Leveraging LS-DYNA® With the Latest Intel Technologies

Nick Meng[1], Ram Ramanujam[1], Jason Wang[2], Robert Lucas[2]

[1] Intel Corporation
[2] Livermore Software Technology Corporation (LSTC)

Abstract

In this paper we discuss Intel's continued optimization efforts with LS-DYNA® and demonstrate the impact of new Intel technologies. Two different approaches to exploit Intel® Advanced Vector Extensions 512 (AVX-512) are shown: LS-DYNA® Explicit using Intel compiler vectorization techniques and LS-DYNA® Implicit using Intel® Math Kernel Library (MKL) for accelerating dense matrix computational kernels. Numerical accuracy of simulation results for LS-DYNA® Explicit comparing Intel® SSE2 and Intel® AVX-512 is also explored. Finally, we reveal the benefits of Intel® Optane™ DC Persistent Memory technology for LS-DYNA® Implicit simulations. For our studies we used the Topcrunch benchmarks, ODB-10M & car2car models, for LS-DYNA® Explicit and AWE benchmarks, CYL1E6 & CL2E6 models, for LS-DYNA® Implicit.

Introduction

The typical problem size of a LS-DYNA® Explicit model has moved to more than 10 million elements, and implicit models with hundreds of millions of degrees of freedom (DOF) have been developed for Turbine Engine and Bio-Engineering analyses. Consequently, LS-DYNA users have increased the number of cores used for simulation workloads to reduce the turnaround time of these larger models, and the scaling of LS-DYNA at the lowest cost and highest performance is becoming more important. The process of improving application performance on single core and then scaling across many nodes is a major focus area for both Intel and LSTC. Herein we introduce the next generation of Intel® technologies and present their impact on the performance of LS-DYNA. These technologies, the 2nd generation Intel® Xeon® Scalable Processor (code named Cascade Lake-SP), Intel® Advanced Vector Extensions 512 (Intel® AVX-512), Intel® 3rd generation Omni-Path Architecture (Intel® OPA) interconnect fabric, Intel® Optane™ DC persistent memory (a new class of memory and storage technology), enhanced Intel® Fortran Compiler 2018, Intel® MPI 2018 library, and Intel® MKL 2018 library with Intel® AVX-512 support, contribute to the most effective integrated scalable solution for LS-DYNA customers both in cost and performance. In particular, we highlight the performance advantages of Intel® AVX-512 for explicit analysis and of the Intel® MKL/AVX-512 library and Intel® Optane™ memory for implicit analysis.

2nd generation Intel® Xeon® Scalable Processor Family

The 2nd generation Intel® Xeon® Scalable Processor (code named Cascade Lake-SP) and Intel® Xeon® Advanced Performance Processor (code named Cascade Lake-AP, specifically 9200 series) introduce many innovative features compared to the previous-generation Intel® Xeon® processor products (References [1][2]). These new features include: higher number of processor cores (up to 56 cores per processor in the 9200 series); higher memory bandwidth (up to 2933 MT/s vs. previous 2663 MT/s in 8200 series; up to 12 memory channels vs. previous 6 memory channels in 9200 series); greater memory capacity

(up to 36 TB system memory using Intel® Optane DC persistent memory); new Intel® Speed Select Technology (configurable core/frequency processor attributes); and, continued improvements to Intel® AVX-512 instruction set architecture (e.g. new Vector Neural Network Instruction set (VNNI)).

The 2nd generation Intel® Xeon® Scalable Processor continues to leverage a mesh interconnect (Figure 1) introduced in the 1st generation (code named Skylake) to reduce memory latency.

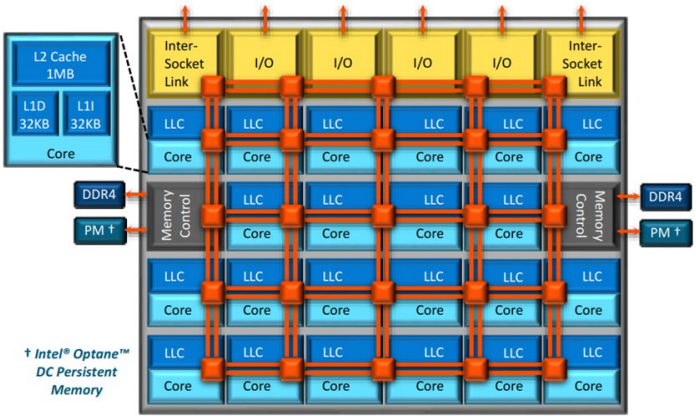

Figure 1: *A representation of generalized CPU layout and mesh interconnect in the 1st and 2nd generation of Intel® Xeon® Scalable Processors intended to help illustrate the concepts (not a definitive representation of the microarchitecture). Source: Reference [3].*

