

Impact of Broadwell BIOS Options On CAE Applications

Authors: Mayura Deshmukh, Ashish K Singh, Neha Kashyap

[Last week's blog](#) on the "Broadwell Performance for HPC" series described the BIOS options and compared performance across generations of processors for molecular dynamic applications (NAMD) and Weather Research and Forecasting (WRF). This blog, third in the series, focuses on BIOS options for some HPC CAE applications for five different Broadwell Intel Xeon E5-2600 v4 series processor models. It aims to answer questions like, which snoop mode works best for my application and processor? Which BIOS System Profile would give the best performance?

There have been a few changes in the BIOS options for Broadwell as compared with the previous generation (Haswell). One of the major additions in the Broadwell BIOS is the "Opportunistic Snoop Broadcast" snoop mode in the Memory settings. This blog discusses performance of the applications for all four snoop modes: Opportunistic snoop broadcast (OSB), Early snoop (ES), Home snoop (HS) and Cluster on die (COD). For more information on the new BIOS options and snoop modes check [blog one](#) of this series.

The Dell BIOS "System Profile" setting can be set to either of the four pre-configured profiles: Performance Per Watt (DAPC), Performance Per Watt (OS), Performance (Perf.) and Dense Configuration or set to Custom. In the pre-configured profiles the Turbo Boost, C States, C1E, CPU Power Management, Memory Frequency, Memory Patrol Scrub, Memory Refresh Rate, Uncore Frequency are preset whereas for Custom the User can choose values for these options. For more information on System Profiles [check the link](#). DAPC and OS have shown to perform similarly in [past studies](#), and Dense Configuration performs lower for HPC workloads, so we will be focusing on DAPC and Performance Profiles in this study. The DAPC (Dell Active Power Control) Profile relies on a BIOS-centric power control mechanism. Energy efficient turbo, C States, C1E are enabled with the DAPC Profile. Performance Profile disables power saving features such as C-states, Energy efficient turbo and C1E. Turbo boost is enabled in both the System Profiles.

This blog discusses the performance of CAE applications with DAPC and Performance profile for each of the four snoop modes for five different Intel Xeon E5-2600 v4 series Broadwell processors. Table 1 shows the application and benchmark details and Table 2 describes the server configuration used for the study.

Table 1 - Applications and benchmarks

Application	Version	Metric	MPI	Benchmark
LS-DYNA®	8.0.0	Elapsed time	Platform MPI 9.1.0	<ul style="list-style-type: none"> car2car with endtime=0.02
STAR-CCM+®	10.04.011	Average Elapsed time	Platform MPI 9.1.3	<ul style="list-style-type: none"> Civil_20m EglinStoreSeparation HIMach10 Kcs Lemans_100m Lemans_17m Reactor9m TurboCharger Vtm
ANSYS® Fluent®	v16	Solver rating	Platform MPI 9.1.2.1	<ul style="list-style-type: none"> truck_poly_14m combustor_12m combustor_71m exhaust_system_33m ice_2m
OpenFOAM	2.4.0	Clock time	Open MPI 1.10.0	<ul style="list-style-type: none"> Cavity 1M Motorbike 11M

Table 2 - Server configuration

Components	Details
Server	PowerEdge R630
Processor	<ul style="list-style-type: none"> 2 x E5-2650 v4 12c, 2.2/1.8 GHz, 105W, 30MB cache 2 x E5-2690 v4 14c, 2.6/2.1 GHz, 135W, 35MB cache 2 x E5-2697Av4 16c, 2.6/2.2 GHz, 145W, 40MB cache 2 x E5-2698 v4 20c, 2.2/1.8 GHz, 135W, 50MB cache 2 x E5-2699 v4 22c, 2.2/1.8 GHz, 145W, 55MB cache
Memory	256GB - 16 x 16GB 2400 MHz DDR4 RDIMMs
Hard drive	6 x 300GB SAS 6Gbps 10K rpm
RAID controller	PERC H330 mini
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 Mode - Opportunistic Snoop Broadcast (OSB), Early Snoop (ES), Home Snoop (HS), Cluster on Die (COD) Node interleaving - Disabled
BIOS	2.0.0
iDRAC Firmware	2.30.30.02

LS-DYNA is a general-purpose finite element program from LSTC capable of simulating complex real-world structural mechanics problems. We ran the car2car benchmark with end time set to 0.02 with both the single precision avx2 and the single precision sse LS-DYNA binaries.

