KNL Performance Comparison: Incompact3d

www.incompact3d.com

March 2017

1. Compilation, Setup and Input

Compilation

Incompact3d was compiled on both the ARCHER KNL system and ARCHER Xeon system with the Intel Fortran Compiler under the default module PrgEnv-intel (version 15.0.2.164 for the ARCHER Xeon system and version17.0.0.098 for the ARCHER KNL system).

Compilation options are: -cpp -xHost -O3 -ipo -heaparrays -safe-cray-ptr

Setup

The ARCHER KNL nodes were used in "quad_100" configuration with all the MCDRAM configured as an additional cache level.

In all cases, jobs were run on fully-populated nodes on both systems.

Input

The archive of the code used for the tests can be downloaded at *http://www.incompact3d.com/download.html.*

It is the version corresponding to the 3D turbulent channel flow.

Simulations were run five times for 50 time steps and an average was made to get the simulation time per time step.

2. Performance Data

The plot below shows a performance comparison between KNL and Xeon system for two simulations with 1024 x 512 x 512 computational cores and with 512 x 256 x 512 computational cores. For each simulation, we have 1 MPI process per core. Maximum number of cores is 512 (8 nodes with 64 cores for the KNL system, 22 nodes with 24 cores for the Xeon System).

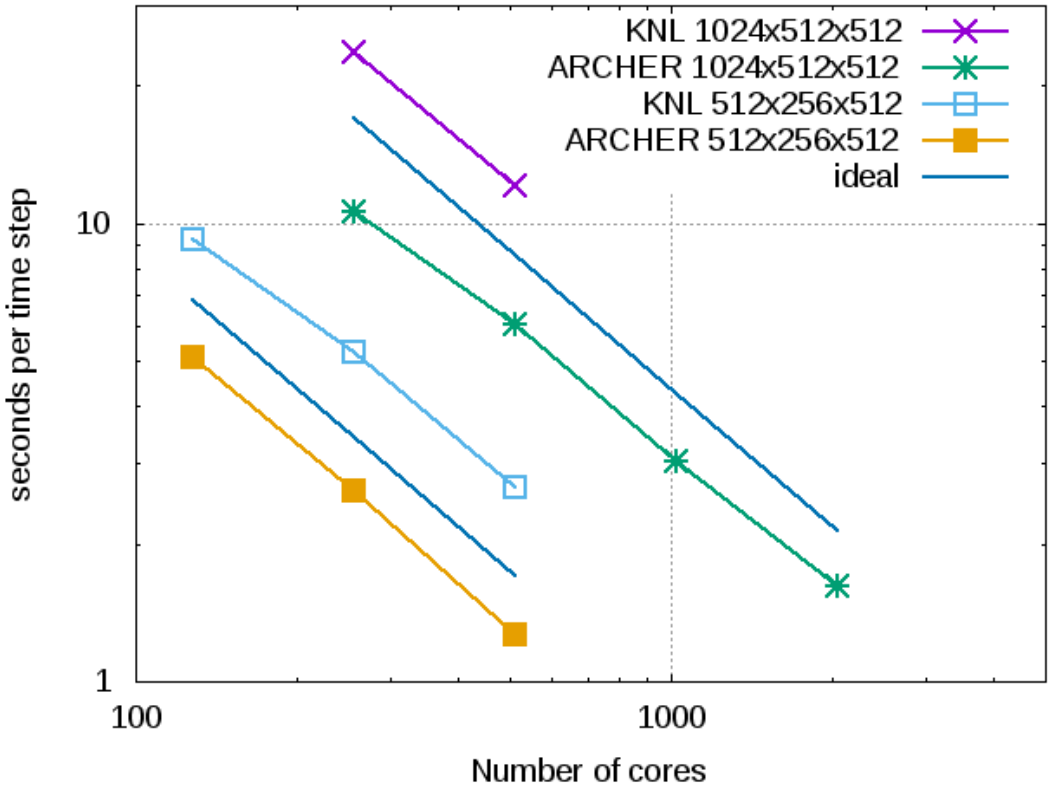

Overall, the KNL system is about twice slower than the Xeon system. Scalability is excellent on both systems.

For the simulation with 1024x512x512 mesh nodes:
KNL 256 cores → 23.80 seconds per time step
KNL 512 cores → 12.12 seconds per time step
ARCHER 256 cores → 10.6 seconds per time step
ARCHER 512 cores → 6.03 seconds per time step
ARCHER 1024 cores → 3.02 seconds per time step
ARCHER 2048 cores → 1.62 seconds per time step

For the simulation with 512x256x512 mesh nodes:
KNL 128 cores → 9.26 seconds per time step
KNL 256 cores → 5.24 seconds per time step
KNL 512 cores → 2.66 seconds per time step
ARCHER 128 cores → 5.09 seconds per time step
ARCHER 256 cores → 2.61 seconds per time step
ARCHER 512 cores → 1.26 seconds per time step

3. Summary and Conclusions

We were expected worse performance for the KNL system. We were very pleased with the results for Incompact3d. Benchmarking is now required on larger KNL systems to understand how the performance varies as the number of nodes is increased. We are expecting excellent performance with hundreds of thousands of KNL cores (Incompact3d can scale with up to one million cores, benchmarks performed on MIRA in the USA in 2016).

We were glad to see that ARCHER is investing in Intel Xeon Phi. At the moment it is not possible to run very large simulations on the Xeon system (with dozens of thousands cores, which is now very common in the US and in China). We are trailing behind France and Germany in Europe and behind the US and China worldwide in terms of number of AUs available for academic research.

As the Xeon system is well oversubscribed and quite small, one option could be to increase the UK resources by purchasing a big KNL system following what was done recently in Italy with the system MARCONI (based on the Lenovo Adam Pass architecture and is equipped with the new Intel Knights Landing BIN1 processors (KNL). It consist of 3600 nodes (1 KNL processor at 1.4GHz and 96 GB of DDR4 ram per node . Each KNL is equipped with 68 cores and 16 GB of MCD RAM –> more than 300,000 cores available). This could be one solution to the problem of the small number of AUs available for academic research in the UK.