For applications that scale well with number of cores as well as need higher memory bandwidth, the 2nd gen Intel® Xeon® Scalable processor 9200 series (code named Cascade Lake-AP) delivers a multi-chip processor (MCP) up to 56 cores with 12 channels of memory with the highest Intel architecture FLOPS per rack, along with the highest DDR4 native memory bandwidth support of any Intel® Xeon® processor platform (Figure 2). For example, 2nd gen Intel® Xeon® Scalable processor 8200 series (code named Cascade Lake-SP) supports 6 channels to DDR4 memory.

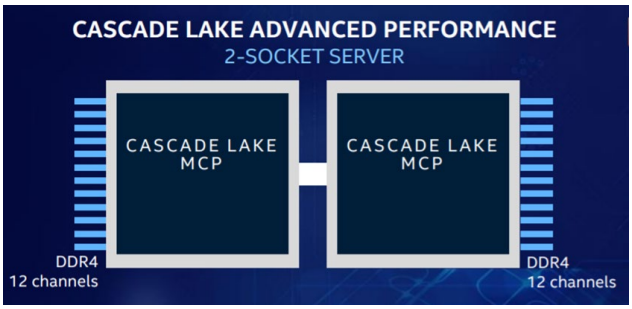

Figure 2: *2nd gen Intel® Xeon® Scalable processor 9200 series (code named Cascade Lake-AP). Source: Reference [4].*

The 1st generation Intel® Xeon® processor Scalable Family (known as Skylake-SP) introduced new Intel® AVX-512 instruction groups (AVX512CD, AVX512F, AVX512BW, AVX512DQ, and AVX512VL) (Reference [1]). Intel® AVX-512 is an extension to the CPU vector extensions, which includes 32 registers each 512-bit wide and eight dedicated mask registers which help in efficient AVX-512 instruction generation for branchy codes. Intel® AVX-512 also doubles the width of the register compared to its AVX2 predecessor, thus

computationally doubling the single and double precision floating point operations per clock cycle. Figure 3, for example, shows addition of eight double precision data elements in two 512-wide source registers in a single AVX-512 VADDPD instruction with masking. Note that, in contrast, Intel® AVX2 and SSE use 256-wide and 128-wide registers respectively. The 2nd generation Intel® Xeon® Scalable Processor provides improvement to Intel® AVX-512 instruction set architecture by introducing new Vector Neural Network Instruction set (VNNI). Intel also continues to enhance Intel compiler vectorization technology to leverage AVX-512 instructions benefiting applications such as LS-DYNA.

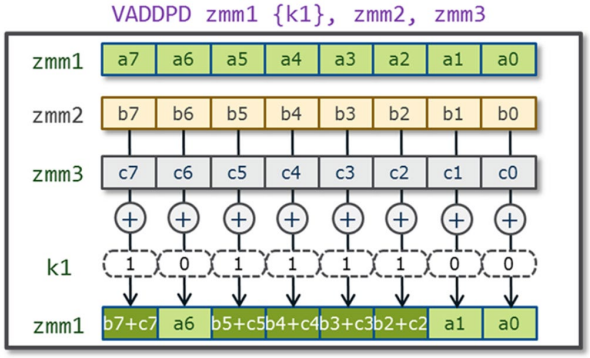

Figure 3: *Vector addition of packed double precision data elements using masking with Intel® AVX-512 instruction set. Source: Reference [1].*

Performance improvement of LS-DYNA using Intel® AVX-512

Intel® and LSTC continue to collaborate to optimize LS-DYNA leveraging SSE, AVX2 and AVX-512 instructions for improved performance. A key focus area has been to make source changes to allow Intel compiler to vectorize performance-critical loops, generating vector instructions including Intel® AVX-512 (e.g., using –xCORE-AVX512 compiler option) for best performance. The result of such optimizations is shown in Figure 4 with LS-DYNA R9.3_SP Explicit using the ODB-10M model on a 192-core cluster of 2nd generation Intel® Xeon® Scalable processors (Cascade Lake-SP). The baseline binary Intel® SSE2 binary is built with Intel® Compiler V13. Switching to Intel® Compiler V16 and still using Intel® SSE2, performance increases by about 4% indicating improvements in compiler vector code generation in newer versions of the Intel compiler. To get even a performance, Intel® Compiler V18 was used yielding a net 22% AVX-512 improvement for LS-DYNA over SSE2 baseline. Intel® has seen similar improvements for newer versions of LS-DYNA as well.

