

M A R V E L L

WHITE PAPER

Energy, Oil & Gas on Marvell ThunderX2

Server Processor Business Unit
Marvell

June 2019

ABSTRACT

Computational modeling plays an important role in the search and extraction of fossil fuels and natural gas. Increases in computational capability are letting scientists simulate reservoirs and plan extraction techniques with increasingly finer resolution grids and with greater accuracy and urgency. In this white paper, we present why the Marvell® ThunderX2® Arm®-based server processor provides significant value for end users deploying applications in the field of Energy, Oil, & Gas.

Marvell ThunderX2 – Designed for High-Performance Computing

Before we dig deeper in to the ThunderX2 architecture, it is important to understand a typical way of categorizing high-performance computing (HPC) deployments. High-performance computing is broadly divided into two categories¹. **Capability computing** refers to using a large-scale HPC installation to ***solve a single problem in the shortest possible time***, for example simulating weather models on a Tier-0 HPC system. **Capacity computing** refers to optimizing system efficiency to ***solve as many mid-size or smaller problems as possible at the same time at the lowest possible cost***, for example automobile manufacturers using rented (on-demand) HPC resources to simulate numerous drive models for their products.

It is clear from the definitions above that an HPC customer must incorporate components of cost/CPU, # of CPUs and performance (execution time)/CPU. Thus, the cost of an HPC application can be summarized as

cost \propto CPU hours

where CPU hours = #cores x execution time

The other important factor driving design decisions of an HPC system is the application suite running on a cluster. Although capability computing targets application runs with the lowest execution time, excessive application scaling may deliver diminishing returns in performance improvement while linearly increasing CPU-hours. This is an unacceptable scenario that leads to inefficient resource utilization. Similarly, although capacity computing targets low-cost HPC computation, excessive slowdown of application runs may have unacceptable impact, e.g., reduction in productivity of engineers waiting for simulation results.

The reality is not all applications are created equal. Some applications scale well, and some don't. In the field of Energy, Oil and Gas, OpenSBLI, PSDNS, OpenFOAM and NEMO are examples of applications with good scalability. Such application variability necessitates making thorough choices in HPC system design components including the processor. The ThunderX2

architecture has been built ground up to strike the right balance between efficiency and throughput and thus provide best in class efficiency per \$ and throughput per \$, making it an easy choice for HPC system architects designing next-generation clusters.

Marvell ThunderX2 Overview

ThunderX2 is Marvell's second generation of Arm-based server processors targeted for the HPC, Cloud/Hyperscale and Enterprise market segments. Based on the 64-bit Armv8-A architecture, the Marvell ThunderX2 processor includes a custom core built using the Arm architectural license. Fully out-of-order, it supports simultaneous multithreading, providing ample compute for data center workloads. In addition, the Marvell ThunderX2 processors support dual socket configurations essential for scaling out applications. The processors are manufactured using a power efficient TSMC 16nm process technology and are fully compliant with Arm's Server Base System Architecture (SBSA) standard. Key features include:

- ✓ Up to 32 cores with support for simultaneous multithreading
- ✓ DDR4 72-bit memory controllers, supporting up to 2666MT/s DRAM
- ✓ Up to 8 DDR4 memory controllers
- ✓ 56 lanes of PCIe and 14 PCIe controllers
- ✓ 2 x SATA 3.0 & USB3 for boot
- ✓ Server class virtualization & RAS features
- ✓ Extensive power management
- ✓ Socketed LGA or BGA for flexibility




Figure 1: High Level Block Diagram of Marvell ThunderX2

Marvell ThunderX2 HPC-Focused Features

Most CPUs and GPUs have flop/s and integer op/s to fully saturate the memory subsystem, hence memory bandwidth – not floating-point capability - is currently the gating factor for

many HPC workloadsⁱⁱ. With 8 memory channels (as opposed to six for Intel Xeon Skylake), ThunderX2 provides a 33% greater memory bandwidth capability to effectively utilize more cores and achieve a higher per core utilization. The resulting higher *operational floating-point efficiency*ⁱⁱⁱ explains why the shorter dual per-core 128-bit vector units on the ThunderX2 cores can compete so effectively against the wider dual per-core Intel Skylake AVX-512 vector units on floating-point performance. ThunderX2 SoCs can deliver over 1 Tflop/s of double-precision and over 2 Tflop/s of single-precision performance in a dual socket configuration.

