

Dell PowerEdge C4130 Performance with K80 GPUs - NAMD

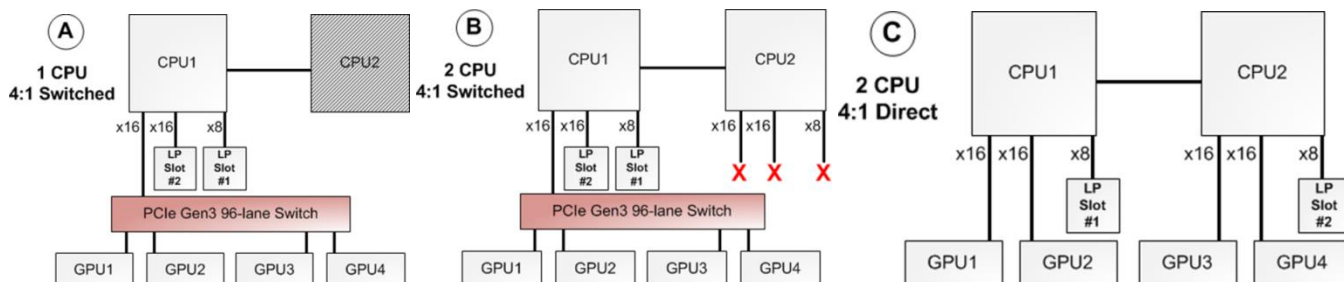
Authors: Saeed Iqbal and Mayura Deshmukh

The field of Molecular Dynamics (MD) has seen a tremendous boost in simulation capacity since the introduction of General Purpose GPUs. This trend is sustained by freely available feature-rich GPU-enabled molecular dynamics simulators like NAMD. For more information on NAMD, please visit <http://www.ks.uiuc.edu/Research/namd/> and for more information on GPUs visit <http://www.nvidia.com/tesla>.

Things only get better for NAMD with the release of the new Tesla K80 GPUs from NVIDIA. K80 offers significant improvements over the previous model the K40. From the HPC prospective *the most important improvement is the 1.87 TFLOPs (double precision) compute capacity*, which is about 30% more than K40. The auto-boost feature in K80 automatically provides additional performance if additional power head room is available. The internal GPUs are based on the GK210 architecture and have a total of 4,992 cores which represent a 73% improvement over K40. The K80 has a total memory of 24GBs which is divided equally between the two internal GPUs; this is a 100% more memory capacity compared to the K40. The memory bandwidth in K80 is improved to 480 GB/s. The rated power consumption of a single K80 is a maximum of 300 watts.

Combining K80s with the latest high GPU density design from Dell, the PowerEdge C4130, results in an extraordinarily powerful compute node. The C4130 can be configured with up to four K40 or K80 GPUs in a 1U form factor. Also the uniqueness of PowerEdge C4130 is that it offers several workload specific configurations, potentially making it a better fit, for MD codes in general, and specifically for NAMD.

The PowerEdge C4130 offers five configurations, noted here as “A” through “E”. Part of the goal of this blog is to find out which configuration is best suited for NAMD. The three quad GPU configurations “A”, “B” and “C” are compared. Also the two dual GPU configurations “D” and “E” are compared for users interested in lower GPU density of 2 GPU per 1 rack unit. The first two quad GPU configurations (“A” & “B”) have an internal PCIe switch module which allows seamless peer to peer GPU communication. We also want to understand the impact of the switch module on NAMD. Figure 1 below shows the block diagrams for configurations A to E