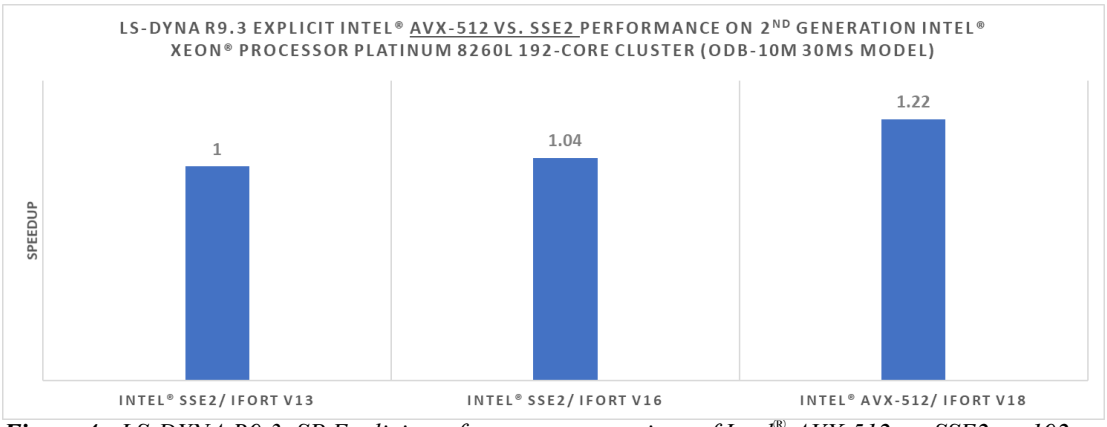

Figure 4: *LS-DYNA R9.3_SP Explicit performance comparison of Intel® AVX-512 vs. SSE2 on 192-core 2nd generation Intel® Xeon® processor (Platinum 8260L) cluster.* *(System configuration: 4 nodes 2-socket Intel® Xeon® Platinum 8260L processor, 384GB DDR4@2667Mhz, single 1.2TB Intel® 3710 SSD, Intel® Omni-Path Architecture fabric, Intel® Turbo Boost enabled, Intel® Hyper-Threading Technology disabled. 48 MPI ranks per node, Intel® 2018 MPI update 1, Linux release 3.10.0-957.5.1.el7.crt1.x86_64)*

Figure 5 demonstrates performance benefits of the recently announced 2nd generation Intel® Xeon® Scalable Processor Platinum 9242 (known as Cascade Lake-AP, 48 cores, 2.3Ghz) vs. Platinum 8260 (known as Cascade Lake-SP, 24 cores, 2.4Ghz) for LS-DYNA R9.3_SP Explicit on a single 2-socket node, almost a linear scaling with 1.95x speedup with the ODB-10M benchmark.

