

Measuring Performance of Intel Broadwell-EP Processors with WRF and NAMD Applications

Authors: Ashish K Singh, Mayura Deshmukh, Neha Kashyap

This blog describes the performance of Intel Broadwell processors with HPC applications, Weather Research and Forecasting (WRF) and NANoscale Molecular Dynamics (NAMD). This is the second blog in the series of four blogs on “Performance study of Intel Broadwell”. [The first blog characterizes the Broadwell-EP processors with HPC benchmarks like HPL and STREAM](#). This study compares five different Broadwell processors E5-2699 v4 @2.2GHz (22 cores), E5-2698 v4 @2.2GHz (20 cores), E5-2697A v4 @2.6GHz (16 cores), E5-2690 v4 @2.6GHz (14 cores), and E5-2650 v4 @2.2GHz (12 cores) in PowerEdge 13th generation servers. It characterizes the performance of the system for WRF and NAMD by comparing five Broadwell processors models with previous generations of Intel processors. For the generation over generation comparison, previous results from Intel Xeon X5600 series Westmere (WSM), Intel Xeon E5-2600 series Sandy-Bridge (SB), Intel Xeon E5-2600 v2 series Ivy-Bridge (IVY), Intel Xeon E5-2600 v3 series Haswell (HSW) and Intel Xeon E5-2600 v4 series Broadwell (BDW) processors were used. This blog also describes the impact of BIOS tuning options on WRF and NAMD performance with Broadwell. Table 1 below lists the server configuration and the application details for the Broadwell processor based tests. The software versions were different for the older generation processors and results are compared against what was best configuration at that time. Due to big architectural changes in servers and processors generation over generation, the changes in software versions is not a significant factor.

Table 1: Details of Server and HPC Applications used with Intel Broadwell processors

Server	Dell PowerEdge R730
Processors	E5-2699 v4 @2.2GHz, 22 cores, 145W E5-2698 v4 @2.2GHz, 20 cores, 135W E5-2697A v4 @2.6GHz, 16 cores, 145W E5-2690 v4 @2.6GHz, 14 cores, 135W E5-2650 v4 @2.2GHz, 12 cores, 105W
Memory	16 x 16GB DDR4 @ 2400MHz (Total=256GB)

Power Supply	2 x 1100W
Operating System	Red Hat Enterprise Linux 7.2 (3.10.0-327.el7.x86_64)
BIOS options	System profile – Performance and Performance Per Watt (DAPC) Logical Processor – Disabled Power Supply Redundant Policy – Not Redundant Power Supply Hot Spare Policy – Disabled I/O Non-Posted Prefetch - Disabled Snoop modes – COD, ES, HS and OSB Node Interleaving - Disabled
BIOS Firmware	2.0.0
iDRAC Firmware	2.30.30.02
Intel Compiler	From Intel Parallel studio 2016 update1
MPI	Intel MPI – 5.1.2
WRF	3.6.1
NetCDF	4.4.0
NetCDF-Fortran	4.4.2
NAMD	2.11
FFTW	2.1.5

Weather Research and Forecasting ([WRF](#)) is an HPC application used for atmospheric research. The WRF model is a next-generation mesoscale numerical weather prediction system designed for both atmospheric research and operational forecasting needs. This serves a wide range of meteorological applications across scales from tens of meters to thousands of kilometers. WRF can generate atmospheric simulations using real data or idealized conditions. We used the CONUS12km and CONUS2.5km benchmarks for this study.

[CONUS12km is a single domain and small size \(48hours, 12km resolution case over the Continental U.S. \(CONUS\) domain from October 24, 2001\) benchmark with 72 seconds of time step. CONUS2.5km is a single domain and large size \(Latter 3hours of a 9hours, 2.5km resolution case over the Continental U.S. \(CONUS\) domain from June 4, 2005\) benchmark with 15 seconds of time step.](#)

WRF decomposes the domain into tasks or patches. Each patch can be further decomposed into tiles that are processed separately, but by default there is only one tile for every run. If the single tile is too large to fit into the cache of the CPU and/or core, it slows down computation due to WRF's memory bandwidth sensitivity. In order to reduce the size of the tile, it is possible to increase the number of tiles by defining "numtile = x" in input file or defining environment variable "WRF_NUM_TILES = x". For both CONUS 12km and CONUS 2.5km the number of tiles are chosen based on best performance. The tiles value depend on workload and hardware configuration. Table 2 shows more detail on the number of tiles used in this study for best performance.

