



Faster Innovation - Accelerating SIMULIA Abaqus Simulations with NVIDIA GPUs

Baskar Rajagopalan
Accelerated Computing, NVIDIA



nVIDIA[®]

AGENDA

- Engineering & IT Challenges/Trends
- NVIDIA GPU Solutions
- Abaqus GPU Computing
- Power Consumption
- Which GPUs & Systems for CAE ?
- Remarks
- Q & A

ENGINEERING & IT CHALLENGES/TRENDS

Simulation Turn-around Time

Short Compute Time

Access Anytime, Anywhere

IT Infrastructure

Centralize Compute & Storage Assets

Faster Deployment

Lower Total Cost of Ownership

IP Protection

AGENDA

- Engineering & IT Challenges/Trends
- **NVIDIA GPU Solutions**
- Abaqus GPU Computing
- Power Consumption
- Which GPUs & Systems for CAE ?
- Remarks
- Q & A

NVIDIA GPU SOLUTIONS

Visualization, Accelerated Computing & Virtualization



QUADRO

Revolutionizing Design & Visualization



TESLA

Accelerating Momentum in HPC and Big Data Analytics




GRID

Enabling End-to-End Enterprise Virtualization


HOW DO GPUS BENEFIT SIMULATIONS ?

Efficient handling of parallel tasks in matrix solutions

CPU
Optimized for
Serial Tasks



GPU Accelerator
Optimized for
Parallel Tasks



2-3X

application speed-up when
paired up with a CPU

WHY TESLA GPU ?