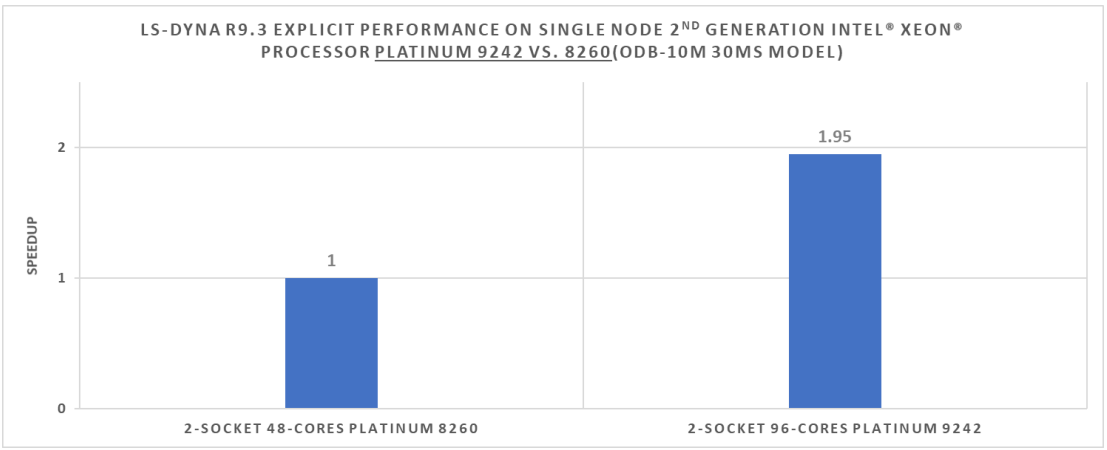

Figure 5. *LS-DYNA R9.3_SP Explicit performance comparison between 2nd generation Intel® Xeon® Scalable Processor Platinum 9242 vs. 8260.* *(System configuration: Single node 2-socket Intel® Xeon® Platinum 8260L processor, 384GB DDR4@2667Mhz, single 1.2TB 3710 Intel® SSD, Intel® Turbo Boost enabled, Intel® Hyper-Threading Technology disabled. 48 MPI ranks, Intel® 2018 MPI update 4, Linux release 3.10.0-957.5.1.el7.crt1.x86_64; Single node 2-socket Intel® Xeon® Platinum 9242 processor, 384GB DDR4@2667Mhz, single 1.2TB Intel® 3710 SSD, Intel® Turbo Boost enabled, Intel® Hyper-Threading Technology disabled, 96 MPI ranks, Intel® 2018 MPI update 4, Linux release 3.10.0-957.5.1.el7.x86)*

A key optimization strategy for HPC applications on Intel platforms is to leverage the Intel® MKL/BLAS library where possible. Intel® MKL is optimized for Intel® SSE, AVX/AVX2 and AVX-512, and can dynamically use the appropriate vector instruction set at run time depending on which Intel Architecture processor it runs on. LS-DYNA R9.1_DP and later versions are linked with an Intel® MKL library supporting for Intel® AVX-512 instructions. To quantify the performance impact of Intel® MKL for LS-DYNA Implicit a performance comparison of Intel® MKL using Intel® SSE/AVX/AVX-512 on 2nd generation Intel® Xeon® Scalable Processor Platinum 8260. LS-DYNA R9.3_DP binary built without Intel® MKL serves as the baseline. LS-DYNA Implicit was run with in hybrid mode with 2 MPI ranks and 16 OpenMP threads per rank using 2.5M DOF CYL1E6 model. Figure 6

shows a significant performance improvement for LS-DYNA R9.3_DP Implicit (double precision), 9.9x, using Intel® MKL AVX-512 vs. without Intel® MKL.

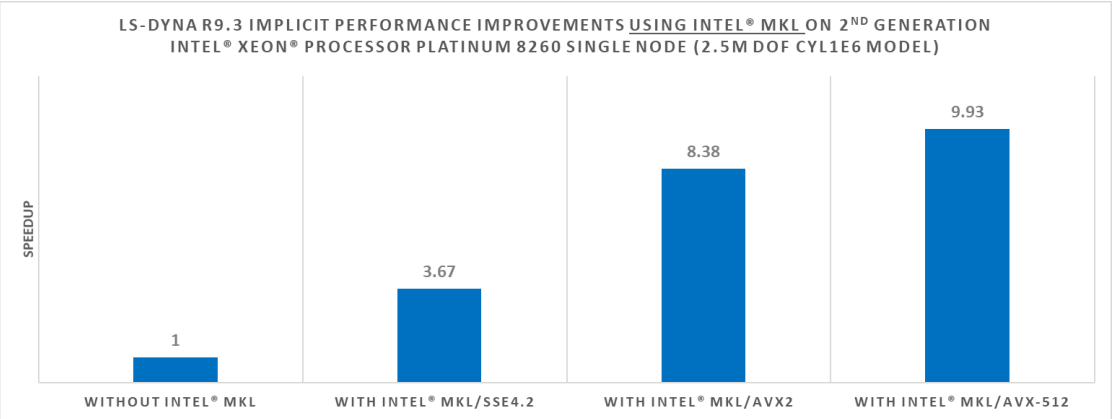

Figure 6: *LS-DYNA R9.3_DP Implicit performance impact using Intel® MKL/AVX-512 on 2nd generation Intel® Xeon® Scalable Processor Platinum 8260.* *(System configuration: Single node 2-socket Intel® Xeon® Platinum 8260 processor, 384GB DDR4@2933Mhz, single 1.2TB Intel® 3710 SSD, Intel® Turbo Boost enabled, Intel® Hyper-Threading Technology disabled. 2 MPI ranks x 16 OMP threads, Intel® MPI 2018.)*

