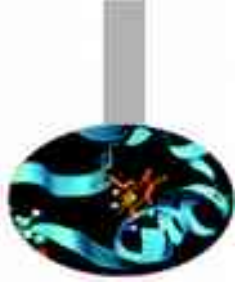

MD on HPC architectures:
GPU (Kepler-Pascal), Intel Xeon (KNL-SKL)

Alessandro Grottesi

Cineca

Today's lecture

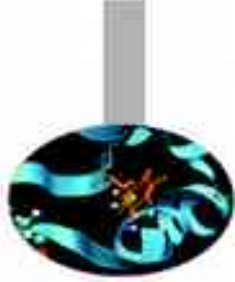

You will learn:

- Gromacs @ CINECA: set up and launch of simulations
- Launch MD code (GROMACS, NAMD, Amber)
- Optimize performance and benchmarking
- Tutorial (later this afternoon...)

Pico

Model: IBM NeXtScale server

Architecture: Linux Infiniband Cluster

Processors Type: Intel Xeon (Ten-Core) E5-2670v2 2.50 GHz (Compute)

Number of nodes: 54 Compute + 4 visualization + 2 Login + 14 other

Number of cores: 1080 (compute)

Number of accelerators: 4 + 2 (only on viz nodes)

RAM: 128 GB/Compute node (2 viz nodes with 512GB)

Galileo

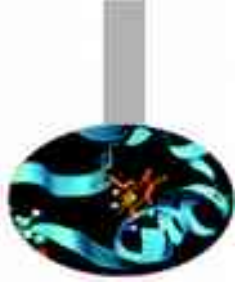

Model: IBM NeXtScale

Architecture: Linux Infiniband Cluster

Nodes: 516

Processors: 8-cores Intel Haswell 2.40 GHz (2 per node)

Cores: 16 cores/node, 8256 cores in total

Accelerators: 2 Intel Phi 7120p per node on 384 nodes (768 in total)

RAM: 128 GB/node, 8 GB/core

Internal Network: Infiniband with 4x QDR switches

Disk Space:2,500 TB of local storage

Peak Performance: ~1 Pflop/s (69th on TOP500)

