

QE dev meeting

Trieste - January, 14th

Summary

- Cineca as IPCC
 - Porting of QE in native mode
 - Implementation of libxphi – offload mode
 - Status and perspectives
- Cineca infrastructure
 - PICO: a data-centric infrastructure
 - GALILEO: the replacement for PLX

Native mode

PW (SCF)

Quantum ESPRESSO* for Intel® Xeon Phi™ Coprocessor

Wed, Jun 4, 2014

[f Share](#) [Tweet](#) [g+Share](#)

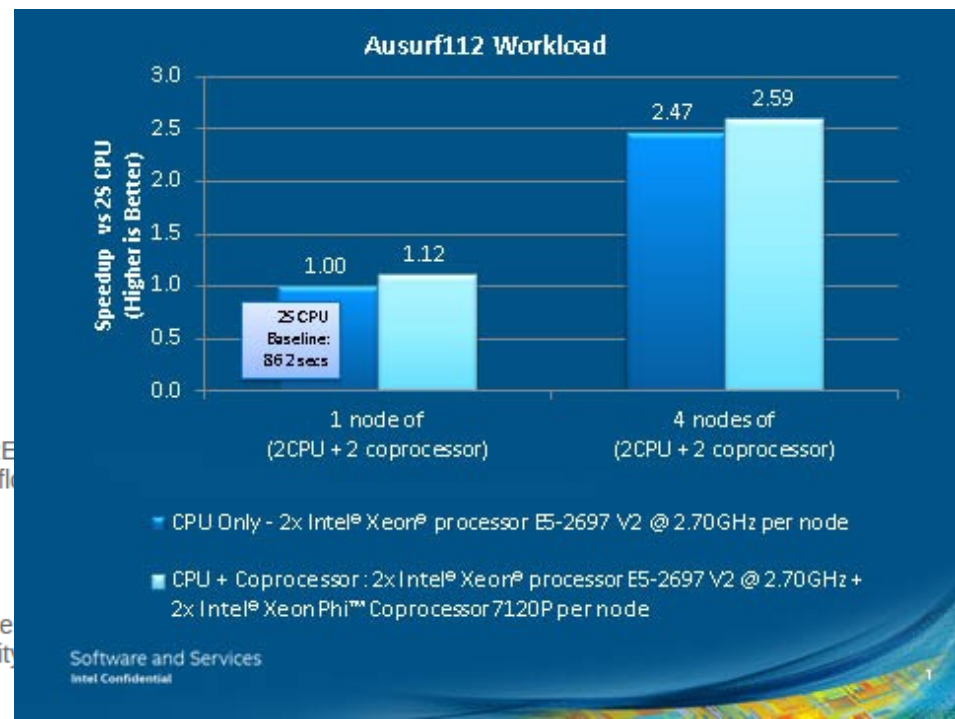


Purpose

This code recipe describes how to get, build, and use the Quantum ESPRESSO Phi™ coprocessor using the Intel® Math Kernel Library with Automatic Offload.