Powerful Accelerators



Performance

Faster DP Performance

Larger Memory

Reduces CUDA Kernel Overload



HPC Solutions

Longer Life-Cycle

Form Factor

Cluster Management

GPU Monitoring



Data Reliability

ECC Protection

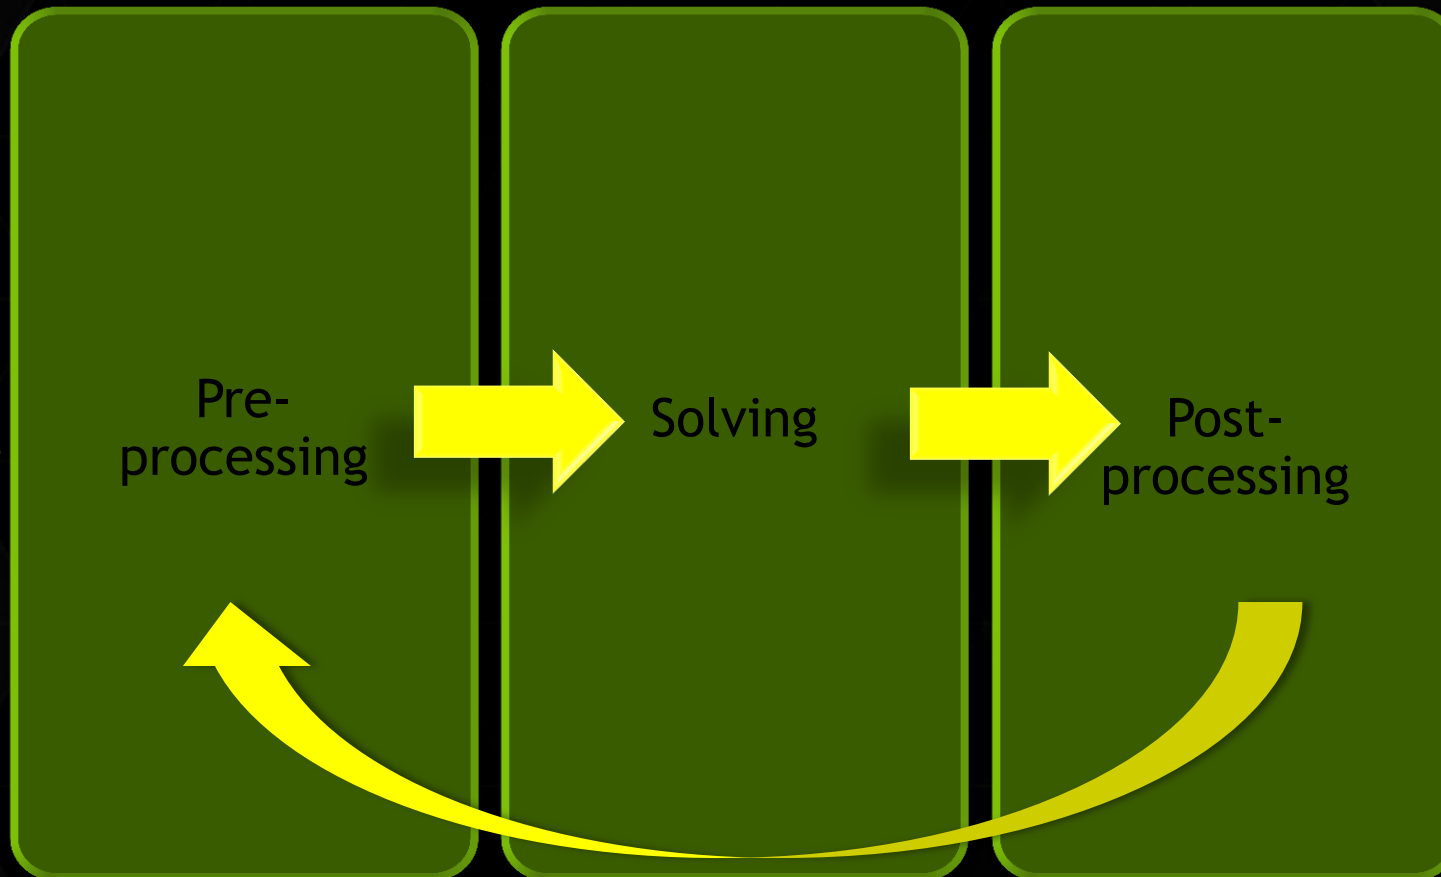

Zero-error Stress Tested

ISV Certification

NVIDIA Support

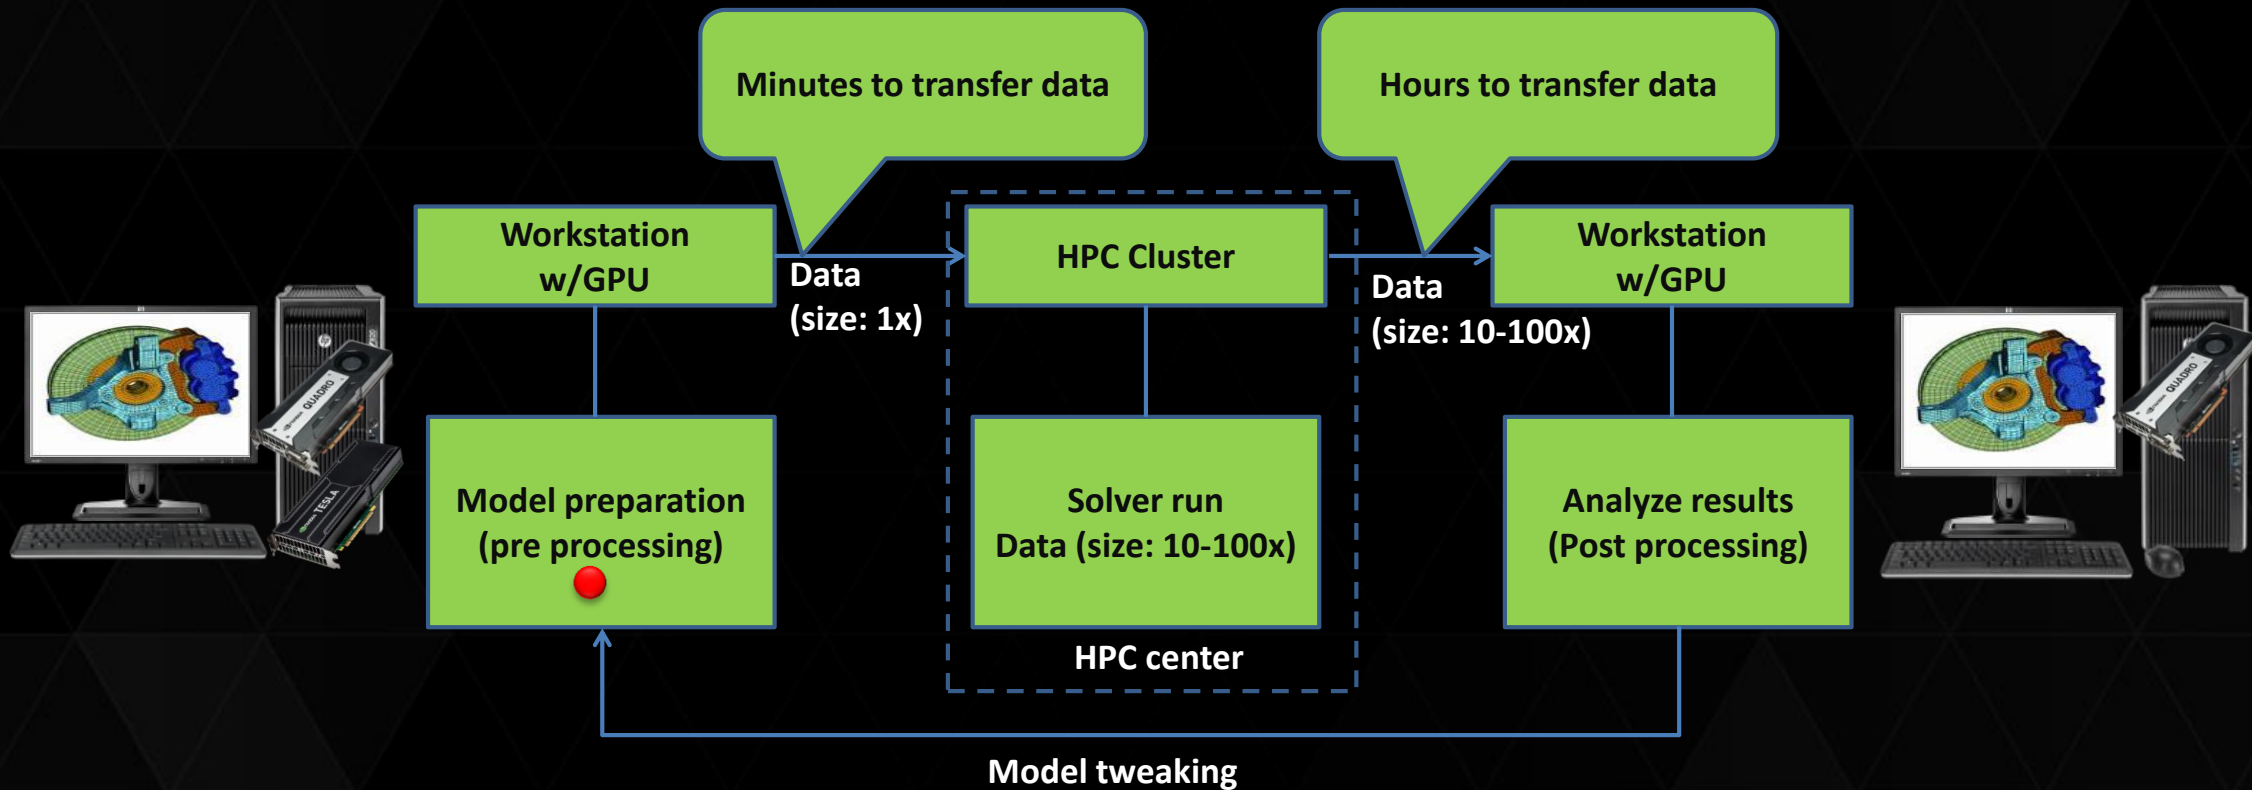
CAE WORKFLOW

Workstation Computing



