Performance Comparing of VASP on CLAIX-2018 and SX-Aurora TSUBASA

Runtime, Energy Consumption, Performance Measurements, Analysis results of VASP on AURORA and on CLAIX18

High Performance Computing | IT Center | RWTH AACHEN UNIVERSITY

Motivation

- **VASP is an important code for users of RWTH Compute Cluster**
- **SX-Aurora TSUBASA is the architecture of interest for us**
  - → Improving, reliability and productivity of High-Performance Computing on the NEC CLAIX system
- **Evaluation of different architectures with representative reduced Benchmark**
  - → Scalability
  - → Power Efficiency
- **Discussing of metrics to make performance comparing of the compute systems more fair and aware**

Compute Systems

- **RWTH Compute Cluster CLAIX-2018**
  - Intel Xeon Platinum 8160 (SkyLake)
  - 2 sockets per node
  - 48 cores per node
  - 2.1GHz
  - Peak performance ~2.24TF per node
  - Intel compiler 19.0, Intelmpi 2018
  - Intel MKL

- **NEC CLAIX**
  - Vector Host
    - Intel Xeon Silver 4108 (SkyLake)
    - 1.80 GHz
  - 8 Vector Engines (cards) Type10
    - 8 cores
    - 2.45TF
    - 1.22TB/s memory bandwidth
  - NEC compiler 3.0.7, NEC MPI 2.7.0, NLC 2.0.0
  - FTRACE

VASP

The Vienna Ab initio Simulation Package(VASP)

A copyright-protected software for atomic scale materials modelling, e.g. electronic structure calculations and quantum-mechanical molecular dynamics, from first principles. The basic methodology is density functional theory (DFT). (reference: https://www.vasp.at/)

VASP Version

- On RWTH Compute Cluster CLAIX-2018
  - Self-built Version 5.4.4
- On SX-Aurora TSUBASA
  - Version 5.4.4, patch from NEC

VASP Benchmarks

- **Big case**
  - Representative but reduced data set from VASP users on CLAIX-2018
  - Scalable case
  - Start data for 488 ions
  - High termination criteria
  - LREAL = Auto
  - VE10: NCORE = 8
  - Xeon: NCORE = 12/48/96
- **Small case**
  - Running test data set for small number of processes
  - Start data for 15 ions
  - Earlier termination of calculation
  - LREAL=.FALSE.
  - VE10: NCORE = 4
  - Xeon: NCORE = 24

VASP Vectorization

FTRACE Analysis results on four Aurora cards

```
Execution Date : Wed Nov  4 11:58:40 2020 CET
Total CPU Time : 5:20'26"309 (19226.309 sec.)
Elapsed Time   : 0:11'37"379

PMMR           : 0x0000000000001000
PMCR0          : 0x0000000000000000


FREQUENCY  EXCLUSIVE       AVER.TIME     MOPS   MFLOPS  V.OP  AVER.    VECTOR L1CACHE CPU PORT VLD LLC PROC.NAME
           TIME[sec](  % )    [msec]                    RATIO V.LEN      TIME    MISS     CONF HIT E.%

   182824  3327.135( 17.3)    18.199   3588.3   2189.0  78.33 198.8   271.374 2099.889    1.223   44.84 HAMIL::HAMILTMU
  1903840  2834.635( 14.7)     1.489  26123.3  19940.6  99.56 241.9  2365.717 379.648  626.206   67.73 FFTBAS_PLAN_MPI
       32  2067.611( 10.8) 64612.831   1315.4     72.5   2.75 255.7     2.897   0.225    0.000   79.12 WAVE::WFINIT
   339008  1454.109(  7.6)     4.289  94179.2  89170.1  99.71 190.6  1453.301   0.736  205.861   56.50 ORTH1
   153544  1156.775(  6.0)     7.534   6769.3   4642.2  90.36  16.0  1128.116  17.265   48.709   48.92 NONL_HIGH::W1_PROJALL
      160  1140.641(  5.9)  7129.009  59648.1  56536.8  99.22 233.5   703.844  14.637    1.320   15.22 DAVID::EDDAV
      384  1002.734(  5.2)  2611.287  18687.9  16397.5  95.91 204.4   141.121  13.610    2.440   38.81 SUBROT::EDDIAG
  2466304   793.562(  4.1)     0.322   6691.6      0.0  98.33 248.4   615.330  48.490   21.043   22.87 MAP_FORWARD
  3063936   571.125(  3.0)     0.186    706.8      0.0  89.70 106.4   167.667  18.289    0.000   55.13 M_ALLTOALL_D
  1341248   548.229(  2.9)     0.409   2879.8      0.0  96.72 201.8   396.774  24.519   12.768   27.16 MAP_BACKWARD
-------------------------------------------------------------------------------------------------------------------
672898380 19226.309(100.0)     0.029  24502.6  21054.5  97.61 170.3  9698.210 2840.884  939.849   42.48 total


  ELAPSED     COMM.TIME  COMM.TIME   IDLE TIME  IDLE TIME  AVER.LEN      COUNT  TOTAL LEN PROC.NAME
  TIME[sec]       [sec]  / ELAPSED       [sec]  / ELAPSED    [byte]                 [byte]

    108.954       0.000                  0.000                 0.0          0       0.0  HAMIL::HAMILTMU
     89.651       0.000                  0.000                 0.0          0       0.0  FFTBAS_PLAN_MPI
     79.602       0.000                  0.000                 0.0          0       0.0  WAVE::WFINIT
     62.367       0.000                  0.000                 0.0          0       0.0  ORTH1
     43.000      22.054                 21.133                47.8K     832384      38.0G DAVID::EDDAV
     37.572       0.000                  0.000                 0.0          0       0.0  NONL_HIGH::W1_PROJALL
     32.152      32.137                 18.028                 4.1M      58400     232.6G DIS_PW_BAND
     31.421       0.585                  0.340               125.0K     204672      24.4G SUBROT::EDDIAG
     27.281      10.886                  4.875                 3.1M    1233600       3.6T MAP_FORWARD
     23.149      15.066                  9.198                 3.1M     671040       2.0T MAP_BACKWARD
-------------------------------------------------------------------------------------------------------------
```

