

Accelerating HPL using the Intel Xeon Phi 7120P Coprocessors

Authors: Saeed Iqbal and Deepthi Cherlopalle

The Intel Xeon Phi Series can be used to accelerate HPC applications in the C4130. The highly parallel architecture on Phi Coprocessors can boost the parallel applications. These coprocessors work seamlessly with the standard Xeon E5 processors series to provide additional parallel hardware to boost parallel applications. A key benefit of the Xeon Phi series is that these don't require redesigning the application, only compiler [directives](#) are required to be able to use the Xeon Phi coprocessor.

Fundamentally, the Intel Xeon series are many-core parallel processors, with each core having a dedicated L2 cache. The cores are connected through a bi-directional ring interconnects. Intel offers a complete set of development, performance monitoring and tuning tools through its Parallel [Studio](#) and [VTune](#). The goal is to enable HPC users to get advantage from the parallel hardware with minimal changes to the code.

The Xeon Phi has two modes of operation, the offload mode and native mode. In the offload mode designed parts of the application are "offloaded" to the Xeon Phi, if available in the server. Required code and data is copied from a host to the coprocessor, processing is done parallel in the Phi coprocessor and results move back to the host. There are two kinds of offload modes, non-shared and virtual-shared memory modes. Each offload mode offers different levels of user control on data movement to and from the coprocessor and incurs different types of overheads. In the native mode, the application runs on both host and Xeon Phi simultaneously, communication required data among themselves as need. A good reference on Xeon Phi and modes can be found [here](#).

The Intel Xeon Phi 7120P coprocessor has the highest performance among the Phi series. It has 61 cores and is rated at 1.2 TFLOPS and can handle 244 threads. The 7120P also has the Intel Turbo Boost technology. Bulk of the compute intensive calculations are done on the coprocessors.

The PowerEdge C4130 offers five configurations "A" through "E". Among these configurations there are two balanced configurations. The two balanced configurations "C" and "D" are considered for acceleration in this blog. Configuration "C" is the balanced four coprocessor option with two coprocessors attached to each host processor, and configuration "D" has a single Xeon Phi attached to the each host processor. Table 1 gives more details of these configurations. The details of the two configurations are shown in the Table 1. The block diagram of configuration "C" and "D" is shown in Figure 1.

This blog shows the results of acceleration observed on the C4130 with Intel Xeon Phi 7120P in configuration "C" and "D".

Table 1: Two balanced C4130 Configurations C and D

4130 Configuration	ACCs	CPUs	Switch Module (SW)	Acc/CPU ratio	Comments
C	4	2	N	2	Balanced
D	2	2	N	1	Balanced

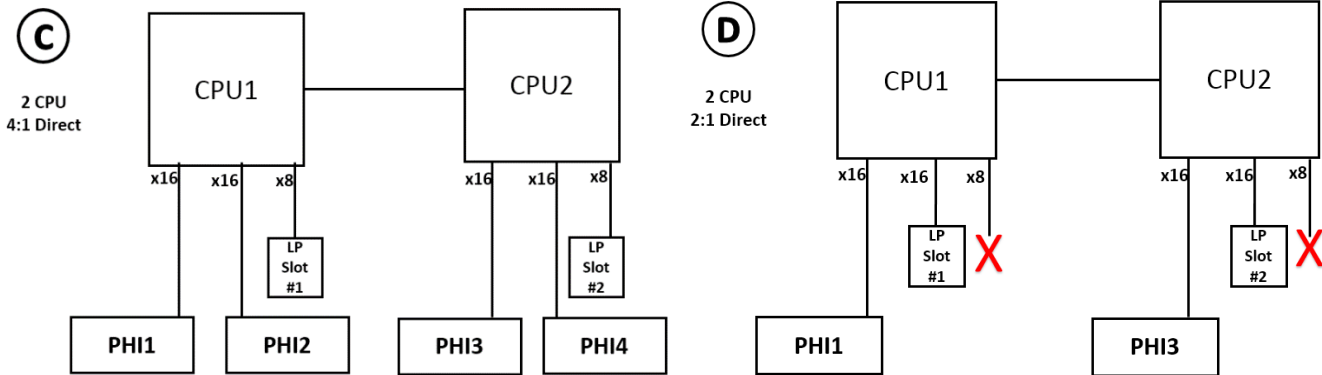


Figure 1: PE C4130 Configuration Block Diagram

Table 2 gives more information about the hardware configuration used for the tests.