The above performance improvements are demonstrated on standard LS-DYNA benchmarks. In addition, Intel and LSTC collaborate in optimizing proprietary customer workloads broadening benefits for the wider LS-DYNA user community.

LS-DYNA Explicit Intel® AVX-512 vs. SSE2 numerical accuracy comparions

Even though performance of LS-DYNA Explicit is improved as shown earlier in the paper, numerical accuracy of LS-DYNA Explicit model is still the top concern for all LS-DYNA users. LS-DYNA users expect to see repeatable, consistent results when they move on from Intel® SSE2 to AVX-512. In theory it's hard to get repeatable, consistent results comparing an SSE2 binary and an AVX-512 binary, as the order of operations changes, and hence the accumulation of round-off error varies. However, in practice correct results can be generated with an AVX-512 binary. This is demonstrated, for example, with the Car2car/120ms model. The simulation was run with Intel® SSE2 and AVX-512 binaries on the 2nd gen Intel® Xeon® Scalable processor Platinum 9242 node. Internal energy (Figure 7), resultant displacement in node 5341465 (Figure 8) and the X-acceleration in node 341513 were checked (Figure 9). As seen from Figures 7-9, the curves are very closely matched. Thus, LS-DYNA R9.3_SP Explicit performance is improved significantly with Intel® compiler V18 using Intel® AVX-512 instructions while the numerical accuracy remains effectively unchanged.

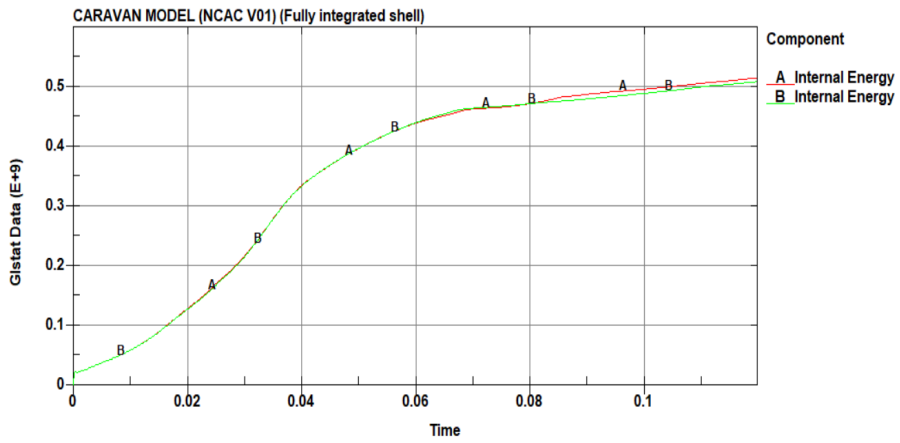

Figure 7: *LS-DYNA R9.3_SP Explicit Internal Energy comparison between Intel® AVX-512 (green) vs.SSE2 version (red). (System configuration: Single node 2-socket Intel® Xeon® Platinum 9242 processor, 768GB DDR4@2933Mhz, single 1.2TB Intel® 3710 SSD, Intel® Turbo Boost enabled, Intel® Hyper-Threading Technology disabled.)*

