



Altair OptiStruct*

Intel® Xeon® Processor E5-2600 v4 Product Family

Intel® SSD Data Center Family for PCIe*

Speeding up Altair OptiStruct* Simulations with the Intel® SSD Data Center Family for PCIe*

Altair OptiStruct* provides engineers and designers with a unified solution from concept to final design by leveraging advanced analysis capabilities and novel, optimization-driven simulation. In this process, the simulation time for one optimization iteration is a critical consideration, since it affects the computational speed and scalability of the entire design process.

Authors**Hongwei Zhou**Altair Engineering, Inc.:
hongwei@altair.com**Nick Meng**Intel Corporation:
nick.meng@intel.com

With ongoing advances in processing power for high-performance computing, a bottleneck has emerged in the memory domain. Consequently, the typically low memory requirements of iterative solvers have made them attractive in many engineering simulation applications. However, for structural solid mechanics applications, the direct solver is still preferred over the iterative solver due to greater robustness and accuracy in handling large, complex industrial models. Large factor matrices are typically involved, which can be as large as a few hundred gigabytes; as a result, memory is the critical impediment for large engineering model simulations using direct solvers.

While external memory is frequently used to circumvent this issue for large models, traditional SATA hard drives are limited by slow read/write rates, reducing the utility of out-of-core methods. Modern operating systems are typically capable of utilizing available system memory for I/O buffering, but while this practice is effective and commonly used, the computational price is high.

Domain-decomposition techniques based on distributed memory programming (DMP) models using message passing interface (MPI) techniques can significantly reduce memory usage on individual computational nodes. This method offers the additional advantage of increased computational scalability. DMP is highly effective because it distributes CPU computation, memory access, and I/O activity. However, a balance must be maintained between placing more MPI ranks and retaining memory for I/O buffering.

The emergence of solid-state drives (SSDs) raises sustained read/write performance, offering significant value for large-dataset I/O operations. SSDs are therefore well suited for OptiStruct simulations. This paper benchmarks two cases using Altair OptiStruct 2017.0: an engine block nonlinear static simulation and a buckling analysis on a full-car model. The tests were conducted using a two-socket system based on Intel® Xeon® processors E5-2667 v4 @ 3.2 GHz and 256 GB PC2400 DDR4 RAM, with five different storage configurations:


- **Single hard-disk drive (HDD)** : 1x Seagate Constellation.2 ST91000640NS (1 TB)
- **Single SATA SSD**: 1x Intel® SSD DC S3710 Series (800 GB)
- **Single NVMe SSD**: 1x Intel® SSD DC P3600 Series (2 TB)
- **RAIDO, SATA SSDs**: 4x Intel SSD DC S3710 Series (800 GB)
- **RAIDO, NVMe SSDs**: 3x Intel SSD DC P3600 Series (1.6 TB)

Case 1: Engine Block Nonlinear Static Analysis


Engine models typically consist of solid tetrahedral or hexahedral elements, generally leading to higher bandwidth in the sparse matrix and introducing a tremendously large factor matrix. For the 20 million degrees of freedom model, storing the factor matrix and internal stack, in addition to other working storage, requires more than 300 GB of space. In such cases, the sparse solver dumps the whole factor matrix onto the disk and retrieves it during the backward and forward substitution stage. On the outer layer, the nonlinear Newton algorithm requires one linear equation solution including matrix decomposition and backward/forward substitution in one iteration. The Newton equation solution process is computationally intensive.

Figure 1 illustrates that the use of RAID0-based SSD or NVMe SSD impacts the linear solver performance dramatically as the I/O bottleneck is resolved. The backward and forward substitution modules (RHSLV) demonstrate speedup by a factor of 10, due to the higher sustained READ I/O bandwidth (see case 1c). In addition, even the numerical factorization module (FCTRIZ) shows some computational advantage. The overall runtime is reduced by 50 percent. Figure 2 reveals that with RAID0 SSD or NVMe SSD, the system I/O wait time is essentially eliminated.

