KNL Performance Comparison: *Quantum Espresso*

31 March 2017

1. Compilation, Setup and Input

Compilation

The open source and well known ab-initio density-functional code Quantum Espresso (based on plane waves and pseudopotentials: http://www.quantum-espresso.org) was compiled on Archer High-Performance Supercomputing facility, as follows:

	Compiler	*mkl*	*FLAGS*	*BLAS, LAPACK and SCALAPACK*
KNL	Intel 2017.0.098	2017.0.098	-D__INTEL -D__DFTI -D__MPI -D__SCALAPACK	$(MKLROOT)/lib/intel64/libmkl_scalapack_lp64.a -Wl,--start-group $(MKLROOT)/lib/intel64/libmkl_intel_lp64.a $(MKLROOT)/lib/intel64/libmkl_core.a libmkl_sequential.a $(MKLROOT)/lib/intel64/libmkl_blacs_intelmpi_lp64.a -Wl,--end-group
Xeon	Intel Composer 2015.2.164	2015.2.164	-D__INTEL -D__DFTI -D__MPI -D__SCALAPACK	$(MKLROOT)/lib/intel64/libmkl_scalapack_lp64.a -Wl,--start-group $(MKLROOT)/lib/intel64/libmkl_intel_lp64.a $(MKLROOT)/lib/intel64/libmkl_core.a libmkl_sequential.a $(MKLROOT)/lib/intel64/libmkl_blacs_intelmpi_lp64.a -Wl,--end-group

Modules loaded and environment variables for KNL (before executing './configure'):
==

```
Module swap PrgEnv-cray PrgEnv-intel
source $INTEL_PATH/linux/bin/compilervars.sh intel64
source /opt/intel/impi/2017.0.098/bin64/mpivars.sh
export FC=ftn
export F90=$FC
export F77=$FC
export MPIF90=$FC
export CC=cc

export FCFLAGS="-O3 -xAVX -fno-alias -ansi-alias -g -mkl -
$MKLROOT/include/fftw"
```

```
export FFLAGS=$FCFLAGS
==================================================
```

Modules loaded and environment variables for Xeon (before executing './configure'):

```
==================================================
Module swap PrgEnv-cray PrgEnv-intel

export FC=ftn
export F90=$FC
export F77=$FC
export MPIF90=$FC
export CC=cc
==================================================
```

Setup

A maximum of 192 cores was considered for calculation done in the following way: up to 192 cores (up to 8 nodes fully occupied) on Archer-Xeon environment; up to 192 cores (up to 3 nodes fully occupied) on Archer-KNL quad_100 environment; up to 128 cores (up to 2 nodes fully occupied) on Archer-KNL quad_0 environment. MPI only, this means no threading or hyperthreading has been used.

Input

The system used in the following tests consists of 13 Platinum atoms included in a model of zeolite NaY supercage (no periodicity, finite system). We have 526 electrons and 379 Kohn-Sham states in total. Spin-polarised framework was used to take into account the magnetic nature of the Platinum clusters. The van der Walls correction was also included. Calculations done at Γ point only. The input file follows:

```
==================================================
&control
   calculation='scf',
   title='Pt13_D9',
   verbosity='high',
   restart_mode='from_scratch',
   prefix='zeolite_D9_test.pbe'
   pseudo_dir = '. /UPF',
   outdir='. /D9_test',
   wfcdir='./wfc/',
   tprnfor = .true.
   forc_conv_thr=1.0d-4
 /
 &system
   ibrav= 1,
   celldm(1)= 38.d0,
   nat=115,
```

```
   ntyp=5,
   ecutwfc = 45.0,
   ecutrho = 360.0,
   occupations='smearing'
   smearing='cold'
   degauss=0.002
   nspin=2
   starting_magnetization(1) = 0.5
   starting_magnetization(2) = 0.5
   starting_magnetization(3) = 0.5
   starting_magnetization(4) = 0.5
   starting_magnetization(5) = 0.5
   london=.true.
   london_rcut = 150
   london_s6 = 0.75
   nbnd=379
 /
 &electrons
   electron_maxstep = 500
   conv_thr =  1.0d-6
   mixing_mode='local-TF'
   mixing_beta = 0.6
 /
 &ions
   pot_extrapolation='second_order'
 /
ATOMIC_SPECIES
 Pt   1.00   Pt.pbe-nd-rrkjus.UPF
 O    1.00   O.pbe-van_ak.UPF
 Si   1.00   Si.pbe-n-van.UPF
 H    1.00   H.pbe-van_ak.UPF
 Al   1.00   Al.pbe-n-van_ak.UPF
ATOMIC_POSITIONS (angstrom)
Pt      9.876907459   9.986946739   7.987279732
Pt      9.762832623   8.236163144  12.063101441
Pt     12.504445913  10.016245831   8.780421482
Pt     13.184325056   8.027215227  10.477505962
Pt      9.002879012  10.028463996  10.448170846
Pt     11.678781008  10.077585228  11.371568835
Pt     10.857519761  11.786641052   9.680438038
Pt     11.961170803   6.111415472   8.795066273
Pt     11.288901189   6.252190302  11.233566490
Pt      9.995585511   7.772251185   9.411474799
Pt     11.276871541   8.012811274   7.136319466
Pt     12.293149439   8.113844914  12.875200734
Pt     13.803762956   7.928555437   8.051258894
O      14.713096220   7.950191680   5.729661550    0   0   0
Si     15.947442410   8.980816480   5.806962750    0   0   0
Si     14.405779710  11.616943480   5.818984550    0   0   0
```