Table2: Parameters used in WRF for best performance

	Total no. of cores	CONUS12km			CONUS2.5km		
		MPI Processes	OMP Processes	TILES	MPI Processes	OMP Processes	TILES
E5-2699 v4	44	22	2	22	44	1	56
E5-2698 v4	40	20	2	20	40	1	80
E5-2697A v4	32	16	2	16	32	1	64
E5-2690 v4	28	14	2	14	28	1	56
E5-2650 v4	24	12	2	12	24	1	56

[NAMD](#) is one of the HPC applications used in molecular dynamics research. It is a portable, parallel and object oriented molecular dynamics code designed for high-performance simulations of large bio molecular systems. NAMD is developed using charm++. Molecular Dynamics simulations of bio molecular systems are an important technique for our understanding of biological systems. This study has been performed with three NAMD benchmarks ApoA1 (92,224 Atoms), F1ATPase (327,506 Atoms) and STMV (virus, 1,066,628 Atoms). In the context of number of atoms, these benchmarks lie in the category of small, medium and large size datasets.

Intel Broadwell processors

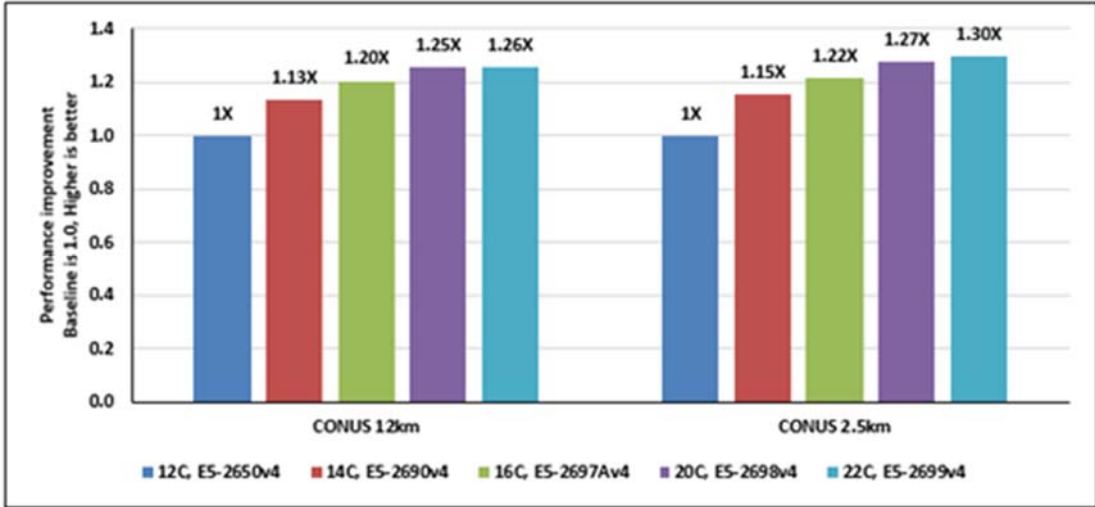


Figure 1: Performance for Intel Broadwell processors with WRF

Figure 1 compares performance among five Broadwell processors by using small and large size of WRF benchmarks. WRF was compiled with the “sm + dm” mode. The combinations of MPI and OpenMP processes that were used are mentioned in Table2.

The “X” value in the graph on top of each bar show the performance relative to the 12 core Broadwell processor (which is set as baseline, 1.0). For the small size dataset CONUS12km, the top bin processor performs 26% better than 12 core processor. While for CONUS2.5km, performance increases up to 30% due to the large dataset size, which can more efficiently utilize larger number of processors. The performance increase from 20 to 22 cores is not as significant due to the lower memory bandwidth per core as explained in the [first blog’s STREAM results](#).




Figure 2: Performance of Intel Broadwell processors with NAMD

Figure 2 plots the simulation speed of NAMD benchmarks with Broadwell processors. The “X” value in the graph on top of each bar show the performance relative to the 12 core Broadwell processor (which is set as baseline, 1.0). As seen from the graph, the relative performance of the different processors models is nearly same irrespective of the NAMD benchmark dataset (small, medium or large). For the top bin processor, the performance improvement is 81 to 84% faster than the 12 core processor. NAMD benchmarks show significant performance improvement with additional cores for Broadwell processors.

BIOS Profiles comparison

This study was performed with all snoop modes: Home Snoop (HS), Early Snoop (ES), Cluster-on-Die (COD) and Opportunistic Snoop Broadcast (OSB) with System Profiles “Performance” and “Performance Per-Watt (DAPC)”. More details on these BIOS profiles are in [first blog of this series](#).