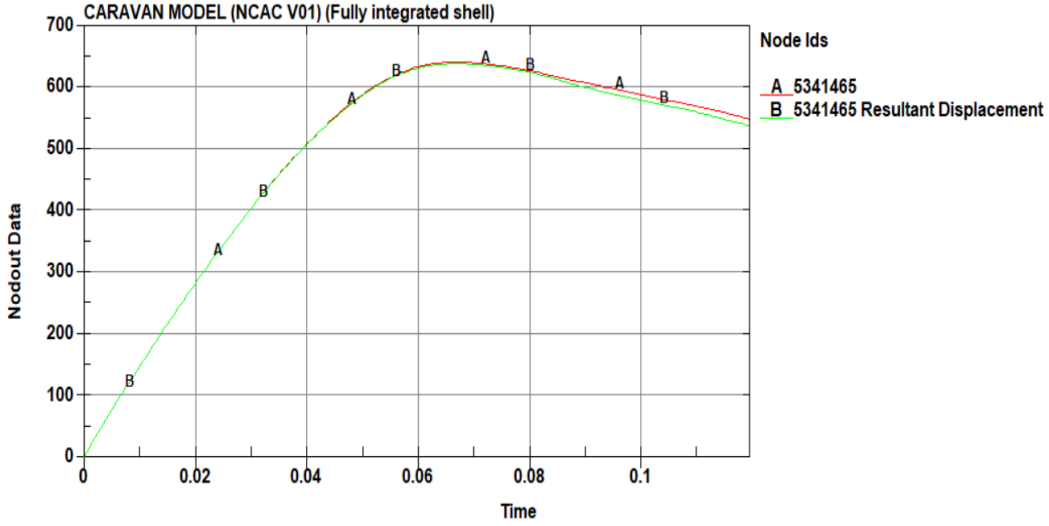

Figure 8: *LS-DYNA R9.3_SP Explicit Resultant Displacement comparison between Intel® AVX-512 (green) vs.SSE2 version (red) at Node 5341465.(System configuration: Single node 2-socket Intel® Xeon® Platinum 9242 processor, 768GB DDR4@2933Mhz, single 1.2TB Intel® 3710 SSD, Intel® Turbo Boost enabled, Intel® Hyper-Threading Technology disabled.)*

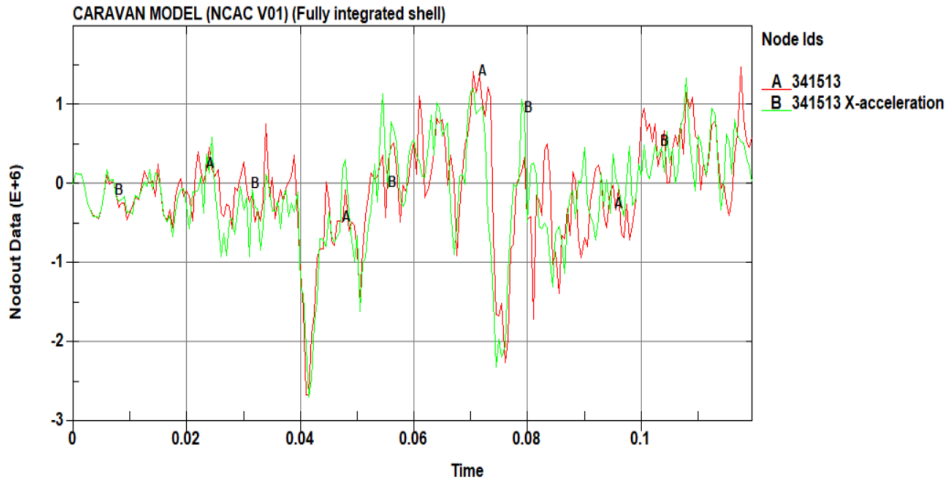

Figure 9: *LS-DYNA R9.3_SP Explicit X-acceleration comparison between Intel® AVX-512 (green) vs.SSE2 version (red)* at Node 341513.*(System configuration: Single node 2-socket Intel® Xeon® Platinum 9242 processor, 768GB DDR4@2933Mhz, single 1.2TB Intel® 3710 SSD, Intel® Turbo Boost enabled, Intel® Hyper-Threading Technology disabled.)*

LS-DYNA® Implicit with Intel® Optane™ DC Persistent Memory

As the problem size of LS-DYNA Implicit models increases rapidly, I/O performance becomes a serious bottleneck in existing systems when the solver uses out-of-core algorithm. A fast local file system and/or larger memory are required to address the I/O performance bottleneck. Simply increasing DDR4 memory capacity results in higher platform cost. Hence leveraging a fast local file system is often a preferred solution. RAID0/w SSD file system is widely accepted by customers today.