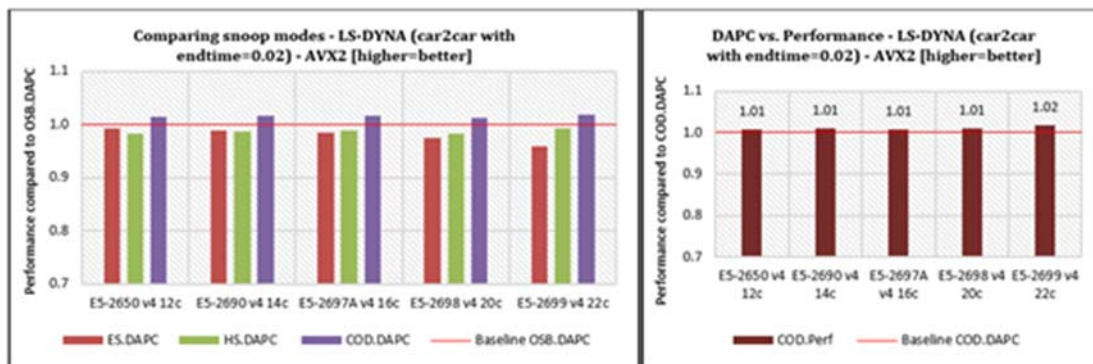


Figure 1: Comparing snoop modes and BIOS Profiles for LS-DYNA

The left graph in Figure 1 shows how better or worse the different snoop modes perform compared to the default setting of snoop mode = OSB and BIOS profile=DAPC (which is set at 1, the red line on the graph). Just changing the snoop mode to COD increases performance by 1-3% with either BIOS profiles across all the processor models. The performance with COD is closely followed by OSB followed by ES for lower core counts and HS for 16, 20 and 22 core processors. With ES mode, the system starts paying the penalty of having lower request tokens per core for higher core counts compared to the other snoop modes (for e.g. for 14 cores $128/14 = 9$ per core Vs. $128/22 = 5$ per

core for 22 core). All the snoop modes with the System Profile set to Performance follow similar pattern as DAPC. As shown in the graph on the right in Figure 1, changing the System Profile from DAPC to Performance can provide up to 2% performance benefit. The COD.Perf is the best option, about 2-4% better compared to OSB.DAPC across all processor models. The total 2-4% improvement with COD.Perf is accounted partially due to the change in snoop mode and partially due to change in the BIOS System Profile to Performance. We ran the car2car benchmark for all the combinations above with the sse LS-DYNA binary as well and noted similar behavior with the Performance System Profile and COD snoop mode being 2-6% better than the default OSB.DAPC. The avx2 binaries performed 12-19% better than the sse binaries across all the processor models.

CD-adapco® STAR-CCM+ is another CFD application widely-used by industry for solving problems involving fluid flows, heat transfer, and other phenomena. The STAR-CCM+ benchmarks results show a pattern similar to LS-DYNA in terms of snoop mode and System Profile.

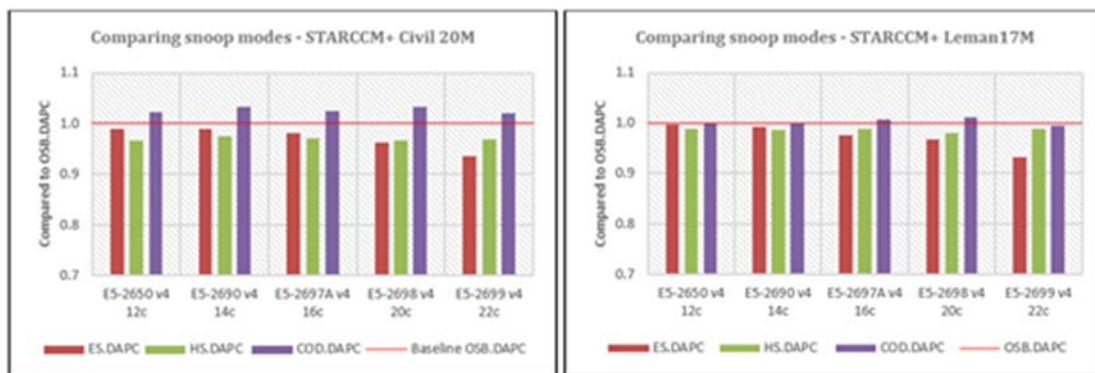


Figure 2: Comparing snoop modes for STAR-CCM+

Figure 2 compares the snoop modes for the Civil_20m and Lemans_17m benchmarks. For simplicity, data for these two benchmarks are shown. The other benchmarks datasets show results similar to the patterns in Figure 2. The BIOS profile in the graphs is set to DAPC and the snoop modes are compared against the default OSB snoop mode (which is set at 1, the red line on the graph). The COD is the best option for the Civil_20m benchmark, it is about 2-3% better for DAPC. For the Performance System Profile COD is 4-6% better for the Civil_20m benchmark (not shown in the graph). COD is followed by OSB and then ES for smaller core counts. Performance with ES though starts reducing as the cores increase similar to what was observed with LS-DYNA car2car benchmark case. The HIMach10 benchmark shows similar pattern to the Civil_20m benchmark. For the HIMach10 benchmark case the COD.Perf option is 2-7% better than the default OSB.DAPC.