Case 1a: Overall Time on Nonlinear Analysis



Case 1b: Time Breakdown on Nonlinear Analysis



Case 1c: Sustained I/O Throughput



Figure 1. Timing comparison between overall wall-clock time and time on FCTRIZ and RHSLV modules. The sustained write and read I/O bandwidth were also measured.




Figure 2. Overview of CPU activities. The dark red area indicates the I/O wait at the system level.

Case 2: Full Car Buckling Analysis

Buckling analysis is another example of a typically I/O-intensive simulation. The Lanczos algorithm is utilized to conduct the eigenvalue solution in this analysis. This state-of-the-art shift-inverted technique requires linear equation backward/forward substitution in each Lanczos iteration. Like case 1, the factor matrix is transferred from disk to memory, which is the main I/O-intensive task. The other critical I/O task is at the Lanczos re-orthogonalization stage, where extensive I/O is required to transfer Lanczos vectors from memory to the disk after the process is completed.

Significant speedup is observed in the extraction of 120 eigenvalues on a 24 million degree-of-freedom model, as shown in Figure 3. This behavior is similar to that observed in the case 1 study.

With more MPI ranks in a node, significant slowdown for a traditional single HDD or SSD is typical as multiple MPI processes compete for I/O bandwidth. Therefore, it is important to ascertain the most efficient number of MPI ranks per node. The RAID SSD resolves such situations appropriately. A performance gain on three MPI ranks per node compared to two MPI ranks per node is also observed.

A significant performance differential exists between the results with the single NVMe drive versus the RAID array of SATA SSDs as shown in Figure 3, despite the fact that both options offer similar maximum bandwidth:

- **Four Intel SSD DC S3710 Series in RAID0 configuration:** 2200 MB/s read and 1840 MB/s write
- **Single Intel SSD DC P3600 Series:** 2600 MB/s read and 1700 MB/s write

The RAID array's faster completion of the full car buckling analysis workload stems from the effect of I/O randomization when multiple MPI jobs run simultaneously; that is, although each computational thread accesses data sequentially,

Case 2: Buckling Analysis




Figure 3. Timing comparisons measured using wall-clock time between five storage systems and different MPI configurations.

varying I/O sizes create quasi-random read-write behavior on the SSDs. For that reason, 4K IOPS performance represents the workload the most, rather than bandwidth. Indeed, the IOPS specifications of the storage options are more illuminating:

- **Four Intel SSD DC S3710 Series in RAID0 configuration:** 39,000 x 4 = 156,000 IOPS (for writes)
- **Single Intel SSD DC P3600 Series:** 56,000 IOPS (for writes)

It would be valuable in future work to compare performance using the Intel® SSD DC P3700 Series. Like the Intel SSD DC S3710 Series, this high-endurance NVMe device is designed for write-intensive workloads; the 2 TB model delivers 175,000 IOPS for writes. Testing with this device in single-SSD and RAID0 configurations would help further illuminate the performance potential of RAID arrays with matching stripes compared to single SSDs for multiple concurrent MPI jobs, due to higher capacity, bandwidth, and IOPS.

Conclusion

SSD technology fundamentally changes the implementation of the direct sparse solver, because there is less distinction between in-core and out-of-core performance. This game-changing technology significantly improves simulation performance on several critical analysis types in the automotive and aerospace industry. Altair OptiStruct continuously optimizes the solver algorithm and parallelization techniques for new storage devices to enhance CAE simulation performance and usability.

Learn more about Intel® Xeon® processors: www.intel.com/xeon

Learn more about Intel SSDs: www.intel.com/ssd

Learn more about Altair OptiStruct: www.altairhyperworks.com/product/OptiStruct

Solution provided by:



Altair



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/performance/datacenter>.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document. You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at www.intel.com.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel and Altair disclaim all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Copyright © 2017 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

*Other names may be trademarks of their respective owners. 0617/KE/MESH/PDF 336079-001US