Intel® Optane™ DC Persistent Memory is an innovative memory technology that delivers a unique combination of affordable large capacity and persistence (i.e., non-volatility), filling a gap between DRAM and NAND (Figure 10). The Intel® Optane™ DC Persistent Memory can be provisioned in two modes. The first is called the Memory Mode. With Memory Mode, applications get a high capacity main memory solution at substantially lower cost and power, while providing performance close to that of DRAM, depending on the workload. No modifications are required to the application. The operating system sees the persistent memory module capacity as system main memory with DRAM memory acting as a cache. In Memory Mode, even though the media is persistent it will look volatile to application software. The second provisioning mode is called App Direct Mode where applications can access persistent memory capacity as non-volatile memory. In particular, the STorage Over App Direct mode (SToAD) allows access to persistent memory capacity using standard file APIs. No modifications to application code, or the file systems that expect block storage devices is required. SToAD provides high-performance block storage, without the latency of moving data to and from the I/O bus.

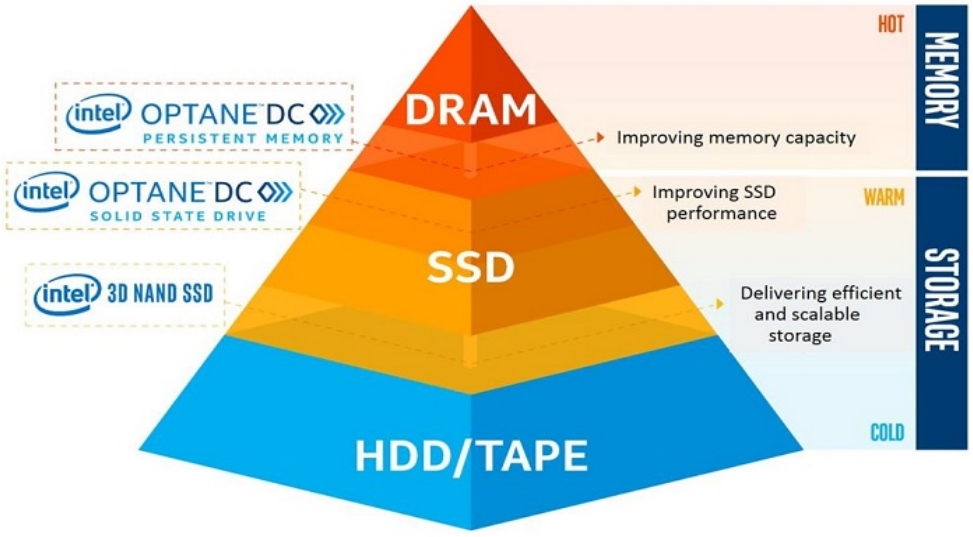

Figure 10: *Storage/memory performance vs. capacity pyramid. Source: Reference [5].*

To demonstrate the impact of Optane memory on the performance of LS-DYNA Implicit, AWE CYL2E6 model (6M DOFs) was benchmarked on 2nd generation Intel® Xeon® Scalable processor Gold 6248M (Cascade Lake-SP) with three modes (RAID0 mode, SToAD Mode, and Memory Mode). The baseline configuration used 192GB DDR4 memory with SSD based RAID0. The second configuration used 192GB DDR memory and 1.5TB of Optane™ DC persistent memory with SToAD Mode. The third configuration used 192GB DDR memory and 1.5TB of Optane DC persistent memory in Memory Mode. The

performance of SToAD Mode is 14% better than RAID0/w SSDs. The performance of Memory Mode is 75% faster than the baseline RAID0/w SSDs. This significant performance jump is primarily due to LS-DYNA Implicit being able to use in-core solver with the large memory capacity of Intel® Optane™ DC Persistent Memory (Figure 11).