O	12.902034020	11.046971480	5.743784150	0	0	0
O	16.532427710	9.010966880	7.307729640	0	0	0
O	14.724842350	12.101801580	7.321825120	0	0	0
Si	17.485628990	9.869090280	8.282639730	0	0	0
Si	15.943966280	12.505217280	8.294661550	0	0	0
O	17.247192840	9.413584880	9.808233540	0	0	0
O	15.436130650	12.510364680	9.822356120	0	0	0
Al	17.523570720	9.861431380	11.329687780	0	0	0
Si	15.968004210	12.521332780	11.341818020	0	0	0
O	16.528388120	11.062458080	11.726833370	0	0	0
O	6.106245030	7.902006380	5.797812750	0	0	0
O	7.882519000	11.018869780	5.783529950	0	0	0
Si	6.373865560	11.571976980	5.882583250	0	0	0
O	5.347790510	10.397400180	5.481760250	0	0	0
Si	4.861816260	8.918753780	5.894741550	0	0	0
O	4.300368080	8.942485980	7.404586170	0	0	0
O	6.073232000	12.053365780	7.390330750	0	0	0
Si	4.865225620	12.443193180	8.382385810	0	0	0
O	3.684025960	11.358198480	8.238602830	0	0	0
Si	3.353176310	9.789970080	8.394544050	0	0	0
O	5.397099180	12.454161380	9.901847700	0	0	0
O	3.620825220	9.337297880	9.916130550	0	0	0
Si	4.889263550	12.459308680	11.429542280	0	0	0
O	4.351415180	10.994285580	11.823253710	0	0	0
Al	3.363577520	9.782156980	11.441810180	0	0	0
O	3.641720170	8.534412380	12.422211790	0	0	0
O	6.108388180	12.862723980	12.402377550	0	0	0
Si	6.427449410	13.347582180	13.905218110	0	0	0
Si	3.393123440	8.023190580	13.929617400	0	0	0
O	4.388306750	8.798060680	14.929500620	0	0	0
O	6.271976450	12.103373580	14.914354210	0	0	0
Al	6.457329890	11.568740880	16.421135568	0	0	0
O	6.156697730	9.986764680	16.470012058	0	0	0
Si	4.948689940	8.921500980	16.433266971	0	0	0
O	7.931195090	13.917554180	13.980418520	0	0	0
O	7.965983320	11.846545180	16.909565554	0	0	0
Si	8.949740870	13.021880980	17.402089811	0	0	0
O	8.778290510	14.300040680	16.437707955	0	0	0
Si	8.920127470	14.784822880	14.908659850	0	0	0
O	10.440986712	14.575662080	14.420134490	0	0	0
O	10.475708154	12.508628980	17.343658209	0	0	0
Si	11.996567400	13.038938580	17.377964271	0	0	0
O	12.138404350	14.318852180	16.411101727	0	0	0
Si	11.966954000	14.801880480	14.884534310	0	0	0
O	12.950711530	13.945655880	13.940672720	0	0	0
O	12.985499760	11.874646780	16.869819758	0	0	0
Si	14.489245440	11.613707380	16.357536837	0	0	0
O	14.644718420	12.150248280	14.848056730	0	0	0
Al	14.459364970	13.392548680	13.841619380	0	0	0