Table 2: Hardware Configuration

Server	PowerEdge C4130
Processor	2 * Intel Xeon CPU E5-2690 v3 @ 2.6 GHz (12 core)
Memory	128GB @ 2133MHz
Accelerator	2 or 4 * Intel Xeon Phi Coprocessor 7120P Memory -16GB # of cores - 61 Processor Base Frequency- 1.238 GHz Max Turbo Frequency - 1.33GHz
Power supply	2 x 1600W
Operating System	RHEL 6.5 – kernel 2.6.32-431.el6.x86_64
BIOS options	System Profile – Performance
	Logical Processor – Disabled
	Snoop Mode-Early Snoop
	Power Supply Redundancy Policy – Not Redundant
MPSS Version	Intel MPSS 3.4

BIOS firmware	1.1.0
iDrac firmware	2.02.01.01
Compilers	Intel C++ Composer XE 2015
MPI	Intel MPI 5.0.1
MKL	Intel MKL 11.2
High Performance Linpack	Version 2.1

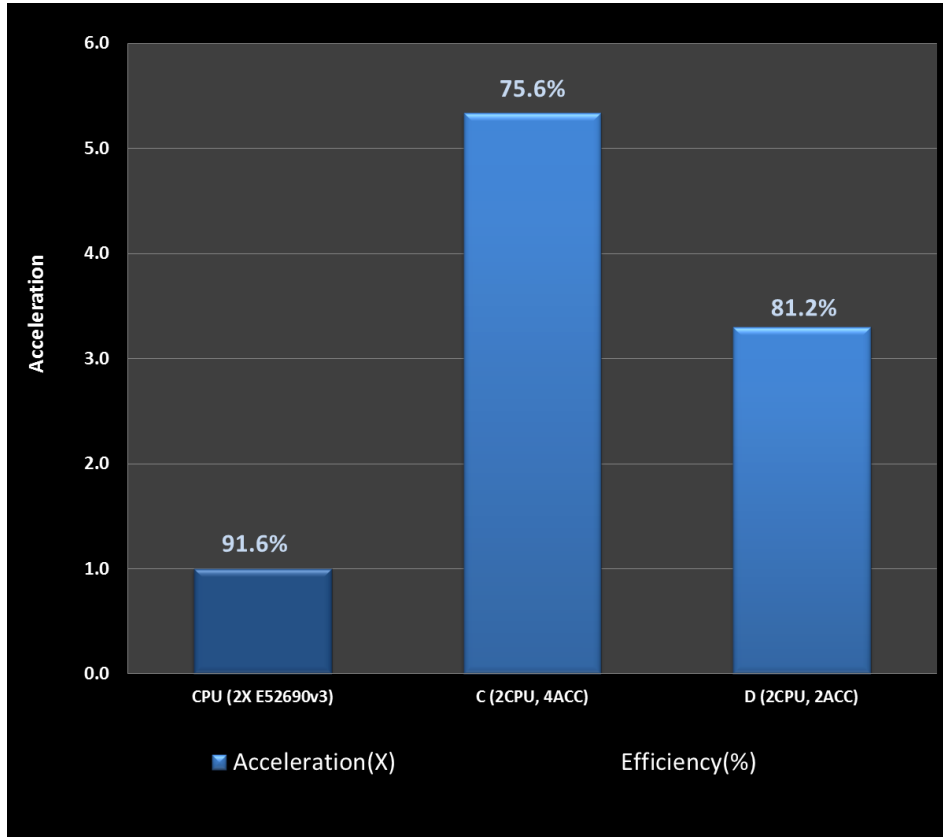


Figure 2: HPL acceleration (FLOPS compared to CPU only) and efficiency on the C4130 configurations.

Figure 2 illustrates the HPL performance on the PowerEdge C4130 Server. The Offload execution mode was used for all the runs. In this mode the application splits the workload where highly-parallel code is offloaded to the coprocessor, and the Xeon host processors primarily run serial code. Configuration C has 2 Phi's connected to each CPU and configuration D has single Phi connected to each CPU. ECC is enabled and the turbo mode is disabled across all the runs.