Marconi A1

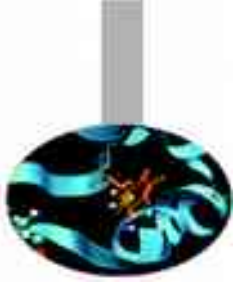

Model: Lenovo NeXtScale

Racks: 21

Nodes: 1.512

Processors: 2 x 18-cores Intel Xeon E5-2697 v4 (Broadwell) at 2.30 GHz

Cores: 36 cores/node, 54.432 cores in total

RAM: 128 GB/node, 3.5 GB/core

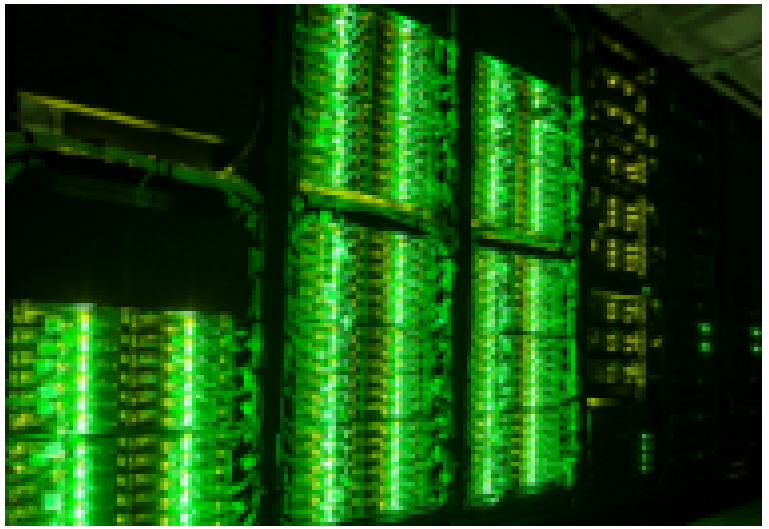

Peak Performance: 2 PFlop/s

Marconi A2

Model: Lenovo Adam Pass

Racks: 50

Nodes: 3.600

Processors: 1 x 68-cores Intel Xeon Phi 7250 CPU (Knights Landing) at 1.40 GHz

Cores: 68 cores/node (272 with HyperThreading), 244.800 cores in total

RAM: 16 GB/node of MCDRAM and 96 GB/node of DDR4

Peak Performance: 11 PFlop/s

Marconi A3

Model: Lenovo Stark

Racks: 21

Nodes: 1.512 + 792

Processors: 2 x 24-cores Intel Xeon 8160 CPU (Skylake) at 2.10 GHz

Cores: 48 cores/node 72.576 + 38.016 cores in total

RAM: 192 GB/node of DDR4

Peak Performance: 7.00 PFlop/s

Compute nodes (Galileo)