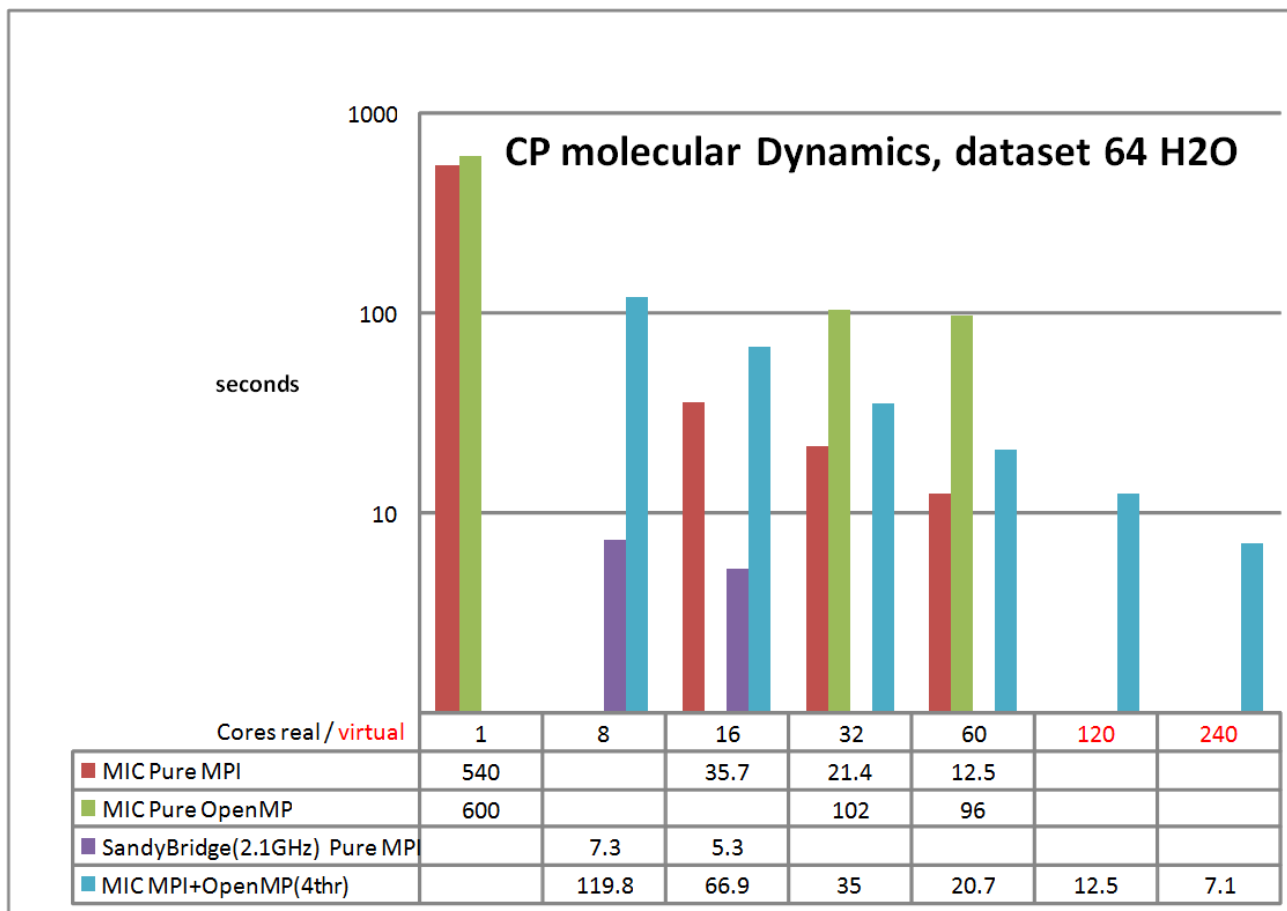
Introduction¹

Quantum ESPRESSO is an integrated suite of open source computer code calculations and materials modeling at the nanoscale. It is based on density functional theory, plane waves, and pseudo potentials.



Native mode

CP (Car-Parrinello)



In native mode it is also possible to use more than one Xeon PHI coprocessor at the same time, like if the parallel system were built using Xeon PHI alone. The good news is that the code works and it is really possible to run the application on more than one coprocessor, the bad news is that the performance are very bad due to communication latency and bandwidth limitation of the low level protocol used to communicate between coprocessor.

CP "native" performance, and scalability, comparison between different processor architectures, using 64 water molecules dataset

DGEMM: Xlibphi

Strategy:

- Implementation of a “double-buffer” technique to hide the latency between host and coprocessor
- Implemented through a dynamic library (plugin) that wraps the MKL function.

Status:

- Currently available on github (<https://github.com/cdahnken/libxphi>)
- Currently working to the documentation and how-tos (submitting scripts, etc.)

Host: Xeon Sandy Bridge E2670

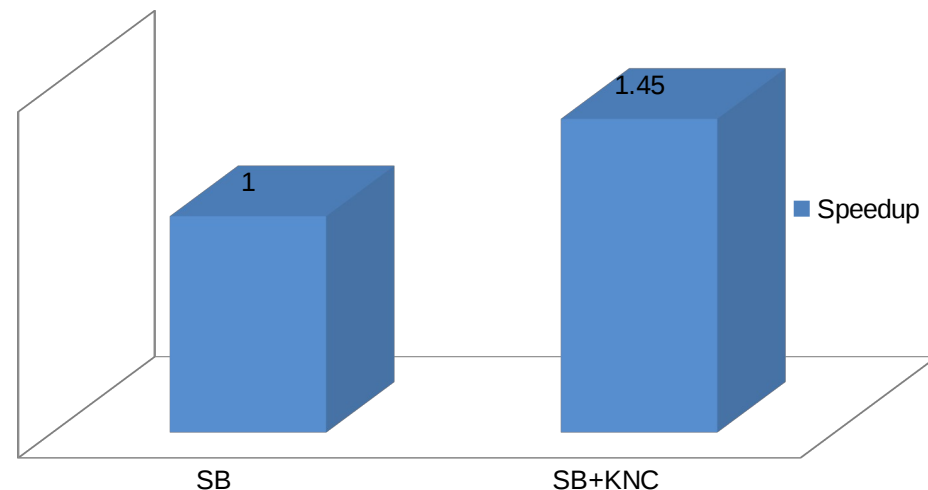
(2 sockets -> 16 cores)