Intel Xeon Phi coprocessor provides more efficient performance for highly parallel applications like HPL. In the above graphs the CPU only performance is shown for reference. The compute efficiency for CPU only configuration is 91.6% whereas Configuration C has a compute efficiency of 75.6% and configuration D has 81.2%. It is already known that the CPU-only configurations in general have higher efficiency when compared to CPU plus Phi configurations. Higher efficiency is observed in configuration D compared C. *Compared to the CPU only*

configuration the HPL acceleration for configuration C with 4 Xeon phis is 5.3X and for configuration D with 2 Xeon phis, it is 3.3X.

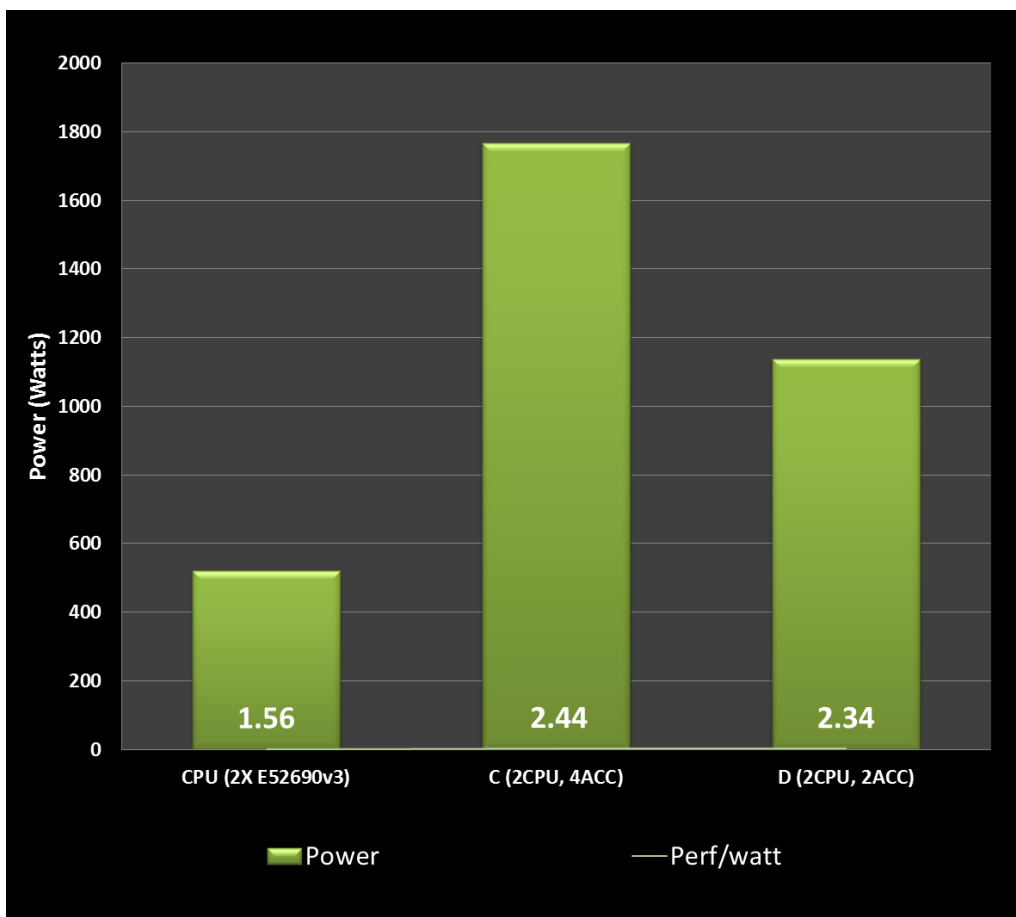


Figure 3: Total power and performance/watt on the C4130 configurations.

Figure 3 shows the associated power consumption data of the HPL runs for CPU-only configuration, Configuration C and D. In general, accelerators can consume substantial power when loaded with compute-intensive workloads. The power consumption of CPU-only configuration is 520W whereas the power consumption increases for configurations C and D. Each Intel Xeon PHI 7120P co-processor can consume power up to 300 watts. The power consumption for configuration C and D is 3.3X and 2.1X respectively when compared to the CPU-only configuration.

The Intel Xeon Phi 7120P co-processor provides high performance, memory capacity and good performance-per-watt metrics. Configuration C shows a performance-per-watt of 2.44 GFLOPS/w and configuration D shows 2.34 GFLOPS/w whereas CPU-only configuration gives 1.56 GFLOPS/w.