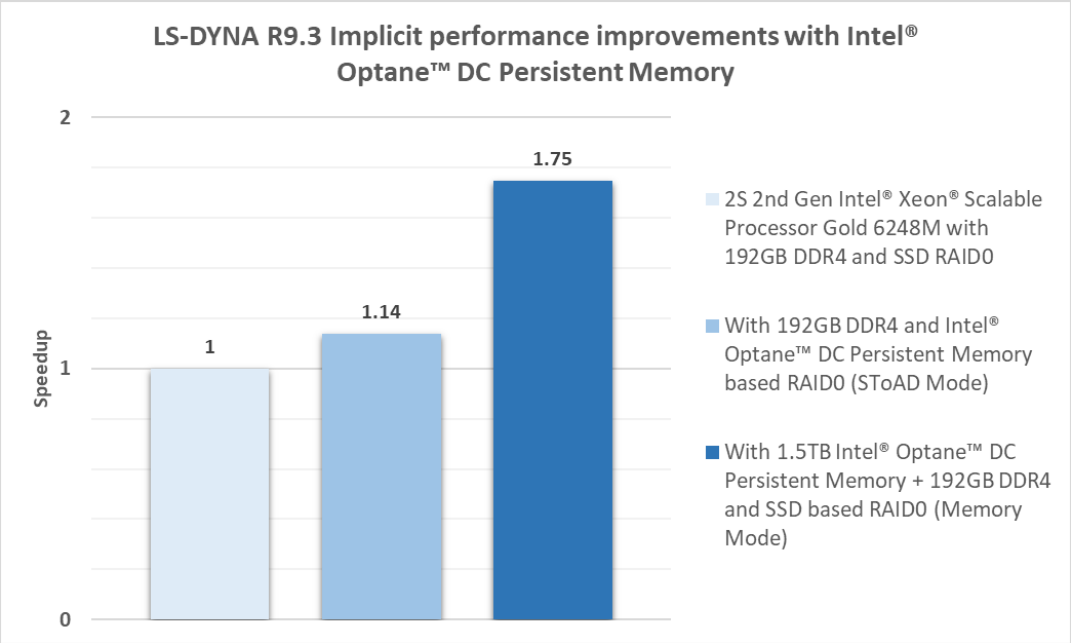

Figure 11: *LS-DYNA R9.3_DP Implicit performance improves 1.75x faster using Intel® Optane™ DC Persistent Memory vs. baseline SSD RAID0.* *(System configuration: 2-socket Intel® Xeon® Gold 6248M node, 192GB DDR4@2933Mhz, 1.5TB Optane™ DC Persistent Memory and RAID0/w 2x 1.2TB Intel® 3710 SSDs, Intel® Turbo Boost enabled, Intel® Hyper-Threading Technology disabled. 2 MPI ranks x 12 OMP threads, Intel® MPI 2018. Linux kernel: 4.18.19-100.fc27.x86_64)*

Conclusions

As seen in this paper, performance of both LS-DYNA Explicit and Implicit continue to improve using latest Intel technologies, specifically the 2nd generation Intel® Xeon® Scalable processors (known as Cascade Lake-SP/AP), Intel® Fortran Compiler V18, Intel® MKL/AVX-512 library, and Intel® Optane™ DC Persistent Memory. In addition, LS-DYNA Explicit numerical accuracy with Intel® AVX-512 is shown to match closely with the SSE2 version. Intel and LSTC remain committed to deeper collaborations optimizing LS-DYNA R9.3, R11 and other versions to bring the best value to LSTC customers on Intel architecture based platforms.

References

1. Intel® Xeon® Processor Scalable Family Technical Overview (https://software.intel.com/en-us/articles/intel-xeon-processor-scalable-family-technical-overview)
2. Second Generation Intel® Xeon® Scalable Processors (https://www.intel.com/content/www/us/en/products/docs/processors/xeon/2nd-gen-xeon-scalable-processors-brief.html)
3. Intel® VTune™ Amplifier Tuning Guide for the Intel® Xeon® Processor Scalable Family, 2nd Gen (https://software.intel.com/sites/default/files/managed/bf/4a/TuningGuide_IntelXeonProcessor_ScalableFamily_2ndGen.pdf)
4. Data Centric Business Update Media Briefing (https://newsroom.intel.com/wp-content/uploads/sites/11/2018/11/cascade-lake-advanced-performance-press-deck.pdf)

5. Intel® Optane™ DC Persistent Memory: A Major Advance in Memory and Storage Architecture (https://software.intel.com/en-us/blogs/2018/10/30/intel-optane-dc-persistent-memory-a-major-advance-in-memory-and-storage-architecture

Notices & Disclaimers

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit http://www.intel.com/benchmarks .

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration.

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice Revision #20110804

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel® Advanced Vector Extensions (Intel® AVX)* provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at http://www.intel.com/go/turbo.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel, the Intel logo, Intel Xeon, Intel Optane, Intel SpeedStep, and VTune are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as property of others.