Coprocessor: KNC 7120

(2 cards -> 2x61 cores)

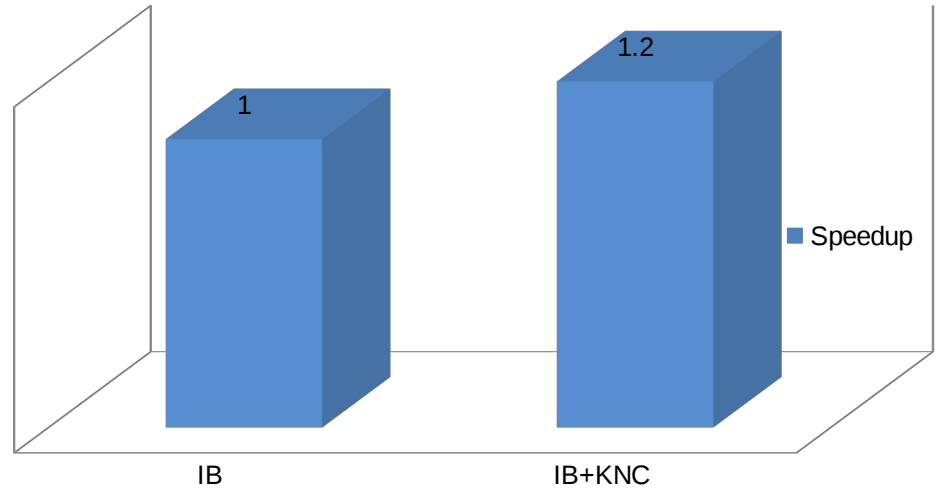
Ausurf 112:

Au complex 2,158,312 G-vectors, 2
k-points, FFT dimensions =
(180,90,288)




Xlibphi – benchmarks

Host: Xeon Ivy Bridge E5-2697
(2 x 12 cores)
Coprocessor: KNC 7120 (2
cards)



8 nodes equipped with:
Host: Xeon Ivy Bridge E5-2697
(2 x 12 cores)
Coprocessor: KNC 7120 (2
cards)



GRIR443:
Carbon-Iridium complex (C200Ir243),
2,233,063 G-vectors, 8 k-points, FFT
dimensions: (180, 180, 192)

Offloading FFT

FFT loop:

QE require to compute a 3D FFT for each electron of the system
-> thousands of parallel 3D FFT are computed at each iteration

Strategy:

decompose the parallel 3D FFT in sequence of 1D FFTs
explicitly manage (outside the FFT) the communication for data
reshuffling
offload the 1D FFTs to MIC
implement software pipelining to mask data transfer latency

Status:

first working implementation completed

```
do ib=1,nbands
  3D fft
enddo
```



```
do ib=1,nbands
  fft along z
end do
mpi transpose
do ib=1,nbands
  fft along xy
enddo
```

Parallel 3D FFT driver
fft along z
Internal mpi transpose
fft along xy



Parallel 1D driver
Global Transpose driver
Parallel 2D driver

(More memory requirements: 30% more)

Next steps:

- To complete the FFT offload. Then testing and preparing docs
- Address the diagonalization part
- Look at the multinode scaling

New Cineca infrastructure: PICO

A machine for data-analytics and post-processing.
Already in production.



	Total Nodes	CPU	Cores per Nodes	Memory (RAM)	Notes
Compute/login node	66	Intel Xeon E5 2670 v2 @2.5Ghz	20	128 GB	
Visualization node	2	Intel Xeon E5 2670 v2 @ 2.5Ghz	20	128 GB	2 GPU Nvidia K40
Big Mem node	2	Intel Xeon E5 2650 v2 @ 2.6 Ghz	16	512 GB	1 GPU Nvidia K20
BigInsight node	4	Intel Xeon E5 2650 v2 @ 2.6 Ghz	16	64 GB	32TB of local disk

New Cineca infrastructure: GALILEO

- Double socket Intel Xeon 2630 v3@2.4GHz, 8 cores each (“Haswell”)
- 524 nodes for a total number of 8384 cores
- 128 GB of RAM memory per node
- 768 Xeon Phi co-processor model 7120p
- 8 nodes for visualization
- Expected peak performance around 1PF

It replaces PLX and EURORA.

It will start production in February 2015