O	14.808306680	10.035200580	16.401506439	0	0	0
O	14.759997120	12.911159880	12.333871930	0	0	0
Si	16.027432020	8.983525080	16.345542700	0	0	0
O	16.565279690	8.866233280	14.833080280	0	0	0
Si	17.553118040	8.102464980	13.817495000	0	0	0
O	17.274974690	8.610737880	12.314260240	0	0	0
O	15.438465210	10.453892580	5.401859750	0	0	0
O	17.134814500	11.433502480	8.132096100	0	0	0
O	10.423727846	5.562231742	5.004919204	0	0	0
O	7.925575110	6.060552280	5.323903550	0	0	0
O	8.704934379	8.481114673	4.861413515	0	0	0
O	10.392243118	10.388863980	5.304069350	0	0	0
O	12.122802869	8.511135824	5.136245336	0	0	0
O	12.907220690	6.088441880	5.284457650	0	0	0
Si	11.960028900	4.872560980	4.815542850	0	0	0
Si	8.878776370	4.855310680	4.839940950	0	0	0
Si	7.340590510	7.485493380	4.851936250	0	0	0
Si	8.866276530	10.162645180	4.839668350	0	0	0
Al	11.913101650	10.179702680	4.815542850	0	0	0
Al	13.468668150	7.519801280	4.803412550	0	0	0
H	16.737762000	8.658714000	5.112958000	0	0	0
H	14.509795000	12.468376000	5.130331500	0	0	0
H	18.537254000	9.697889000	8.009198000	0	0	0
H	16.309290000	13.507547000	8.026574000	0	0	0
H	16.775543000	13.261793000	11.439804000	0	0	0
H	6.249430700	12.422130000	5.195739000	0	0	0
H	4.064258000	8.587766000	5.213311000	0	0	0
H	4.484503300	13.441347000	8.120206000	0	0	0
H	2.299332900	9.606982000	8.137774000	0	0	0
H	4.075137600	13.190689000	11.540367000	0	0	0
H	5.722090700	14.141773000	14.191100000	0	0	0
H	3.502409000	6.931745000	14.012087000	0	0	0
H	2.341637100	8.140566000	14.230615000	0	0	0
H	5.255875600	7.944338000	16.834251000	0	0	0
H	4.151672400	9.185825000	17.143826000	0	0	0
H	8.697535500	13.301827000	18.435543000	0	0	0
H	8.654738000	15.849580000	14.832106000	0	0	0
H	12.261952000	13.321781000	18.407316000	0	0	0
H	12.219157000	15.869536000	14.803881000	0	0	0
H	15.194603000	12.148034000	17.010960000	0	0	0
H	15.737596000	8.003019000	16.751253000	0	0	0
H	16.832594000	9.256818000	17.043413000	0	0	0
H	17.457375000	7.009870500	13.901587500	0	0	0
H	18.607857000	8.231628000	14.101816000	0	0	0
H	12.166068000	4.512358000	3.796818300	0	0	0
H	12.107302000	3.944635600	5.387614000	0	0	0
H	8.659757000	4.493230000	3.824594500	0	0	0
H	8.749569000	3.925992700	5.414113000	0	0	0
H	6.884428000	7.526468300	3.851818000	0	0	0

```
H       8.594681000  10.527301000   3.838039000    0   0   0
K_POINTS gamma
================================================
```

2. Performance Data

Figure 1 and Figure 2 show the performance of Quantum Espresso on Xeon and KNL environments (the latter in both QUAD_100 'cache mode' where 100% of 16GB of on-chip MCDRAM is used to cache accesses to system DDR memory and QUAD_0 'flat mode' where the cache memory is fully used as additional system memory) for different total number of cores used to achieve a full convergence of a ab-initio Self-Consistent electronic cycle (SCF). The relative times (CPU and WALL times respectively) calculated with respect to the fastest time in the overall set of calculation are reported on the y-axis.