All the other benchmarks (EglinStoreSeparation, Kcs, Lemans_100m, Reactor9m, TurboCharger, Vtm) show similar pattern to Lemans_17m. The COD and OSB perform similarly, there is only ~1% difference between OSB and COD

across the benchmark cases across all processor models. After COD and OSB, ES option is better for lower core counts and HS for 16, 20 and 22 core processors. As mentioned previously, the system in ES mode starts paying the penalty of having lower request tokens per core for higher core counts compared to the other snoop modes.

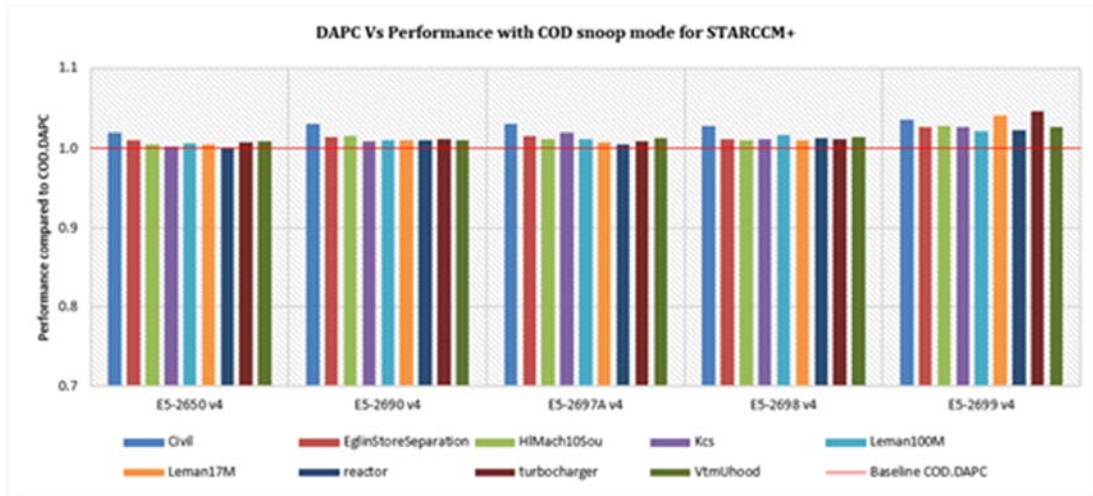


Figure 3: DAPC vs. Performance with COD snoop mode for STARCCM+

The graph in figure 3 compares the System Profile BIOS options DAPC and Performance. We are comparing the performance of COD.Perf with respect to COD.DAPC, which is the red baseline set at 1 in the graph. The Performance profile provides 2-4% benefit over the DAPC for the Civil_20m benchmark for all the processor models. Also for the high core count, E5-2699 v4 the Performance profile performs 2-5% better across all the benchmarks. For all the other processor models there is not a significant gain (only about 1%) with the Performance profile for all the benchmarks (except Civil_20m).

ANSYS Fluent is a computational fluid dynamics application. Fluent provides multiple benchmark cases. We picked four representative cases from the v16 benchmark suite: combustor_12m, combustor_71m, exhaust_system_33m and ice_2m and one from the older v15 benchmark suite: truck_poly_14m, to allow us to compare our data with previous generation processor models. The Fluent benchmarks exhibit a similar pattern as LS-DYNA and STAR-CCM+ benchmarks.

