

GPU Computing for Small & Medium-sized Problems with Abaqus 2016

William VILLERS

Director of Engineering, TEN TECH LLC

Abstract: Every year, significant development effort is put into GPU Acceleration of the Abaqus solver. However, the majority of published benchmarks, only highlights very large dynamics models. In this document, we discuss our findings in terms of performance benefits of GPU computing for a range of medium-sized problems with Abaqus 2016 as typically encountered in a Small Business Environment.

Keywords: Abaqus, GPU Computing, Scalability, HPC, Cloud Computing, NVIDIA, Rescale.

1. Introduction

Following a token usage discussion with one of our Abaqus customers, we wanted to have a clear view of how much benefit GPU Computing in Abaqus 2016, specifically for smaller models. We created a relatively small model and went on to test GPU Computing with Abaqus 2016. The recorded results were very favorable: 30%-50% speed increase over a range of linear and non-linear analysis with the help of GPU Computing, as can be seen on the figure below.




Figure 1: 2015 Preliminary Abaqus GPU Computing Study

These positive results and the interest this original study generated (see original article at <http://tinyurl.com/tentechllc-gpu>) encouraged us to explore Abaqus 2016's GPU Computing further with different models and hardware configurations, which is the object of this document.

2016 Science in the Age of Experience
<http://www.3ds.com/events/science-in-the-age-of-experience>

2. Benchmark Model & Hardware

TEN TECH LLC is an Aerospace & Defense contractor, specializing in Defense Electronics. It is no surprise that our test model would be representative of our daily activities.

Hardware-wise, our current in-house workstations and workgroup analysis server were put to the task. To push the study a little further, we also employed the cloud-based simulation HPC solver provided by Rescale.

2.1 Benchmark Model

For all of our studies, a typical Defense Electronics Subsystem, a VPX Single Board Computer (SBC), was utilized. The 3D CATIA CAD assembly of the SBC was brought into Abaqus/CAE via the CATIA Associative Interface, thus creating a ready-to-mesh analysis assembly.

The Finite Element Model created is composed of brick and tetrahedron solid elements, for a total of 3 million degrees of freedom, a relatively modest model size, as can be seen on Figure 2.




Figure 2: Single Board Computer & Abaqus/CAE Model

Standard MIL-STD-810G harsh environment shock and vibration analysis as well as classic low-cycle fatigue thermal expansion analysis were selected as test cases. Three separate cases were created to test the relative benefit of GPU computing for different types of solutions: linear statics, non-linear statics and non-linear statics with surface-surface contact and friction.

2.2 Hardware Configurations

In order to thoroughly investigate the benefits of GPU computing, the test models were solved using Abaqus 2016 on several in-house hardware configurations: HP Z820 workstation and HP Z800 workgroup server, both running 64-bit Windows 10 OS. The detailed hardware specifications can be found in Table 1.

	Processor	RAM	NVIDIA GPU
HP Z820	Single 4-core Xeon E5 @ 3.3GHz	64GB	Quadro K4000
HP Z800	Dual 6-core Xeon X5680 @ 3.3GHz	96GB	Quadro 4000

Table 1: In-house benchmarked hardware

For problems that exceed the limits of our in-house hardware, TEN TECH LLC utilizes the [Rescale ScaleX](#) cloud-based HPC resource. This provides our engineers with an ability to run Abaqus on a large cluster available on demand. An example of our use of Abaqus on Rescale's cluster can be found at <https://www.tentechllc.com/resources/Rescale-Abaqus-Case-Study.pdf>. Apart from the abundance of cores and RAM available, Rescale's "Tungsten" configuration also provides access to high-end NVIDIA Tesla GPU accelerators. Thanks to Rescale, a third hardware configuration employing a 16-core and a 32-core environment augmented by an NVIDIA Tesla K40 were utilized, to provide one more data point.

As we are examining the results, it is important to remember the context of our benchmark: modest-sized models for modest-size engineering groups (with modest IT budget). To that effect, the different price/performance points of the GPU employed need to be compared by Year of Introduction, Single & Double Precision Peak Performance as can be seen in Table 2.

GPU	Year	Cores	SP Perf.	DP Perf.	Street Price
Quadro 4000	2010	256	0.49 TFlops	0.24 TFlops	\$200
Quadro K4000	2013	768	1.24 TFlops	N/A	\$800
Tesla K40	2015	2,880	4.29 TFlops	1.43 TFlops	\$3,500

Table 2: GPU Performance/Price Comparison

3. Benchmark Results

In this section, we will discuss the results of each benchmark case, utilizing the following nomenclature to describe the test cases performed on the SBC model:

- **Case 1:** linear statics analysis
- **Case 2:** non-linear analysis
- **Case 3:** non-linear statics with surface-surface contact

All the test cases were solved using the same version of Abaqus 2016. The in-house tests were performed on Windows 10 while the Rescale cluster is using Linux.