VASP Scalability

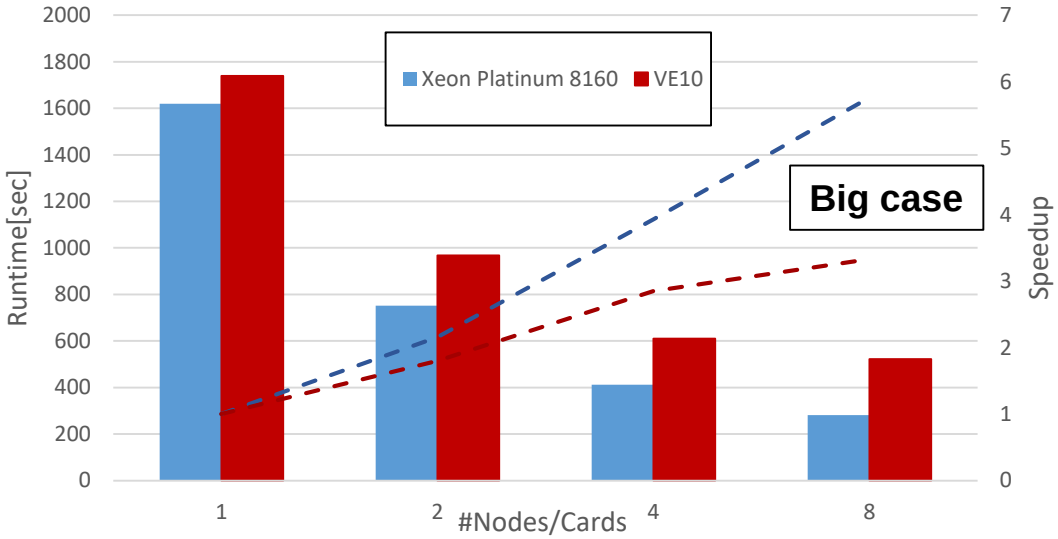

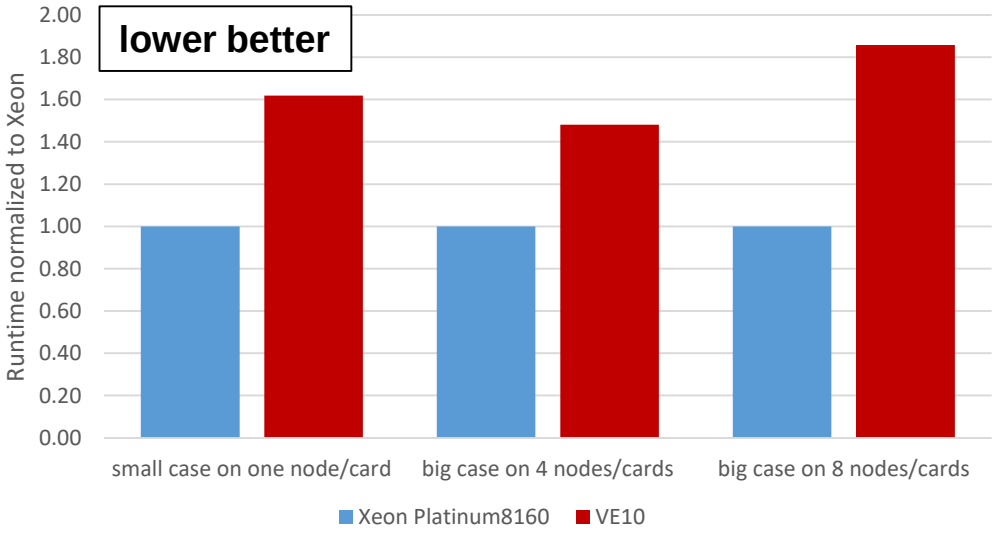

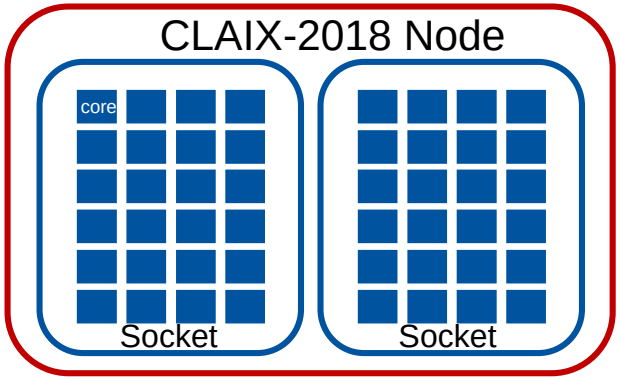

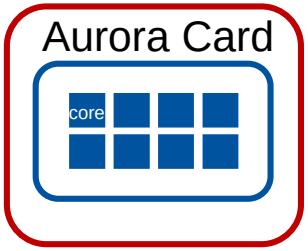

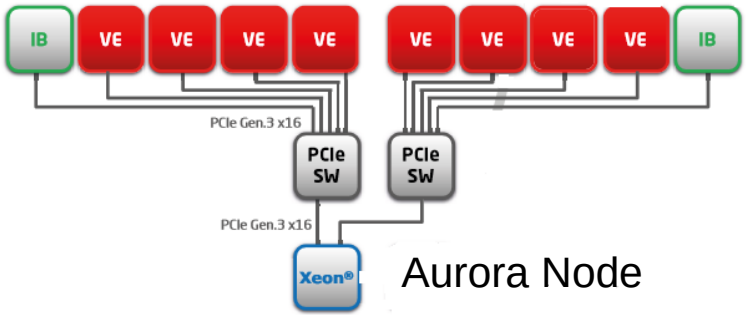

Comparing of Energy Consumption of VASP

- **VASP needs some more time on Aurora**
- **VASP can be more efficient on Aurora in consumption of energy**

- **Energy consumption measurements**
  - → No using of energy measuring devices
  - → Only using of performance monitoring tools

- **CLAIX-2018**
  - → Likwid measurements on one compute node with following groups
    - → ENERGY for energy consumption
    - → FLOPS_DP for compute performance
- **Aurora**
  - → veperf for energy consumption and compute performance on VEs
  - → Likwid for energy consumption of VH

VASP Energy Consumption and Power Efficiency

$$\boldsymbol{Power} = \frac{Energy\ per\ node/card}{Runtime}$$

$$\boldsymbol{Efficiency} = \frac{Number\ MFLOPS}{Power}$$

SX-Aurora TSUBASA

$$Energy\ per\ card = VE\ Energy + \frac{VH\ Energy}{8}$$

$$VE\ Energy\ = Energy\ of\ one\ VE$$
(from **veperf)**

$$VH\ Energy = SUM(Energy\ STAT + DRAM)$$
(from **Likwid)**

$$MFLOPS = Avg(MFLOPS\ over\ all\ cores) * 8$$
(from **veperf)**

Intel Xeon Platinum 8160

$$Energy\ per\ node = SUM(Energy\ STAT + DRAM)$$
(from **Likwid**)

$$MFLOPS = SUM(MFLOPS\ over\ cores\ of\ one\ node)$$
(from **Likwid**)