Figure 3: BIOS profile comparison with WRF for CONUS 2.5 km

Figure 3 compares the available snoop modes and two BIOS profiles for the large WRF dataset “CONUS2.5km”. The left graph compares snoop modes with the default BIOS setting “OSB” snoop mode with “DAPC” System Profile, which is shown as the red line set at 1. As per the graph, the COD snoop mode performs 2 to 3% better than default OSB snoop mode. As WRF is a memory sensitive application, the ES snoop mode performance is less than the other snoop modes, up to 8% lower at 22 cores, due to having less request tokens per core compared to other snoop modes (e.g. for 14 core 128/14 = 9 per core vs. 128/22 = 5 per core for 22 core). The right graph compares “Performance” with the “DAPC” system profile for the better performing “COD” snoop mode with COD.DAPC as the baseline. There is not a significant performance difference with “Performance,” only up to 1% better than “DAPC”.




Figure 4: BIOS profile comparison with ApoA1 (92,224 Atoms)




Figure 5: BIOS profile comparison with ATPase (327,506 Atoms)

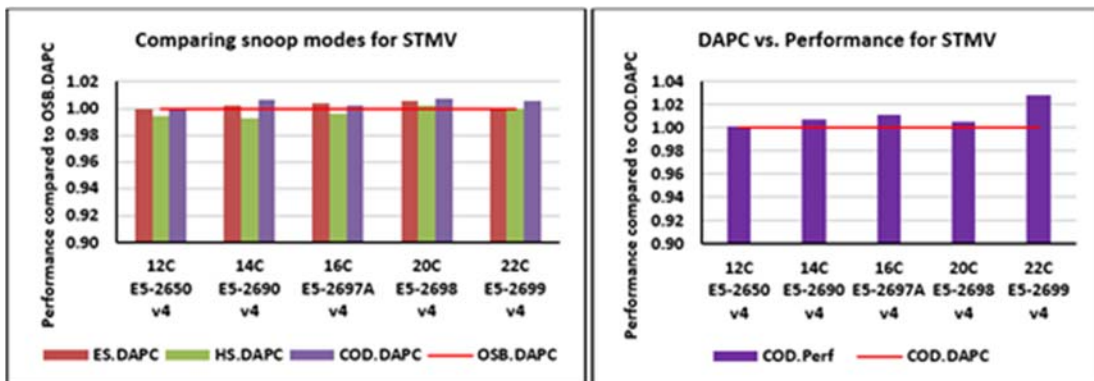


Figure 6: BIOS profile comparison with STMV (1,066,628 Atoms)

Figure 4, 5 and 6 show the performance characteristics of snoop modes available for Broadwell processors with three (small, medium and large) NAMD benchmarks. The left graphs compare snoop modes with the default BIOS Profile (OSB snoop mode with DAPC system profile, which is shown as the red line set at 1). The performance of all NAMD benchmarks across all snoop modes are almost the same for all the datasets. COD is about 1% better for some of the processors for all the data sets but it is not significantly different compared to the performance of the other snoop modes. The right graph compares “Performance” system profile with the default “DAPC” system profile, which is the baseline with COD snoop mode (red line set at 1 in the graph). It can be seen from the graph, NAMD performed up to 3% better with “Performance” profile and COD snoop mode compared to “DAPC” system profile. As seen from these three graphs, the performance with COD snoop mode and “Performance” system profile improves more with the larger datasets specifically for the 22core part.