A higher operational floating-point efficiency explains why the shorter dual per-core 128-bit vector units on ThunderX2 can compete so effectively against the wider dual per-core Intel Skylake AVX-512 vector units on floating-point performance that dominates HPC applications.

In summary, Marvell ThunderX2 can deliver competitive or superior performance on parallel and floating-point dominated HPC workloads because it:

- (1) Has a higher per SoC core count that translates to an overall increase in vector units since each core has two 128-bit vector units.
- (2) Is a balanced architecture that delivers a higher *operational flop/s* on real HPC applications because the memory system can better supply data to the Arm cores and dual per-core vector units. This is seen in the benchmark and optimized performance results discussed below.
- (3) Can better support those HPC applications that benefit from shorter 128-bit vector operations as opposed to longer 512-bit vector operations.
- (4) Does not underclock. Briefly, the size of the dual per-core vector units on the die means that the Marvell ThunderX2 SoC can preserve floating-point performance when using the per-core vector units rather than underclocking to keep the SoC within thermal design limits.

The reality is not all applications are created equal. Some applications scale well, and some don't. In the field of Computational Fluid Dynamics (CFD), OpenSBLI, PSDNS, OpenFOAM and NEMO are examples of applications with good scalability. Such application variability necessitates making thorough choices in HPC system design components including the processor. The ThunderX2 architecture has been built ground up to strike the right balance between efficiency and throughput and thus provide best in class efficiency/\$ and throughput/\$, making it an easy choice for HPC system architects designing next-generation clusters.

HPC Synthetic Benchmark Results

Synthetic benchmark results demonstrate the design choices to create a more balanced processor for HPC workloads. Specifically, ThunderX2 matches or exceeds the performance or scaling efficiency of the latest x86 processors on STREAM, HPCG, and HPL synthetic benchmarks.

The HPL benchmark shows the benefit of ThunderX2 processors' balanced design, including the internal interconnect bus clocked at core clock rate. In contrast, Intel processors adjust their frequency according to workload. While this can provide power efficiency, an independent CERN presentation notes, "Intel processors adjust their frequency according to workload. Highly threaded, vectorized code may run in a lower frequency range. This behavior can confuse scaling studies, and it may reduce the benefit of AVX and AVX-512 vectorization."^{iv}

Intel processors adjust their frequency according to workload. Highly threaded, vectorized code may run in a lower frequency range. This behavior can confuse scaling studies, and it may reduce the benefit of AVX and AVX-512 vectorization.

Summary from the independent, third-party CERN presentation, Turbo Boost Up, AVX Clock Down: Complications for Scaling Tests

As a more balanced processor, Marvell ThunderX2 can deliver superior performance on many HPC workloads to compete very effectively against the latest AVX-512 powered Intel processors.

STREAM

The [STREAM](#) benchmark is the industry standard for measuring memory bandwidth. A 1.27x higher STREAM benchmark result (shown in Table 1) demonstrates that a pair of 32-core Marvell ThunderX2 processors can process more data per second compared to a dual-socket Intel Xeon Skylake Gold 6148. Users can verify that the Intel results reported here are consistent with those reported by third-party studies, including Colfax Research^v.

HPC users will appreciate the comparison as The Next Platform describes the 20-core Intel® Xeon® SP 6148 Gold processor as, "a typical chip for HPC customers."