CAE WORKFLOW

Workstation + Server/Cluster Computing




Traditional workflow

DASSAULT SYSTÈMES & NVIDIA GRID

Visualizing complex industrial design from a continent away

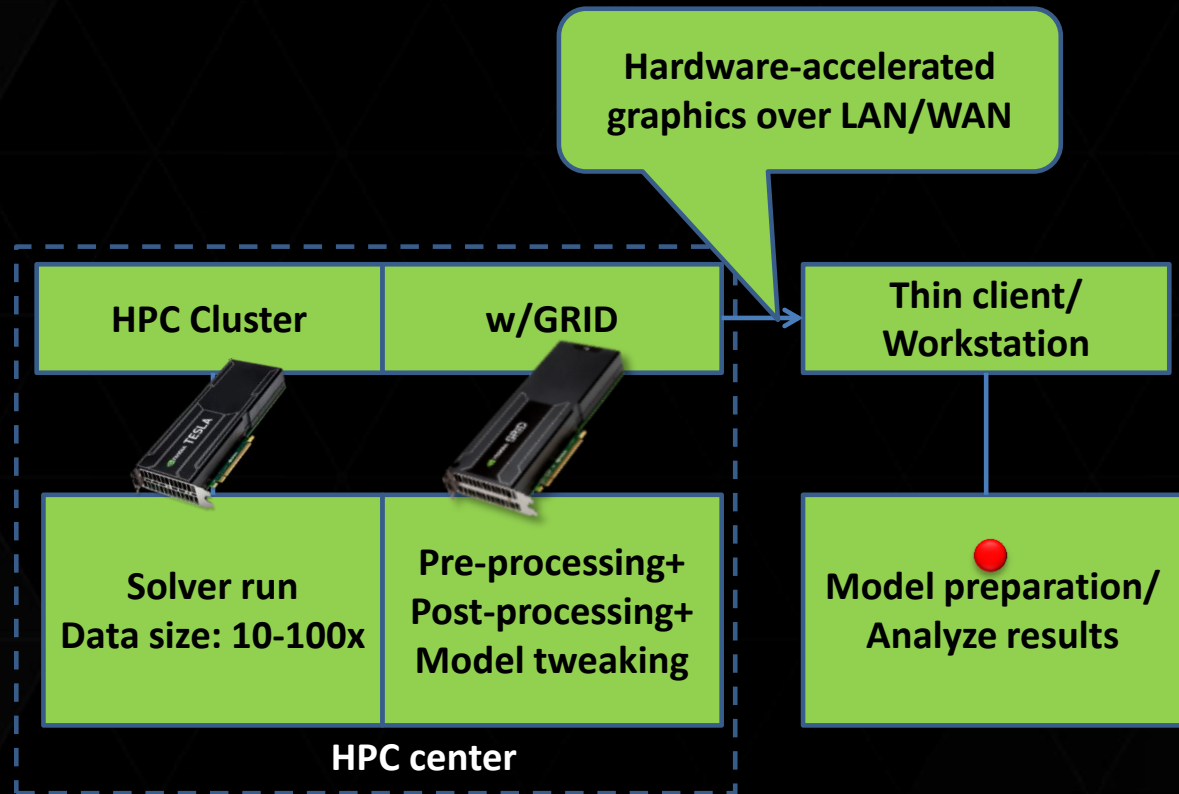
- ▶ Showcase of a Proof Of Concept at SIGGRAPH 2014 with NVIDIA
- ▶ Remote graphics from DS Cloud with GRID K2 with H264 HW encoding
- ▶ Real-time interactive crash test visualization with no data transport integrated into the 3DEXPERIENCE platform



<http://blogs.nvidia.com/blog/2014/08/21/visualizing-complex-industrial-design-from-a-continent-away/#sthash.NEMQyf6F.dpuf>

CAE WORKFLOW WITH NVIDIA GRID

Remote Client + Server Computing



AGENDA

- Engineering & IT Challenges/Trends
- NVIDIA GPU Solutions
- **Abaqus GPU Computing**
- Power Consumption
- Which GPUs & Systems for CAE ?
- Remarks
- Q & A