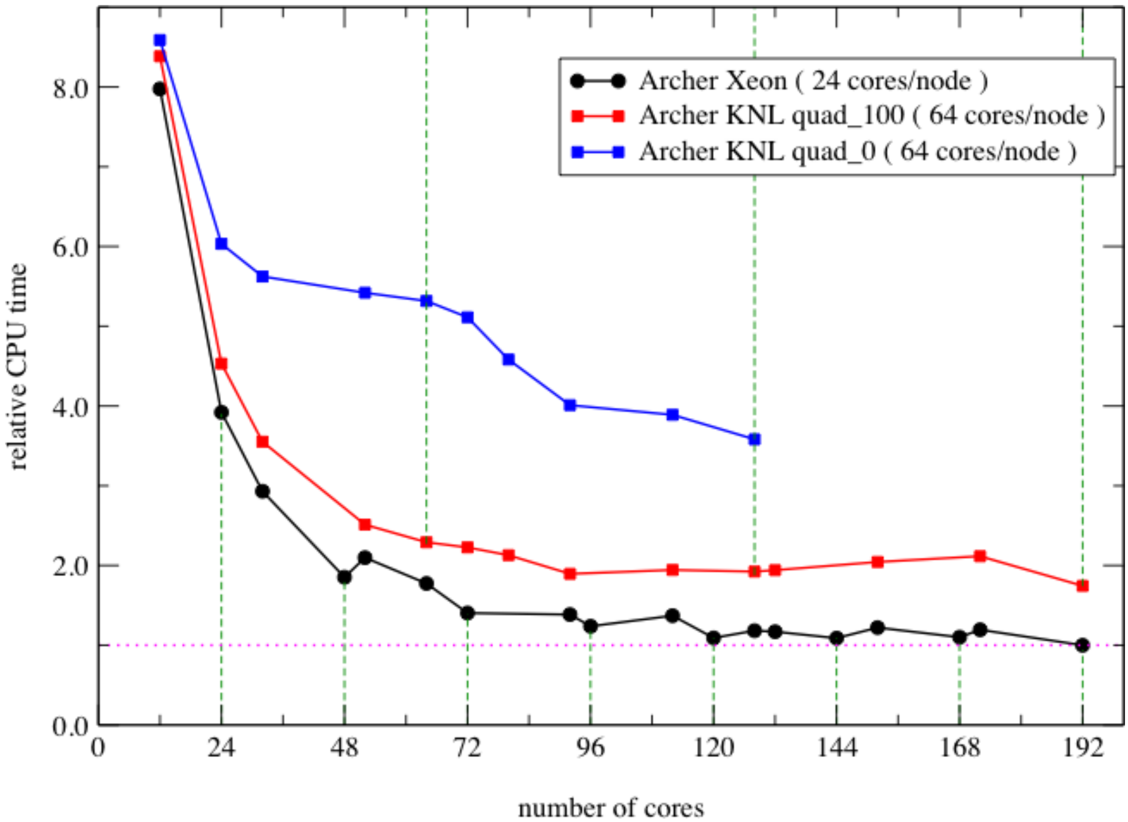

Figure 1: relative CPU time (time actually spent by CPU executing the calculation) was calculated as: t_{ncores}/t_{192_xeon}, where t_{ncores} is the time spent by the calculation to converge and complete a Self-Consistent electronic cycle (SCF) as CPU time and t_{192_xeon}=586.80 s is the fastest CPU time, achieved on Xeon environment with 192 cores. Dashed vertical green lines represent the value of the relative CPU time for calculations on full nodes (multiple of 24 cores on Xeon or multiple of 64 on KNL). Horizontal magenta dotted-line represents the relative CPU time when $t_{ncores}= t_{192\_xeon}$

