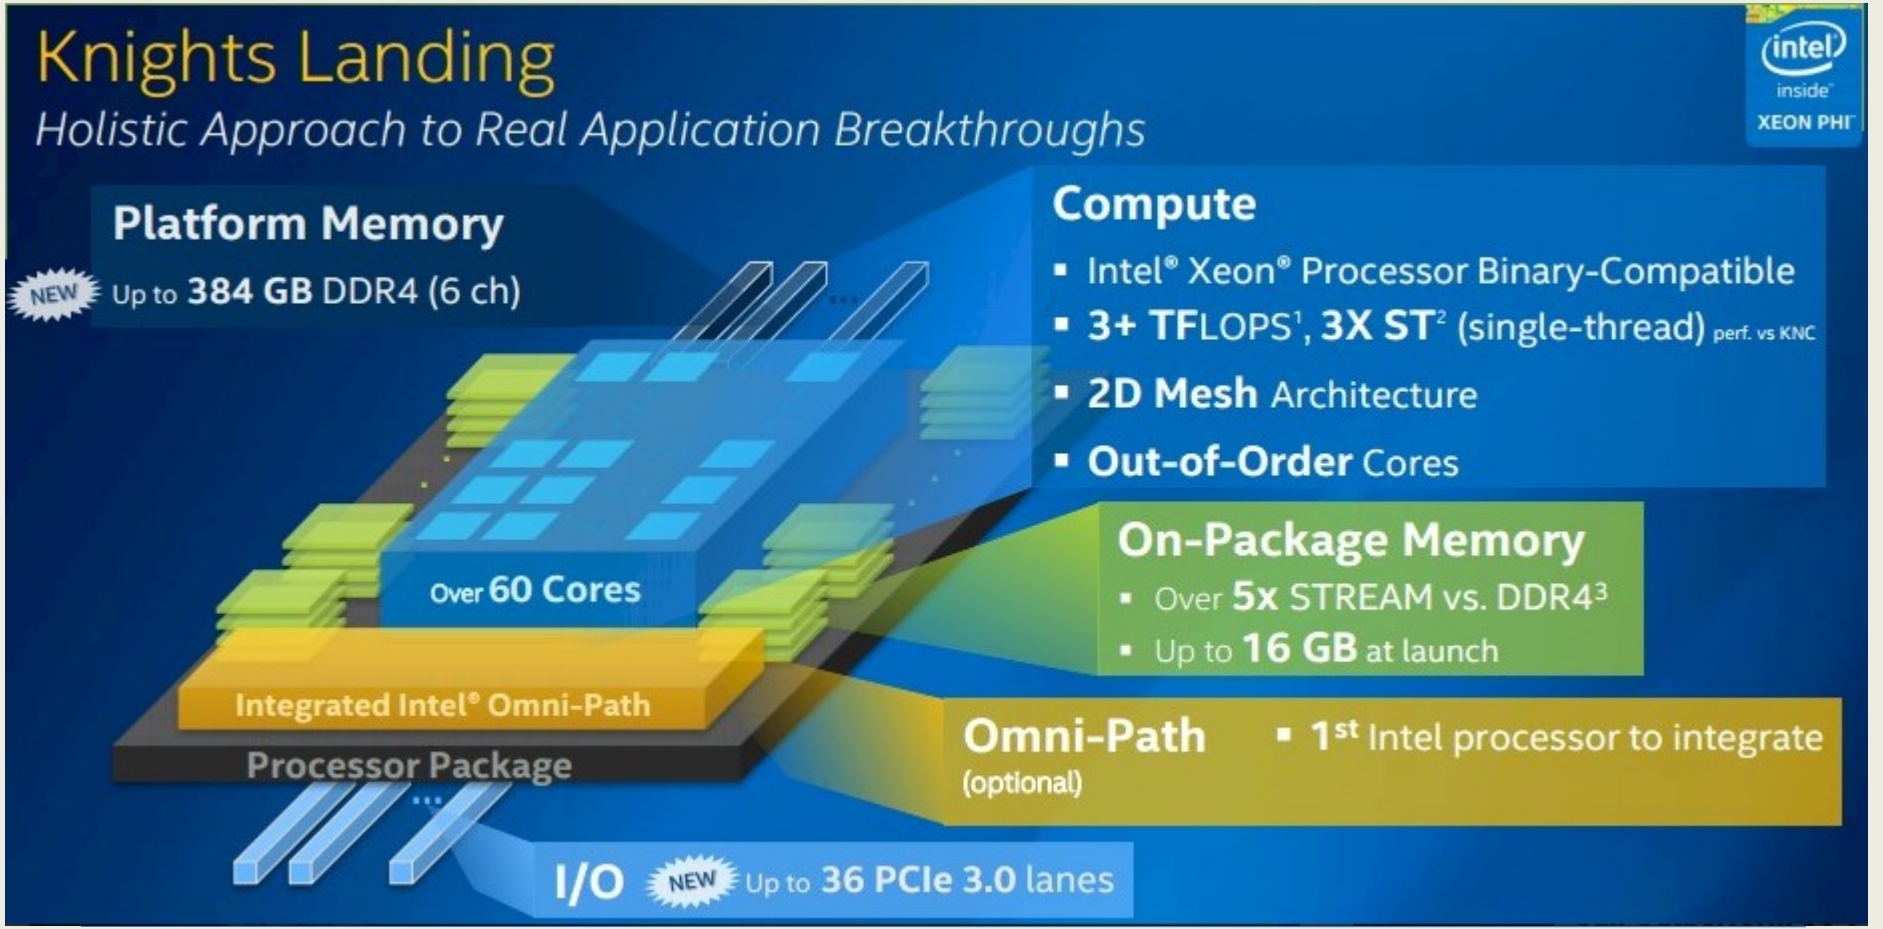

KNL in a nutshell

- 72 atom-based cores out-of-order
- mesh internal interconnect
- stand-alone processor
- up to 16Gb high-bandwith on-package memory (MCDRAM)
- 6 channels DDR4, up to 384Gb
- > 3TF DP
- AVX512

Challenges and opportunities

- Increase the parallelism
  - MPI: hierarchically, suitable to fit different architectures (node, cluster..)
  - OpenMP threads in order to exploit shared memory intranode
  - OpenMP tasks: to improve the scheduling, tackling latencies
- Exploit high-throughput simulations: ensemble techniques, data screening, etc.