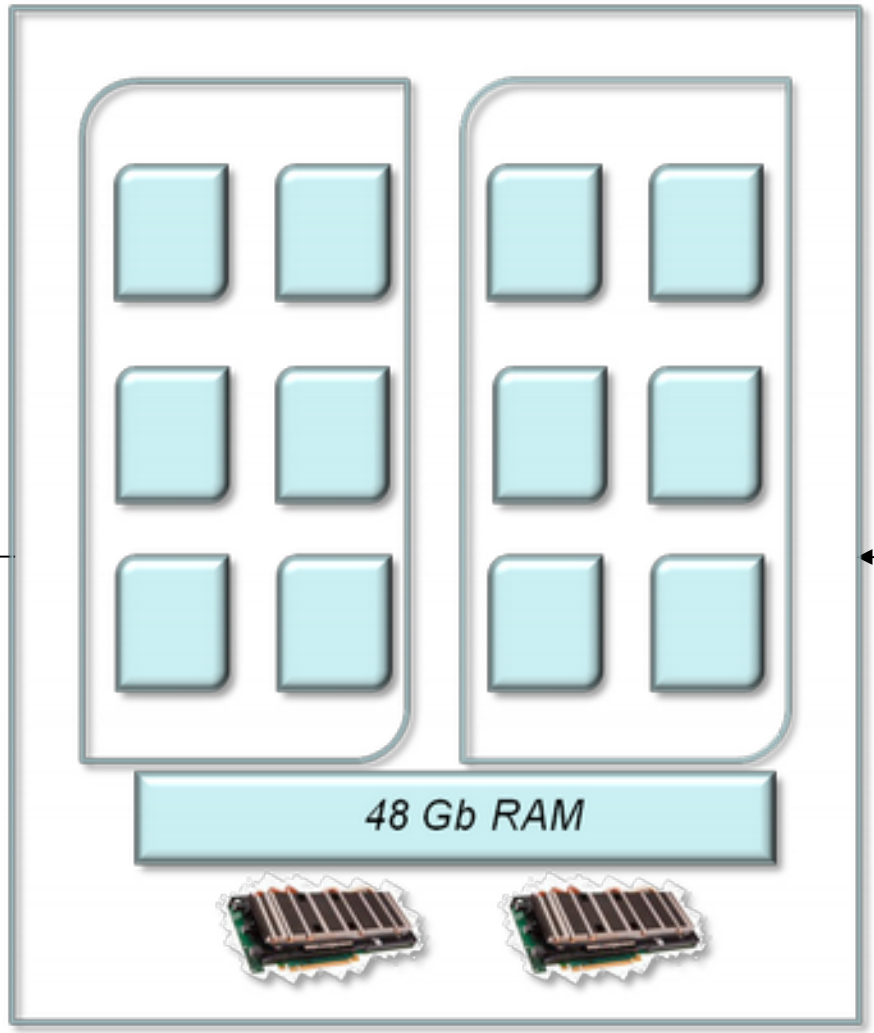

Infiniband connection

Molecular Dynamics and accelerators

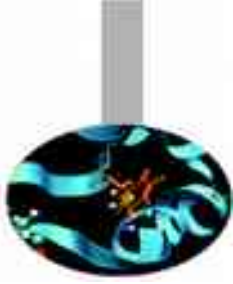

Intel Xeon Phi

Nvidia K80

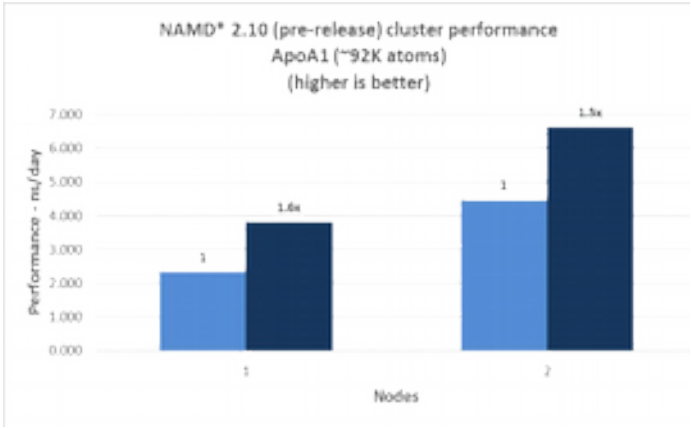

– GROMACS
  - Under development, currently (Jan 2015) only native-mode version available.

– NAMD
  - Pre-release version of NAMD 2.10 has Xeon PHI support but still under development. Speed-ups < 2 for ApoA1 and STMV benchmarks.

– LAMMPS
  - Xeon PHI support available in current downloads for non-bonded calculations. Reported speed-ups of about 1.78x compared to non-accelerated code (one coprocessor/node) for Rhodopsin benchmark. Higher speed-ups obtained with materials simulations.

– AMBER
  - Xeon PHI-enabled version released 6 Aug 2014. Waiting for benchmarks.

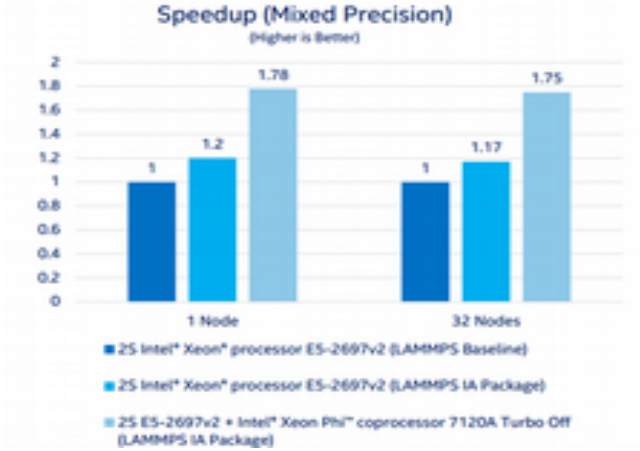

Intel Xeon PHI overview

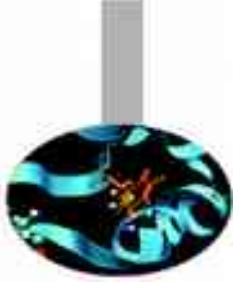

- Intel product line based on Intel's Many Integrated Core (MIC) technology where many low, power cores (>50) are packed on a single chip.
- Currently available device (Knight's Corner or KNC) can be seen as a co-processor, in direct competition to NVIDIA GPU for HPC.
  - connection to host CPU via PCI-eXpress link.
- But unlike GPU technology is not too dissimilar from host CPU → not essential to rewrite code, in principle just re-compile (with Intel compilers). Should lead to shorter development path in most cases.
- Doesn't mean though that code does need not porting – to obtain peak performance some optimisation is needed.