Workload	Marvell ThunderX2	Intel SKL Gold 6148	ThunderX2 Improvement over SKL
STREAM	251 GB/s	198 GB/s	1.27

Table 1: STREAM benchmark numbers. (gcc7.2 compiler on ThunderX2, icc18 compiler on SKL)

The consistency of the ThunderX2 memory subsystem performance across all aspects of the STREAM benchmark can be seen in the table below.

STREAM	Memory Bandwidth (GB/s)
Copy	240
Scale	236
Add	252
Triad	252

Table 2: Complete STREAM benchmark numbers. (gcc7.2 compiler on ThunderX2, icc18 compiler on SKL)

Comparing single and dual socket performance demonstrates the efficacy of the Marvell CCPI2™ interconnect as both single and dual socket configurations deliver high performance. The CCPI2 interconnect provides full cache coherency between the ThunderX2 processors in a dual socket system.

HPCG

The High-Performance Conjugate Gradients (HPCG) benchmark is based on an iterative sparse-matrix conjugate gradient kernel with double-precision floating-point values. HPCG is representative of HPC applications governed by differential equations, which tend to have much stronger needs for high memory bandwidth, low latency, and accessing data using irregular patterns.

The following benchmark results were independently determined by HPE.

Workload	Marvell ThunderX2	Intel SKL Gold 6148	ThunderX2 Improvement over SKL
HPCG	35 GF/s	36 GF/s	0.97

Table 3: HPCG benchmark numbers. (gcc7.2 compiler on ThunderX2, icc18 compiler on SKL, results provided by HPE)

The rapid evolution of Arm for HPC applications and capability of the memory system design can be seen in the HPCG results as ThunderX2 delivers effectively equivalent floating-point

performance compared to a latest generation x86 processor that has a significantly higher peak-floating point capability.

HPL

The HPL benchmark has low memory bandwidth utilization and is a flop/s dominated synthetic benchmark. Results show a ThunderX2 processor pair was able to more effectively use all its cores and 128-bit vector units to deliver 91.82% of its peak theoretical floating-point performance on HPL without underclocking or dynamic clock scaling. As a result, ThunderX2 was able to deliver 1.077 TF/s of HPL performance on the HPL benchmark.

The Marvell ThunderX2 processor pair was able to more effectively use all its cores and 128-bit vector units on the HPL benchmark to deliver 92% of its theoretical peak floating-point performance.

In contrast, the Intel Xeon Skylake Gold 6148 processor was only able to deliver 72% efficiency using its dual per-core AVX-512 vector units as measured independently by Fujitsu^{vi}. The Fujitsu results confirm that other processors in the Intel Xeon Scalable Processor family deliver lower HPL efficiency than ThunderX2.

Workload	Marvell ThunderX2	Intel SKL Gold 6148
HPL	91.82% efficiency ^{vii}	72% efficiency ^{viii}

Table 4: Comparative HPL benchmark results.
(gcc7.2 compiler on ThunderX2, icc18 compiler on SKL)

Benefits of Running Energy, Oil and Gas Workloads on Marvell ThunderX2 Systems

Comparing datasheet specs, ThunderX2 offers higher throughput, capacity and Perf per \$ compared to the volume Intel Xeon Skylake SKUs as shown in Figure 2.