- Beware of the numerical stability
  - now the number of processes is going up to O(100.000)...
- New algorithms?
- Don't forget energy efficiency...

Conclusions & perspective

- We are entering ~~the exascale~~ a new phase:
  - more nodes
  - more cores per node
  - poor single core performance
- We should live with it
- Enforce interactions within developers, users' community, computing centers, vendors
- Exploit tools, techniques, ask for support, never stop learning...

Good luck!

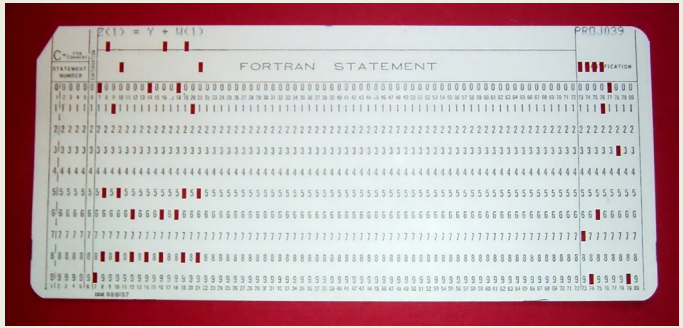

OpenACC
Directives for Accelerators

NVIDIA

python

intel

THE C++ PROGRAMMING LANGUAGE

OpenMP™

AiiDA