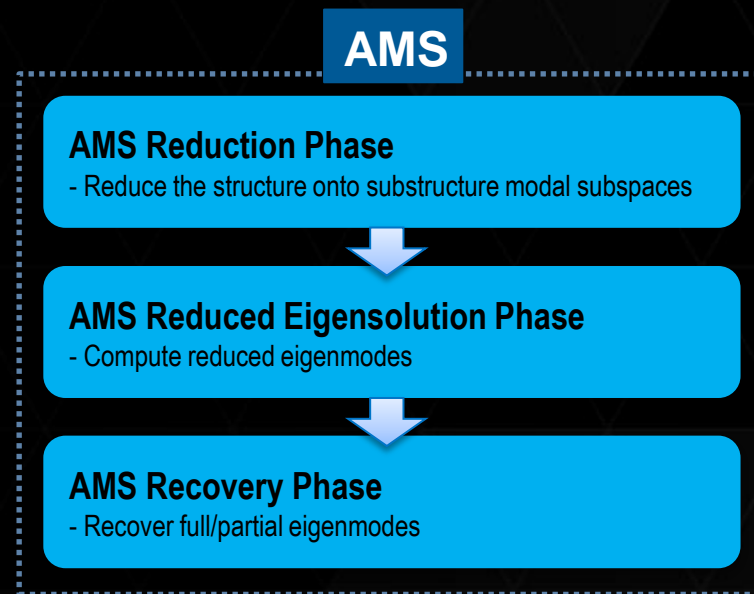
ABAQUS/STANDARD GPU COMPUTING

| | 2011 | 2012 | 2013 | 2014 |
|-------------------------|---|--|-----------------------------------|--|
| Abaqus Release | 6.11 | 6.12 | 6.13 | 6.14 |
| GPU Acceleration | <p>Direct Sparse solver</p> <p>Single GPU</p> | <p>Multi-GPU/node</p> <p>Multi-node DMP clusters</p> <p>Flexibility to run jobs on specific GPUs</p> | <p>Un-symmetric Sparse Solver</p> | <p>Direct Sparse solver</p> <p>DMP Split, less memory</p> <p>AMS Solver</p> <p>Reduced Eigen Phase</p> |
| GPUs | Tesla 20-series Quadro 6000 | Fermi + Kepler Hotfix | Tesla K20/K20X/K40 | Tesla K20/K20X/K40/K80 |

ABAQUS/STANDARD GPU COMPUTING

- ▶ Abaqus 6.14, July 2014
 - ▶ Direct Sparse Solver
 - ▶ Relaxation of memory requirements for GPU
 - ▶ Improved performance / DMP split
 - ▶ AMS eigensolver
 - ▶ Reduced eigen solution phase
- ▶ Abaqus 2016, Fall 2015
 - ▶ AMS: Reduction Phase
 - ▶ Mode-based steady state dynamics

AMS: Automatic Multi-level Substructuring



ABAQUS 6.14 GPU SUPPORT

Supported & recommended features

- ▶ General, Linear, and Nonlinear Analyses
 - ▶ Static Stress & Displacement
 - ▶ Dynamic Stress & Displacement
 - ▶ Heat transfer (steady-state & transient)
 - ▶ Multi-Physics
 - ▶ Thermo-electrical-structural
 - ▶ Pore-fluid flow-mechanical-thermal
- ▶ Linear Perturbation Analysis
 - ▶ Static Stress & Displacement
 - ▶ Linear Static
 - ▶ Dynamic Stress & Displacement
 - ▶ Steady-state dynamics (direct)

ABAQUS 6.14 GPU SUPPORT

Supported & recommended features

▶ Solution Techniques

- ▶ Parallel execution on both shared memory & distributed memory parallel (cluster) systems
- ▶ Parallel direct sparse solver with dynamic load balancing
- ▶ Parallel AMS eigenvalue solution
- ▶ GPGPU-accelerated sparse solver

▶ Abaqus/AMS

- ▶ High-performance automatic multi-level substructuring eigensolver

ABAQUS GPU LICENSING

A single GPU board is treated as one core for token count

- ▶ 5 tokens unlocks a single CPU core
- ▶ 1 additional token unlocks additional CPU core OR unlocks 1 entire GPU
- ▶ GPUs help in reducing consumption of licensing tokens