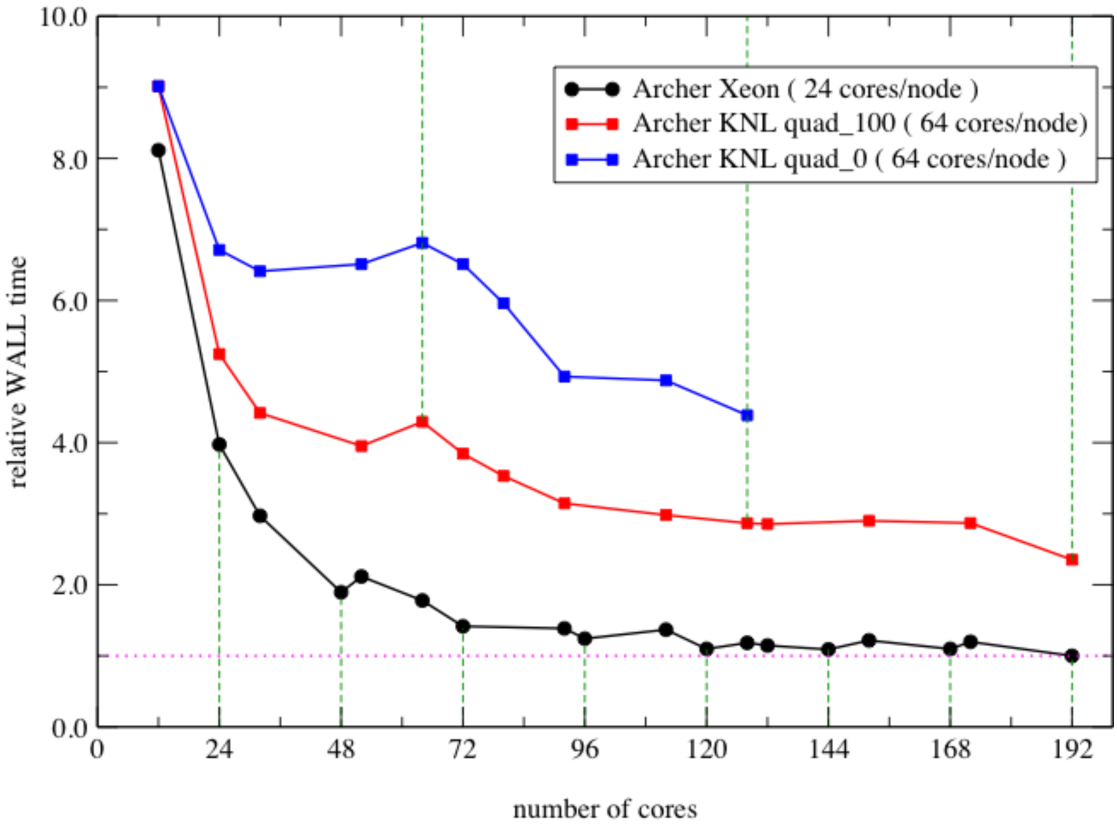

Figure 2: relative WALL time (total time actually spent by CPU executing the calculation and by the system as I/O time and communication channel delay) was calculated as: t_{ncores}/t_{192_xeon}, where t_{ncores} is the time spent by the calculation to converge and complete a Self-Consistent electronic cycle (SCF) as WALL time and t_{192_xeon}=598.96 s is the fastest WALL time achieved on Xeon environment with 192 cores. Dashed vertical green lines represent the value of the relative WALL time for calculations on full nodes (multiple of 24 cores on Xeon or multiple of 64 on KNL). Horizontal magenta dot-line represents the relative WALL time when t_{ncores}= t_{192_xeon}

3. Summary and Conclusions

In the above tests where calculations were executed using MPI only parallelisation on different number of cores up to 192 cores, the speed performances reveal that:

1) Calculations executed on Archer Xeon 24 cores/node are faster than any combination of cores and available cache memory arrangement on KNL 64 cores/node. In particular, the former architecture is two times more foster than the latter in QUAD_100 and four times faster with respect to the QUAD_0 arrangements, respectively.

2) On Xeon architecture the best scaling can be observed up to 2 nodes (48 cores), while flat behavior can be observed after 120 cores. Fully occupied nodes perform better than partial occupied ones.

3) KNL in QUAD_100 'cache mode' shows better scaling up to 50/64 cores (one full node), after that a quasi-flat behavior is observed.

4) KNL in QUAD_0 'flat mode' shows that there is a jump in scaling when the calculation is executed in one or two nodes (64 or 128 cores respectively). Interestingly, the performance of KNL in 'flat mode' is the same when the node is occupied for half of its load or bigger.