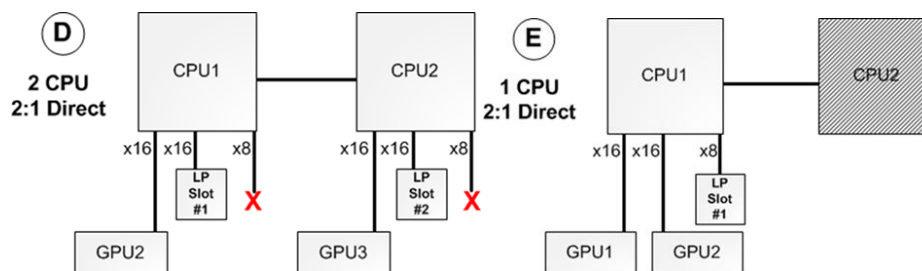


Figure 1: C4130 Configuration Block Diagram

In this blog we evaluate improved NAMD performance with Tesla K80 GPUs. Two proteins F1ATPASE and STMV, which consist of 327K and 1066K atoms respectively, have been selected due to their relatively large problem size. The performance measure is in “days/ns”, which shows the number of days required to simulate 1 nanosecond of real-time. Table 1 gives more information about the hardware configuration and application details used for the tests.

Table 1: Hardware Configuration and Application Details

| Server | C4130 Prototype System |
|-------------------------|---|
| Processor | 1 or 2 x Intel Xeon CPU E5-2690 v3 @ 2.6 GHz (12 core) |
| Memory | 64GB or 128GB @ 2,133MHz |
| GPU | 2 or 4 x NVIDIA Tesla K80 (4,992 CUDA cores, base clock 562 MHz, boost clock 875MHz, power 300W) |
| Power supply | 2 x 1,600W |
| Operating System | RHEL 6.5 – kernel 2.6.32-431.el6.x86_64 |
| BIOS options | System Profile – Performance Logical Processor – Disabled Power Supply Redundancy Policy – Not Redundant Power supply Hot Spare – Disabled |
| CUDA Version and driver | CUDA 6.5 (340.46) |
| BIOS firmware | 1.1.0 |
| iDRAC firmware | 2.02.01.01 |
| NAMD | 2.9 |




Figure 2: NAMD performance and acceleration on the five C4130 configurations.

Figure 2 shows the performance of NAMD on the PowerEdge C4130. Configurations “A”, “B” and “C” are four GPU configurations. However the acceleration on NAMD also seems to be sensitive to number of CPUs, e.g., there is a significant difference in the acceleration between “A” and “B”. “B” has an additional CPU compared to “A”. Among the three quad GPU configurations the current version of NAMD performs best on configuration “C”. The difference in two highest performing configurations “C” and “B” is the manner in which GPUs are attached to the CPU. The balanced configuration “C” has 2 GPUs attached to 2 CPUs resulting in 7.8X acceleration over the CPU-only case. The same four GPUs attached via a Switch module to a single CPU results in about 7.7X acceleration.

On the dual GPU configurations, “D” performs better with 5.9X acceleration compared to 4.4X in configuration “E”. The fact that “D” does better is in line with the assumption that a 2nd CPU is helpful for NAMD, as we saw “B & C” (the dual CPU quad GPU configurations) performing best.