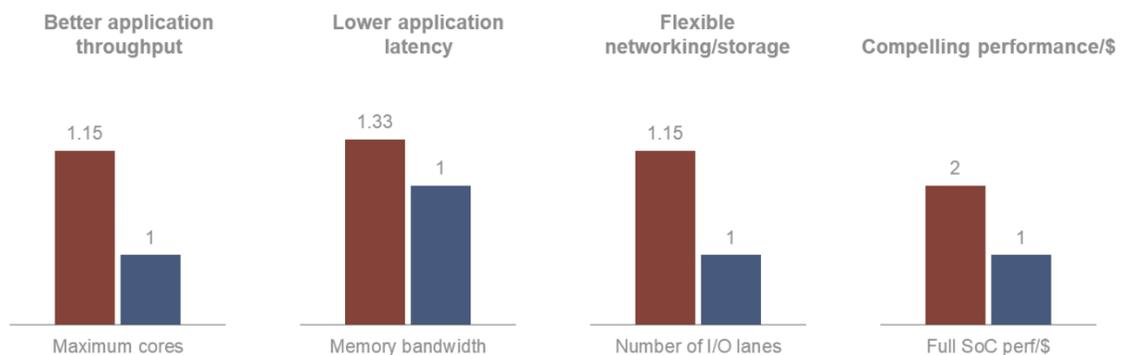


Figure 2: comparison of throughput, capacity and performance per \$; Marvell ThunderX2 represented in red and Intel Xeon Skylake represented in blue

However, to prove the value of ThunderX2 in solving real world energy, oil and gas problems, we ran representative benchmarks and the results are discussed below.

AWP-ODC

AWP-ODC is a seismic wave propagation modeling code and demonstrates that finite difference codes map nicely and run efficiently on ThunderX2. AWP-ODC is written in Fortran 77 and Fortran 90. It uses the standard MPI-2 and MPI-IO HPC message and parallel IO frameworks.

Computationally, AWP-ODC utilizes a staggered-grid finite difference scheme to approximate the 3D velocity-stress elastodynamic equations to simulate the dynamic rupture and wave propagation that occurs during an earthquake. Dynamic rupture produces friction, traction, slip, and slip rate information on the fault. The moment function is constructed from this fault data and used to initialize wave propagation for earthquake study and analysis.

Workload	Marvell ThunderX2	Intel SKL Gold 6148	ThunderX2 Improvement over SKL
AWP-ODC	23.11 seconds	27.84 seconds	1.2x

Table 5: Comparative AWP-ODC benchmark results. (Arm 18.1 Compiler + Arm Perf Libraries on ThunderX2, icc18 compiler on SKL; based on Marvell internal tests; note that lower is better)

The AWP-ODC result demonstrates that superior performance can be achieved on Fortran as well as C/C++ codes.

OpenFOAM®

Originally developed as an alternative to early simulation engines written in Fortran, OpenFOAM is a modular C++ framework aiming to simplify writing custom computational fluid dynamics (CFD) solvers. The code makes use of object-oriented features in C++, which enables a powerful, extensible design methodology. OpenFOAM's flexibility comes at the cost of additional code complexity; for example, components are compiled into separate dynamic libraries, which are then combined at runtime based on the user's input to form the required executables. OpenFOAM is almost entirely memory bandwidth bound. Table 6 shows that ThunderX2 with gcc performs well relative to Intel Skylake with gcc.

Workload	Marvell ThunderX2	Intel SKL Gold 6148	ThunderX2 Improvement over SKL
OpenFOAM	92 seconds	116 seconds	1.26x

Table 6: Comparative OpenFOAM benchmark results on a single node (gcc on ThunderX2, gcc compiler on SKL; based on Marvell internal tests; note that lower is better)

OpenSBLI

OpenSBLI is a grid-based finite difference solver used to solve compressible Navier-Stokes equations for shock-boundary layer interactions. The application uses Python to automatically generate code to solve the equations expressed in mathematical Einstein notation and uses the Oxford Parallel Structured (OPS) software for parallelism. As a structured grid code, it should be memory bandwidth bound under the Roofline model, with low computational intensity from the finite difference approximation. Measurements are made with the ARCHER benchmark which solves a Taylor-Green vortex on a grid of $1024 \times 1024 \times 1024$ cells (around a billion cells). Table 7 shows that ThunderX2 with gcc performs well relative to Intel Skylake with gcc.