Generation over Generation comparison

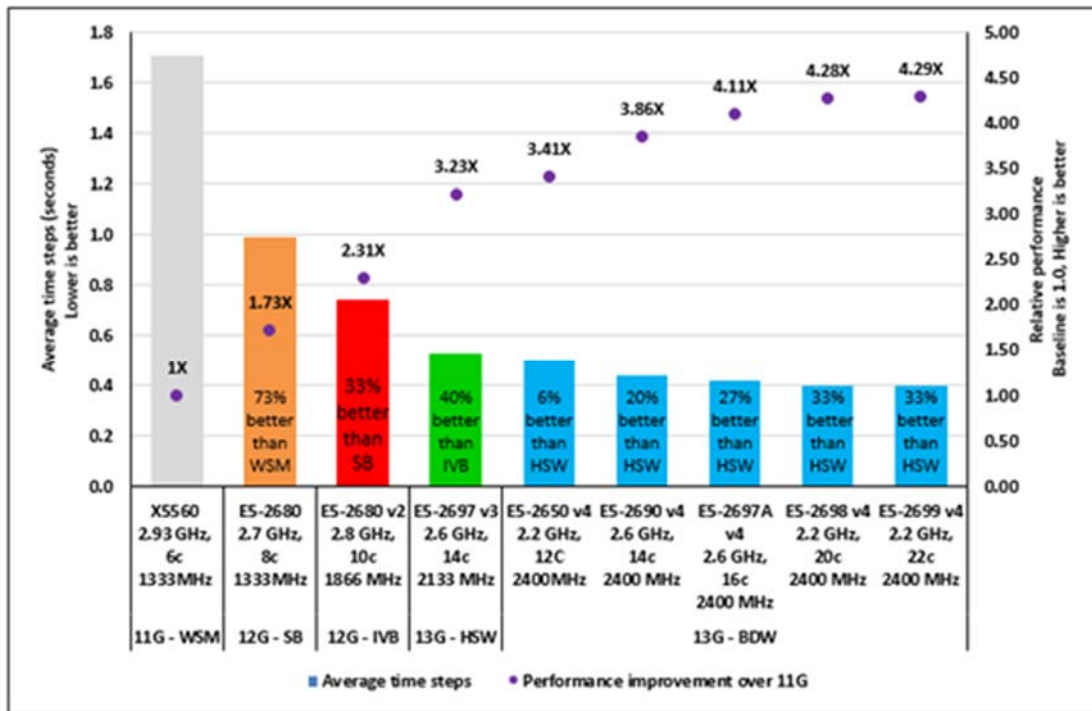


Figure 7: Generation over generation comparison of Intel processors for CONUS12km WRF benchmark

Figure 7 plots the performance characteristics of the CONUS 12km WRF benchmark over multiple generations of Intel processors. Bars in the graph show the average time step result of the CONUS12km benchmark in seconds and purple

dots show the performance relative to WSM processor. It can be easily seen from the graph, the performance of the 14 core Broadwell processor is 20% better than the 14 core HSW processor. The performance of all the Broadwell processors is better than the Haswell E5-2697 v3. The performance improves up to 33% for top bin processor relative to Haswell E5-2697 v3. The performance for the 20 and 22 core Broadwell processors is the same and that is likely because of the lower memory bandwidth per core.




Figure 8: Comparing two generations of Intel processors with WRF Conus 2.5

Figure 8 shows the performance comparison among two generations of Intel processors: Haswell and Broadwell. In this graph, the bar shows the average time step value and the purple dots show the performance improvement relative to the 12 core Haswell processor. The 12 core Broadwell processor has 13% higher memory frequency than the 12 core Haswell processor (2400 MT/s vs. 2133 MT/s in Haswell), but it also has 17% lower AVX base frequency. Due to these performance parameters, the 12 core Broadwell processor performs 6% lower than the 12 core Haswell processor. As per the graph, the top bin 22 core Broadwell processor performs 14% better than the 14 core Haswell processor. Similar to what we saw earlier, there is not a significant performance improvement from the 20 core to 22 core Broadwell processors due to lower memory bandwidth per core.




Figure 9: Performance comparison of multiple generations of Intel processors with ApoA1 benchmark

Figure 9 shows the comparison of multiple generations of Intel processors (92,224 Atoms) NAMD benchmark named ApoA1. The bars show the NAMD performance of the processors in “days/ns”. As seen from the graph, HSW performs 40% better than IVB. While, BDW’s performance improvement varies from 23 to 52% except for the 12 core BDW. The 12 core BDW processor performs 18% slower than the 14 core HSW processor due to 22% lower base frequency. The dots in the graph show the performance improvement over IVB processor. The graph shows that the performance increases with increasing number of cores for the BDW processors. The top bin 22 core BDW performs 112% better than 12 core IVB.




Figure10: Performance comparison of multiple generations of Intel processor with F1ATPase benchmark

Figure 10 compares performance of multiple generations of Intel processors with the medium sized (327,506 Atoms) NAMD benchmark, named F1ATPase. It can be seen from the graph that HSW performance improvement is 33% better than IVB and BDW performance is up to 62% better than HSW.




Figure 11: Performance comparison of multiple generations of Intel processors with STMV benchmark

Figure 11 plots the performance comparison graph among multiple generations of Intel processors for the large size (1,066,628 Atoms) NAMD benchmark. As per this graph, the BDW processors are performing up to 63% better than HSW processor.

It can be seen from figures 9, 10 and 11 that the larger datasets make better use of the computation power and the relative performance with additional cores is better as compared with the smaller dataset.

Conclusion

This blog characterizes five Intel Broadwell processors and shows performance improvement for real time HPC applications. Additional cores, along with the higher memory frequency support in Intel Broadwell processors, improve the performance of HPC workloads specifically for compute sensitive workloads like NAMD. The performance of memory bandwidth sensitive workloads like WRF increase up to the 16 core processors, but the performance improvement for the 20 and 22 core processors is not as significant due to the lower memory bandwidth per core.