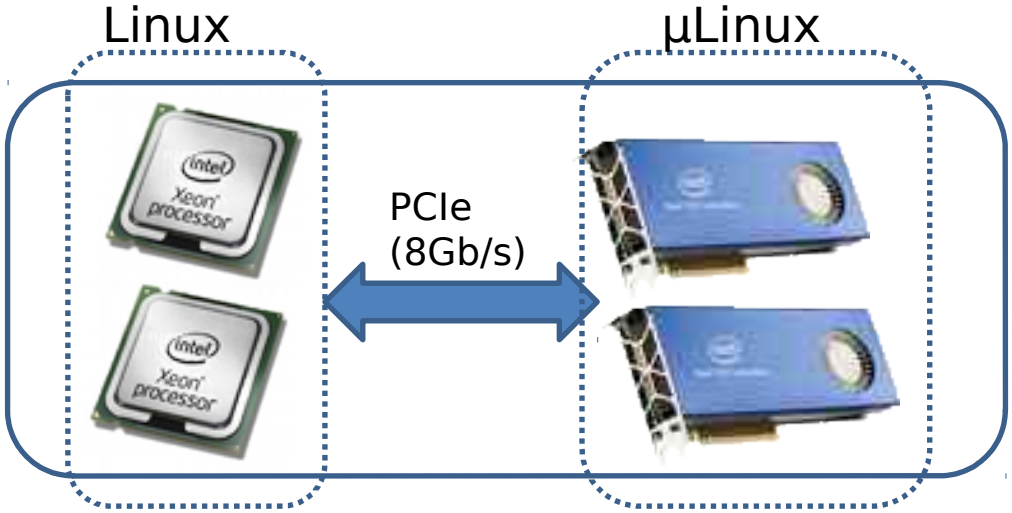

Classical Molecular Dynamics and Intel Xeon PHI

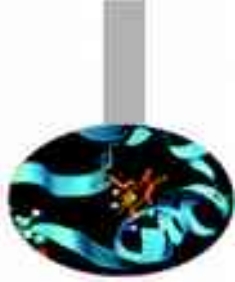

- Development some way behind GPU-CUDA versions of classical MD programs (which started about 4 years ago).
- But given that there is no need to rewrite in new languages (e.g CUDA) development path should be shorter.
- Most current Xeon PHI versions seem to be based on the off-load model to exploit CUDA developments.
- Off-loaded calculations invariably involve non-bonded dispersion interactions may also include PME, energy calculation etc.
- Intel maintains a list of “recipes” for building Xeon PHI applications (not just MD):

https://software.intel.com/en-us/articles/namd-for-intel-xeon-phi-coprocessor

Intel Xeon Phi: Knights Landing

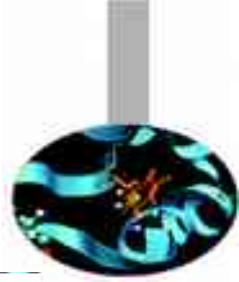

Knights Landing Overview

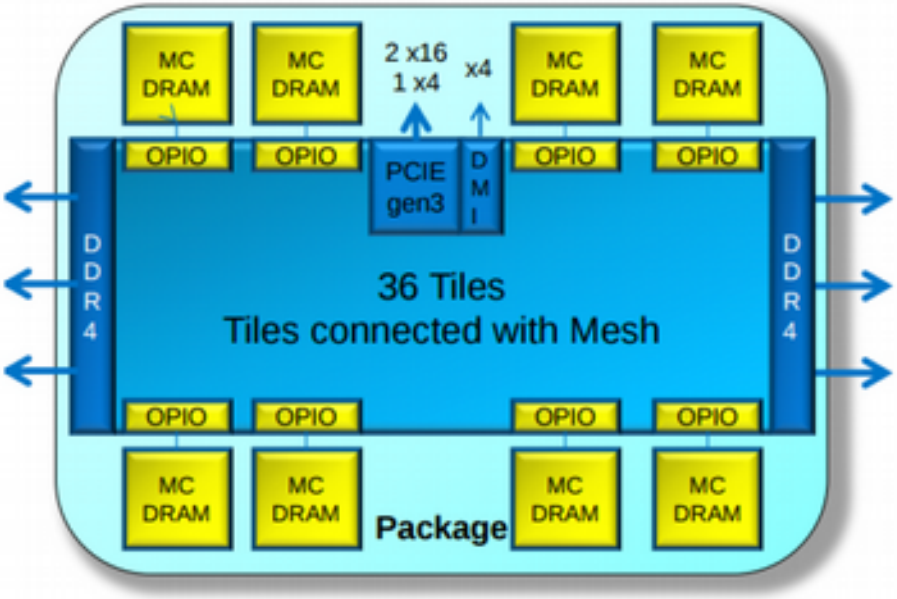

- Stand-alone, Self-boot CPU
- Up to 72 new Silvermont-based cores
- 4 Threads per core. 2 AVX 512 vector units
- Binary Compatible[1] with Intel® Xeon® processor
- 2-dimensional Mesh on-die interconnect

- MCDRAM: On-Package memory: 400+ GB/s of BW[2]
- DDR memory
- Intel® Omni-path Fabric

- 3+ TFLops (DP) peak per package
- ~3x ST performance over KNC

It's not a GPU. It's not an accelerator.
It's very different from a KNC.