3.1 Case 1: Linear Statics

This case is the simplest and fastest to solve, and is meant to confirm the usefulness of GPU computing with Abaqus 2016, even for the smallest of problems. We will also utilize this model to compare GPU performance between two versions of Abaqus: 2016 and 6-14.

This initial case results, run on our HP Z800, are presented on Figure 3.




Figure 3: Case 1 HP Z800 Benchmark Results

The nominal case, with 6-CPU is a short 900s Wall Clock. Adding a GPU to the mix drops this to 590s, a 34% improvement. Quite remarkable considering the relatively modest performance expected out of a 2010 GPU. Equally remarkable is the performance gain between Abaqus 6-14 and 2016 versions of around 20%. This alone should encourage everyone to upgrade to Abaqus 2016.

On such a small problem, it seems more advantageous to run the model with a lower count of cores but with GPU as our example shows 6-core with 1GPU solves slightly faster than with a 12-core configuration, while 12-core still brings benefit over 6-core, solving 30% faster.

3.2 Case 2: Non-linear Statics

The second test case being a more complicated one to solve leads to longer solution time and therefore should provide more granular information with regards to GPU computing benefits.

The in-house results of our test, presenting 4-core and 8-core configurations, with and without GPU are presented on Figure 4.




Figure 4: Case 2 HP Z820 Benchmark Results

This second set of results doesn't seem to be as positive as the first series of benchmark. Somewhat inexplicably, a mere 7% performance increase over a 4-core solve can be obtained while an 8-core configuration is 17% faster. But, as the number of cores increases, the benefit of GPU also increases as the GPU performance gain doubles as we double the number of cores: 14% performance increase over an 8-core solve if we add one GPU. This set of results could be greatly attributed to the fact that the Quadro K4000 does not offer full Double Precision performance.

To further study that effect, the same model was solved on the Linux cluster environment provided by Rescale, which includes a high-end Tesla K40, which offers 1.43TFlops of Double Precision performance. Results of these test cases can be found on Figure 5.




Figure 5: Case 2 HPC Benchmark Results

The use of GPU hardware offering full Double Precision acceleration seems to bring the results back in line with the observations of the first benchmark: an 8-core + 1GPU configuration is 41% faster than an 8-core, and 8% faster than 16-core, while a 16-core configuration is 36% faster than its 8-core counterpart.

Once more, it seems that for a medium-size problem like Case 2, GPU Computing performs very well and is more advantageous than doubling the number of cores, provided the GPU hardware can provide full Double Precision acceleration.

3.3 Case 3: Non-linear Statics with Contact

With Case 3, we introduce surface-surface contact, which should increase the solution time and provide very definitive data points with regards to GPU benefits. As our Quadro K4000 was lacking Double Precision, this series of test will utilize our in-house HP Z800.

On Figure 6, we present the benchmark results for 6-core and 12-core solutions, with or without GPU for our HP Z800 workgroup server.




Figure 6: Case 3 HP Z800 Benchmark Results

The same trend can be observed than for Case 1 & 2: 6-core + 1GPU performs 47% better than 6-core and 11% than 12 core. Pushing the study a little further, adding GPU to a 12-core solve still brings another 26% performance gain. In the case of a 6-CPU solve, the addition of 1 token for GPU computing saves us over 3h30min of solution time. And this is with using a \$200 street price 2010 GPU!

4. Conclusions

While this set of benchmarks is by all means not comprehensive, certain trends can be captured to help making more educated software and hardware decisions when it comes to taking full advantage of the computing power and solver technology available:

- Significant benefits can be obtained, even with the smallest problems and older hardware,
- Abaqus 2016 GPU computing performs around 20% better than previous release (6-14),
- Adding one GPU is roughly equivalent to doubling the core count,
- GPU Computing still performs in HPC environment

While GPU computing consumes less tokens than adding extra cores, its inclusion in everyday Abaqus analyses will allow for greater analysis throughput and better utilization of hardware resources.

2016 Science in the Age of Experience

<http://www.3ds.com/events/science-in-the-age-of-experience>

5. Acknowledgements

TEN TECH LLC would like to acknowledge the following companies and individuals for their contribution in performing these test cases by providing guidance, technical data and hardware resources:

- Rescale | Cloud HPC Simulation
 - Sarah Dietz, Business Development Manager
 - Mulyanto Poort, VP of Engineering
- NVIDIA
 - Baskar Rajagopalan, Sr. Alliance & Marketing Manager for Accelerated Computing