Workload	Marvell ThunderX2	Intel SKL Gold 6148	ThunderX2 Improvement over SKL
OpenSBLI	392 seconds	562 seconds	1.43x

Table 7: Comparative OpenSBLI benchmark results on a single node (gcc on ThunderX2, gcc compiler on SKL; based on Marvell internal tests; note that lower is better)

PSDNS

PSDNS is a highly parallelized application code used for performing direct numerical simulations (DNS) of three-dimensional unsteady turbulent fluid flows, under the assumption of statistical homogeneity in space. A system of partial differential equations expressing the fundamental laws of conservation of mass and momentum is solved using Fourier pseudo-spectral methods in space and explicit Runge-Kutta integration in time. The pseudo-spectral approach requires multiple FFTs to be taken in three directions per time step, resulting in communication-intensive operations due to transposes involving collective communication among processors. Table 8 shows that ThunderX2 with gcc performs well relative to Intel Skylake with icc.

Workload	Marvell ThunderX2	Intel SKL Gold 6148	ThunderX2 Improvement over SKL
PSDNS	236 seconds	270 seconds	1.14x

Table 8: Comparative PSDNS benchmark results on a single node (gcc on ThunderX2, icc18 compiler on SKL; based on internal Marvell tests; note that lower is better)

Conclusion

HPC deployments are driven by cost metrics that are directly proportional to the processor cost, performance and number of processors. By providing a balance of compute, throughput and efficiency, the Marvell ThunderX2 delivers best-in-class perf per \$ and efficiency per \$ for

running HPC workloads with different scaling capabilities. For customers looking to run cost effective Energy, Oil and Gas applications, the Marvell ThunderX2 is an ideal processor choice.

References

ⁱDarko Zivanovic, Milan Pavlovic, Milan Radulovic, Hyunsung Shin, Jongpil Son, Sally A. McKee, Paul M.Carpenter, Petar Radojković and Eduard Ayguadé, 2016. Main Memory in HPC: Do We Need More, or Could We Live with Less? ACM Trans. Embedd. Comput. Syst. V, N, Article 000 (2016), 25 pages.

ⁱⁱ See the ExaNoDe Report on the HPC application for a more detailed discussion: <http://exanode.eu/wp-content/uploads/2017/04/D2.5.pdf>.

ⁱⁱⁱ See section 5.1 Memory bandwidth vs. FLOPs analysis in the ExaNoDe Report on the HPC application bottlenecks at <http://exanode.eu/wp-content/uploads/2017/04/D2.5.pdf>.

^{iv} See the 12/17 presentation by Steve Lantz “Turbo Boost Up, AVX Clock Down: Complications for Scaling Tests” at

<https://indico.cern.ch/event/668302/contributions/2732551/attachments/1576588/2489822/TurboBoostUpAVXClockDown.pdf>, and the 11/17 Cloudflare blog post, “On the dangers of Intel’s frequency scaling” at <https://blog.cloudflare.com/on-the-dangers-of-intels-frequency-scaling/>.

^v Colfax, A Survey and Benchmarks of Intel® Xeon® Gold and Platinum Processors at <https://colfaxresearch.com/xeon-2017/>.

^{vi} Efficiency as reported by Fujitsu (<https://sp.ts.fujitsu.com/dmsp/Publications/public/wp-performance-report-primergy-rx2540-m4-ww-en.pdf>).

^{vii} Efficiency provided by Marvell.

^{viii} Efficiency as reported by Fujitsu (<https://sp.ts.fujitsu.com/dmsp/Publications/public/wp-performance-report-primergy-rx2540-m4-ww-en.pdf>).

Marvell Semiconductor, Inc.
5488 Marvell Lane
Santa Clara, CA 95054, USA

Tel: 1.408.222.2500
www.marvell.com

Copyright © 2019 Marvell. All rights reserved.

Marvell, the Marvell logo, ThunderX2 and CCPI2 are registered trademarks or trademarks of Marvell and/or its affiliates in the US and/or elsewhere. Other names and brands may be claimed as the property of others.