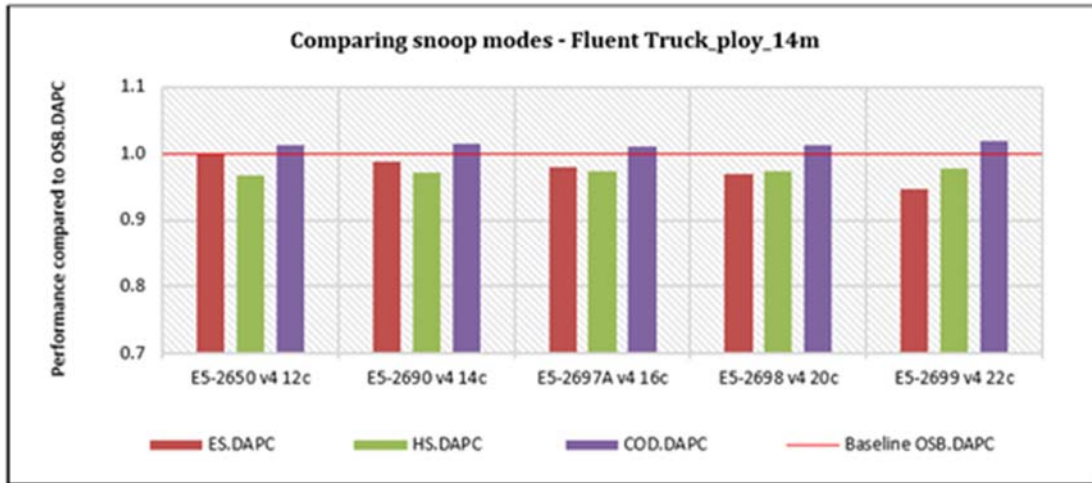


Figure 4: Comparing snoop modes for ANSYS Fluent

The graph in Figure 4 shows the performance of truck_ploy_14m for all the snoop modes compared to the default OSB.DAPC which is shown as the red baseline in the graphs. All the other benchmarks show a similar pattern. COD performs up to 2% better than OSB for truck_ploy_14m, combustor_12m and ice_2m. COD is about 5% better for combustor_71m and 6% better for exhaust_33m. COD is followed by OSB, followed by ES for lower core counts and HS for higher core count processors for all the benchmarks.

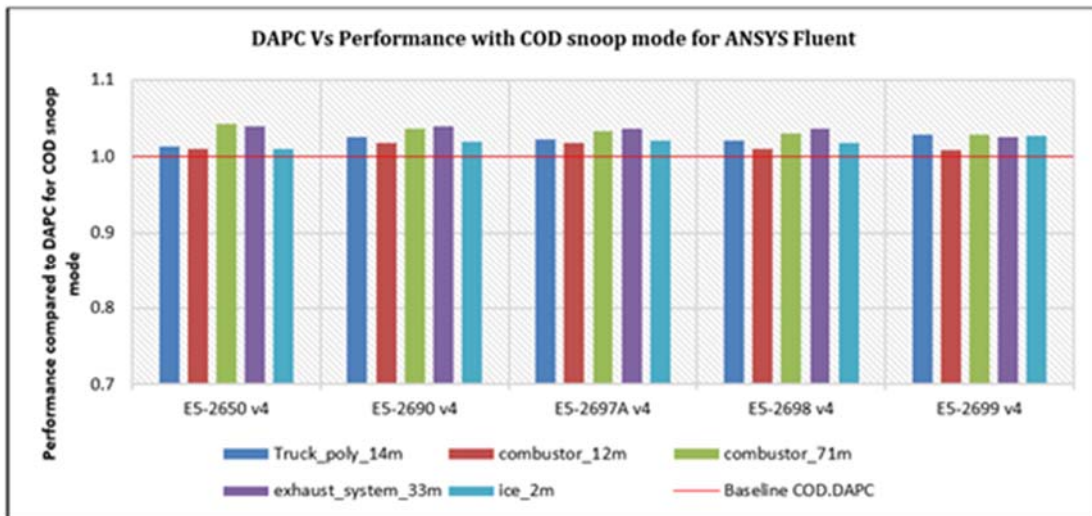


Figure 5: DAPC vs. Performance with COD snoop mode for ANSYS Fluent

Figure 5 shows the performance for Performance profile with respect to DAPC with COD set as the snoop mode for both options. DAPC is shown as the red baseline in the graph. The Performance BIOS profile option is about 4%

better for all the processor models for the larger combustor_71m and exhaust_33m benchmark cases. The Performance profile is 1-3% better for the other benchmark cases.

OpenFOAM (Open source Field Operation And Manipulation) is a free, open source software for computational fluid dynamics (CFD). OpenFOAM was compiled with -march=native / Broadwell option. We used the cavity-1M and motorBike-11M datasets which are modifications of the OpenFOAM tutorials/incompressible/icoFoam/cavity and tutorials/incompressible/simpleFoam/motorBike models respectively.