VASP Power Efficiency

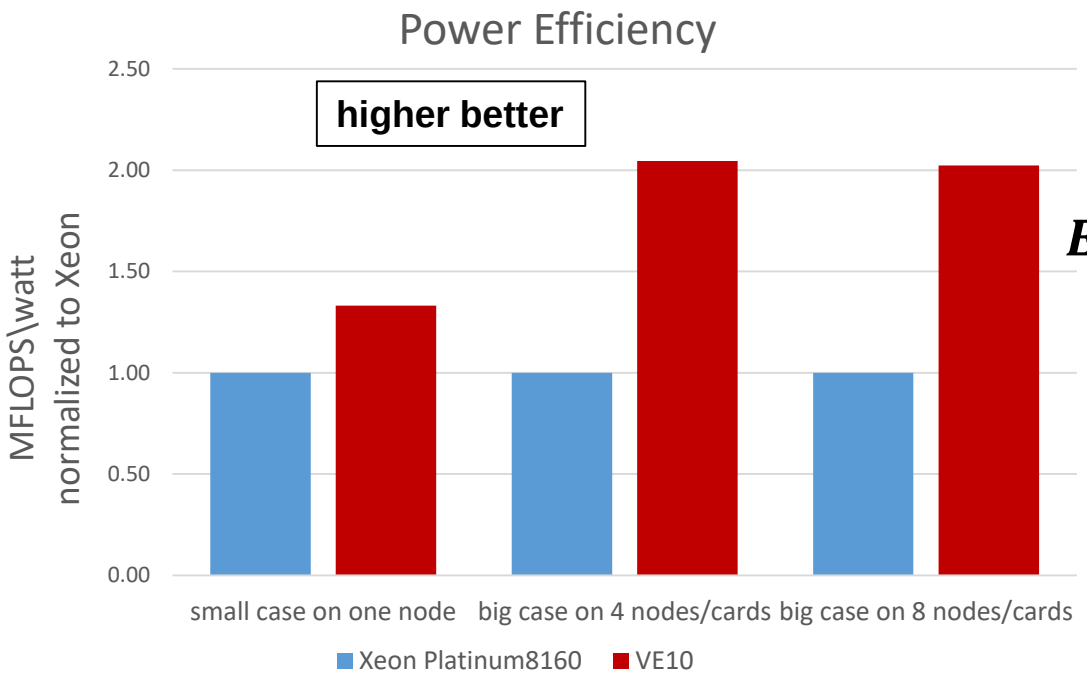

$$Efficiency = \frac{Number\ MFLOPS}{Power}$$

$$Power = \frac{Energy\ per\ node/card}{Runtime}$$

Aurora

- Power 95-120 Watt per card incl. VH (Xeon is ~340 Watt per node)
- Power efficiency (MFLOPS/Watt per card) is better (1.3-2.05x)
- Energy consumption measurements include energy for VE and VH (CPU and DRAM)

Energy Consumption Measurements on Aurora

#Cards	Runtime Speedup	Energy Cards [kJ]	Energy Host [kJ]	Energy Cards + Host part [kJ]	Power per Card excl. Host [W]	Power per Card incl. Host [W]	Power total [W]
1	1.00	195	128	211	112	121	121
2	1.80	208	73	227	108	117	234
4	2.85	240	52	266	98	109	436
8	3.25	364	43	407	85	95	761

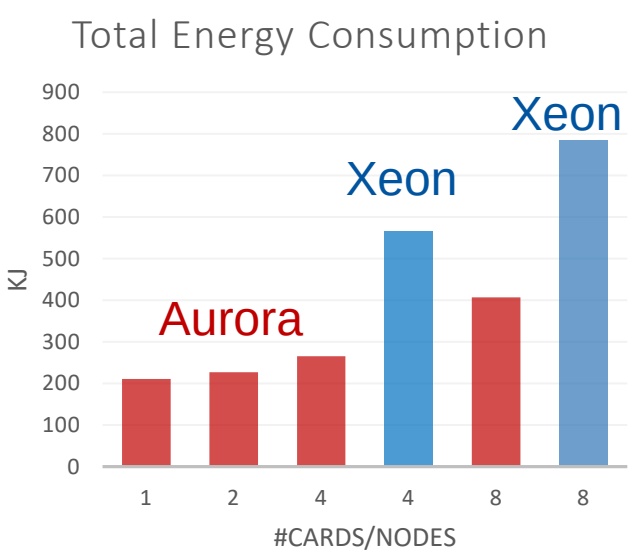

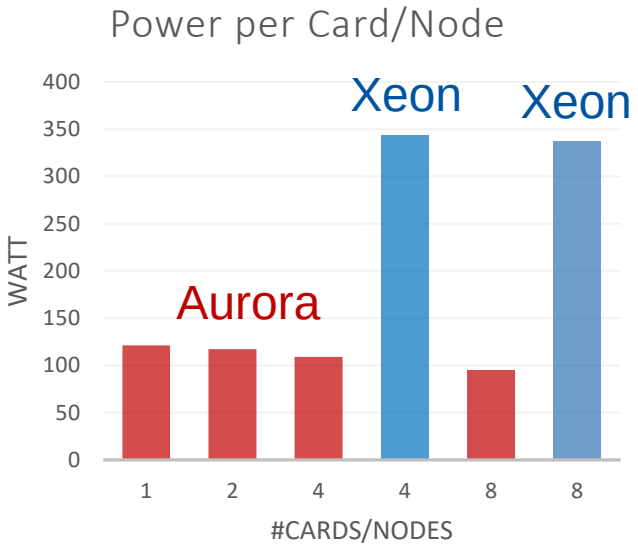

Conclusion

- **VASP on Aurora**
  - → Not much more time for solution
  - → Very high vector operation ratio and average vector length
  - → Much lower energy consumption
  - → Higher Power Efficiency [MFLOPS/WATT]

Thank you for your attention