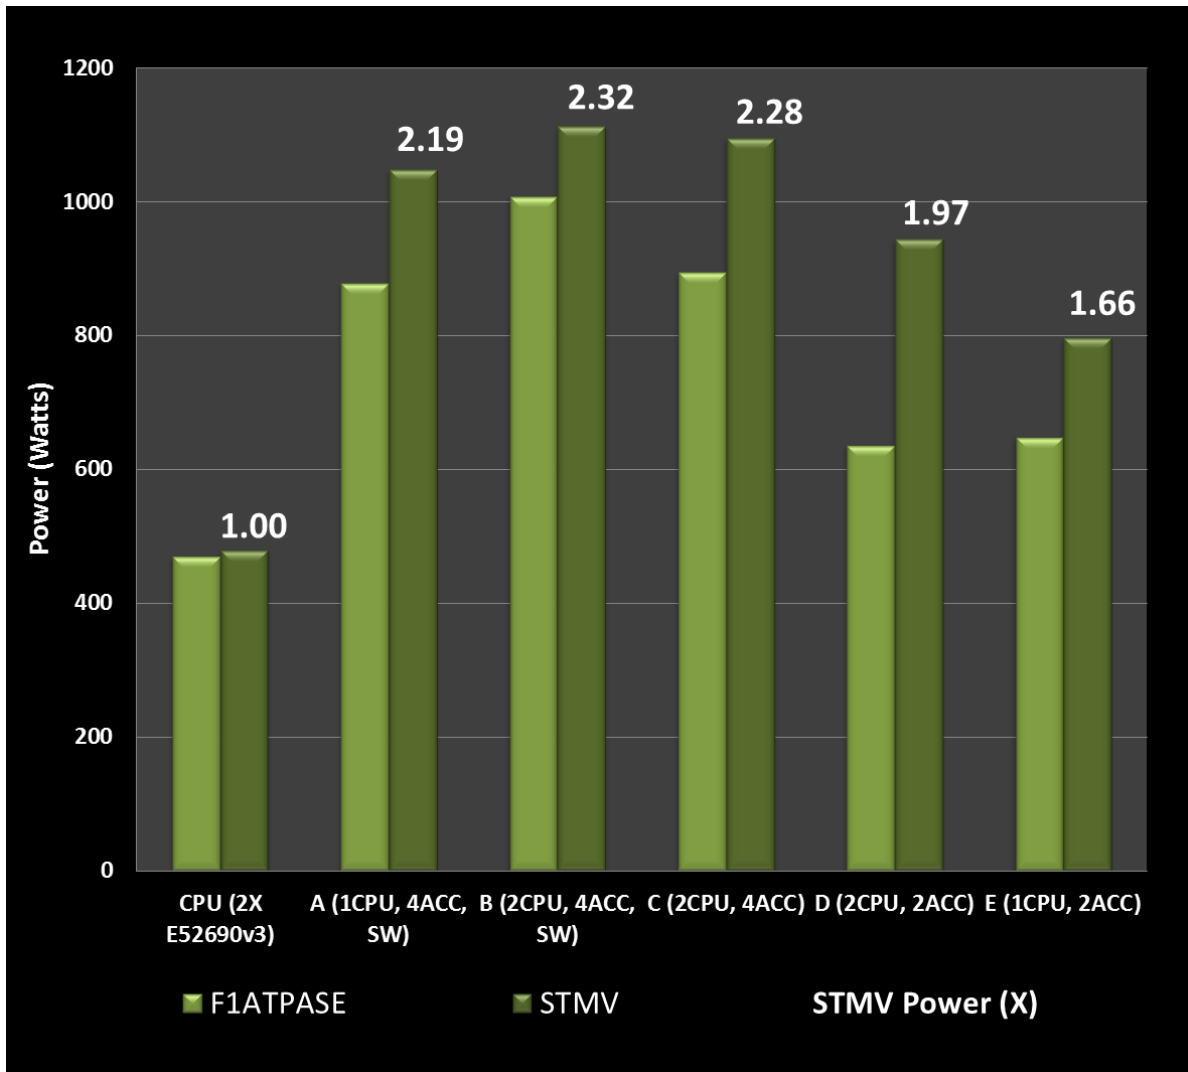


Figure 3 shows the power consumption results of running NAMD for the runs in Figure 2.

As shown, the power consumption for quad GPU configurations is about 2.1X to 2.3X resulting in accelerations from 4.4X to 7.8X. Configuration "C" does the best from performance per watt perspective (an acceleration of 7.8X for 2.3X more power). The power consumption of "D" is higher than "E" due to higher due to the additional CPU and improved utilization of GPUs resulting in better acceleration.

In conclusion, first, using K80 GPUs can substantially accelerate NAMD and it does that in a power efficient way too. Second, the balanced configurations seem to do better with NAMD. The configuration "C" and "D" are best for NAMD, the particular choice depends on required GPU density/U.