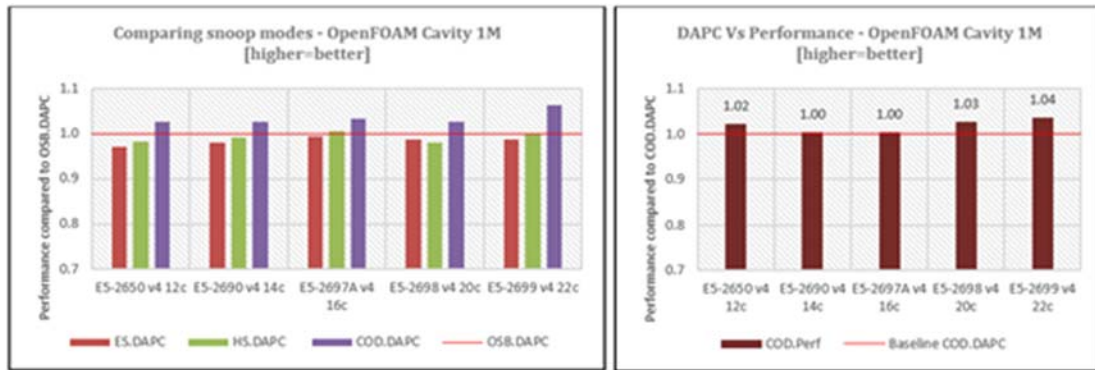


Figure 6: Comparing snoop modes and BIOS Profiles for OpenFOAM Cavity 1M benchmark

As shown in left graph of figure 6 for DAPC System Profile, the benchmark performance increases by 3-6% when in COD snoop mode when compared to OSB. ES and HS options perform up to 3% lower than OSB across all the processor models. The pattern is similar for the Performance System Profile, where COD is better by 3-7% followed by OSB. HS is lower than OSB but better than ES for all the processors models except for the 20core E5-2698 v4 where ES is 1% better than HS for DAPC profile and 7% better than HS for Performance System Profile. There is not a lot of difference in performance for DAPC Vs Performance profile especially for the higher frequency processors 14core E5-2690v4 and the 16core E5-2697A v4. For the other models the Performance profile shows up to 4% benefit as shown in the right graph of figure 6.

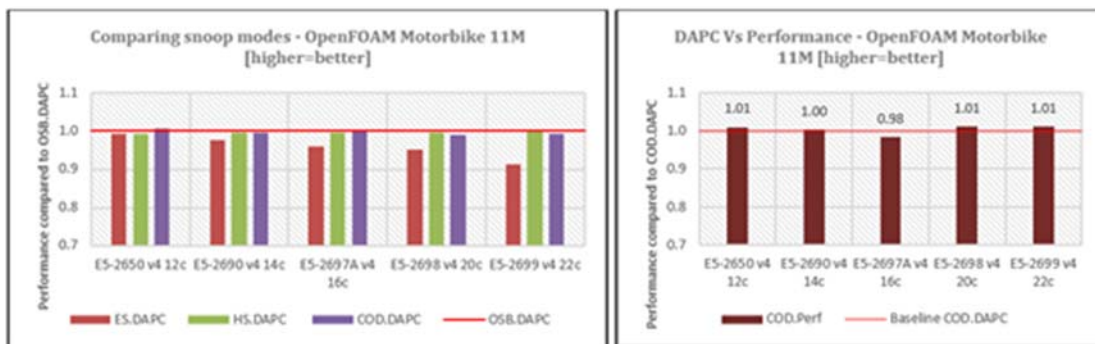


Figure 7: Comparing snoop modes and BIOS Profiles for OpenFOAM Motorbike 11M benchmark

For the openFOAM motorbike 11M benchmark the OSB, COD and the HS snoop modes perform similarly with about 1% variation. The performance for ES is low across all the processor models and it keeps on dropping as the number of cores increase as shown in the left graph of figure 7. The snoop modes with BIOS System Profile set to Performance follow exactly similar trend. As shown in the right graph on figure 3, the DAPC and Performance profiles show similar performance with Performance about 1% better in most cases except for the E5-2697A where the DAPC.COD was 2% better.

Conclusion

Most of the data sets used in this study show advantage of COD mode, but COD benefits codes which are highly NUMA optimized and where the dataset fit into the NUMA memory (that is half of each sockets memory capacity). OSB is a close second and a good option for codes with varying level of NUMA optimization; OSB is also the default memory snoop BIOS option. HS and ES perform slightly lower than COD and OSB. ES is better than HS for lower core counts but as the core counts increase ES starts paying the penalty of having lower request tokens per core for higher core counts compared to the other snoop modes. In terms of System Profile, Performance Profile performs slightly better than DAPC in most of the cases.

Be sure to check back next week for the last blog in the series which will compare the performance of HPC CAE applications across generations (Ivy-bridge vs. Haswell vs. Broadwell)