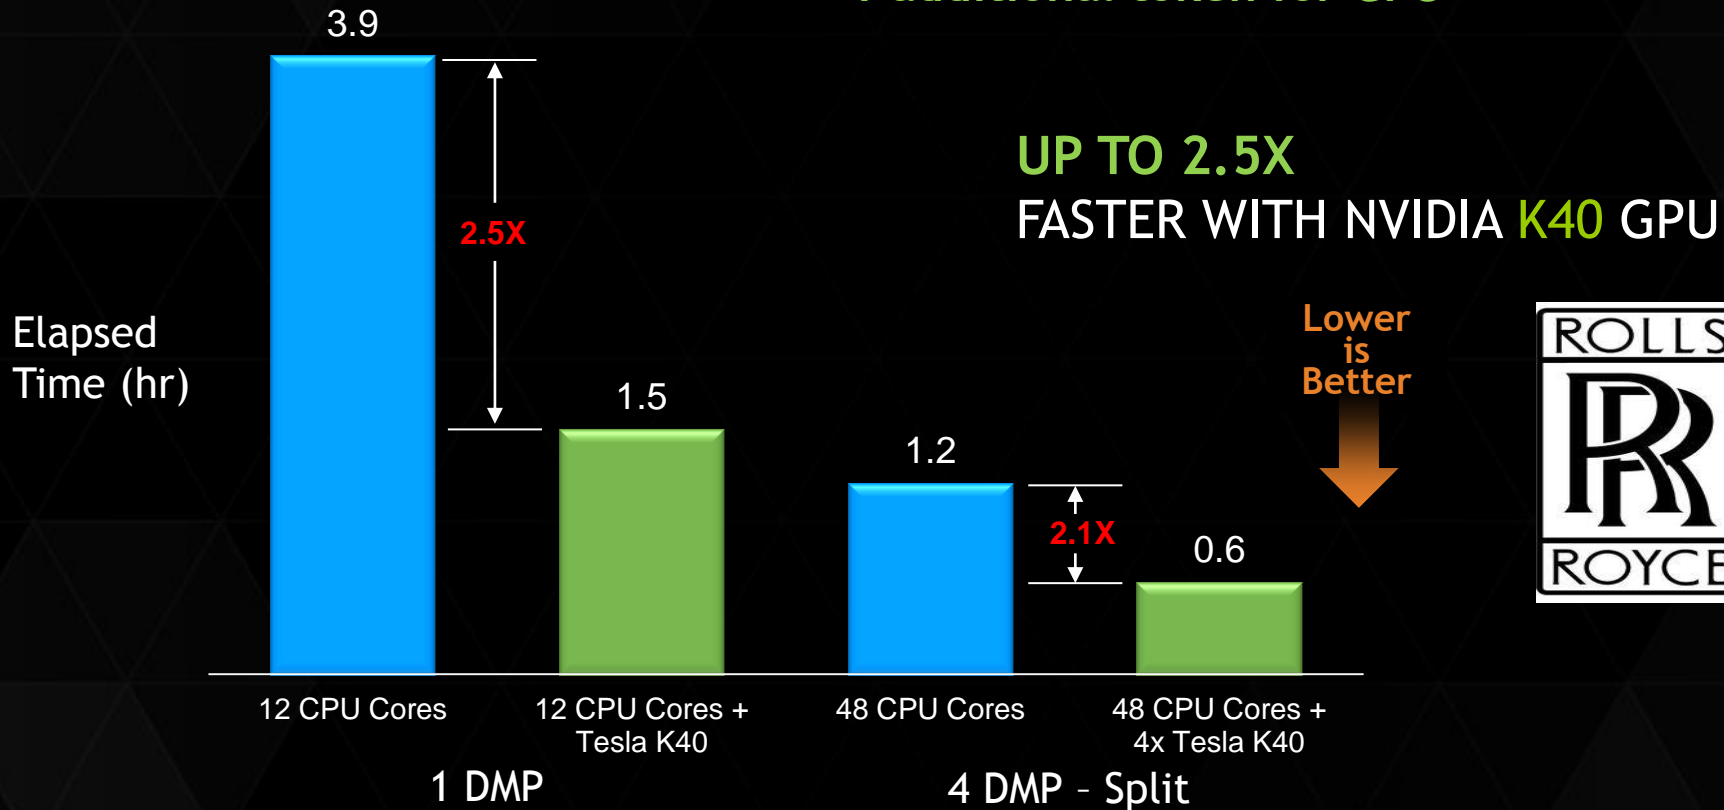
| Cores | CPU Tokens* | GPU | GPU Tokens* (1) | GPU | GPU Tokens* (2) |
|-------------|-------------|-----|-----------------|-----|-----------------|
| 1 | 5 | 1 | 6 | 2 | 7 |
| 2 | 6 | 1 | 7 | 2 | 8 |
| 3 | 7 | 1 | 8 | 2 | 9 |
| 4 | 8 | 1 | 9 | 2 | 10 |
| 5 | 9 | 1 | 10 | 2 | 11 |
| 6 | 10 | 1 | 11 | 2 | 12 |
| 7 | 11 | 1 | 12 | 2 | 12 |
| 8 (1 CPU) | 12 | 1 | 12 | 2 | 13 |
| 9 | 12 | 1 | 13 | 2 | 13 |
| 10 | 13 | 1 | 13 | 2 | 14 |
| 11 | 13 | 1 | 14 | 2 | 14 |
| 12 | 14 | 1 | 14 | 2 | 15 |
| 13 | 14 | 1 | 15 | 2 | 15 |
| 14 | 15 | 1 | 15 | 2 | 16 |
| 15 | 15 | 1 | 16 | 2 | 16 |
| 16 (2 CPUs) | 16 | 1 | 16 | 2 | 16 |

* # of Tokens = INT(5*cores^0.422)

ABAQUS PERFORMANCE WITH GPU


No additional tokens for GPU

1 additional token for GPU



Large Model (~77 TFLOPs), 4.71M DOF, Nonlinear Static, Direct Sparse Solver
 Abaqus 6.14-2 with Intel Xeon E5-2697v2, 2.70 GHz CPU, 128 GB memory; Tesla K40 GPU


ABAQUS PERFORMANCE WITH GPU



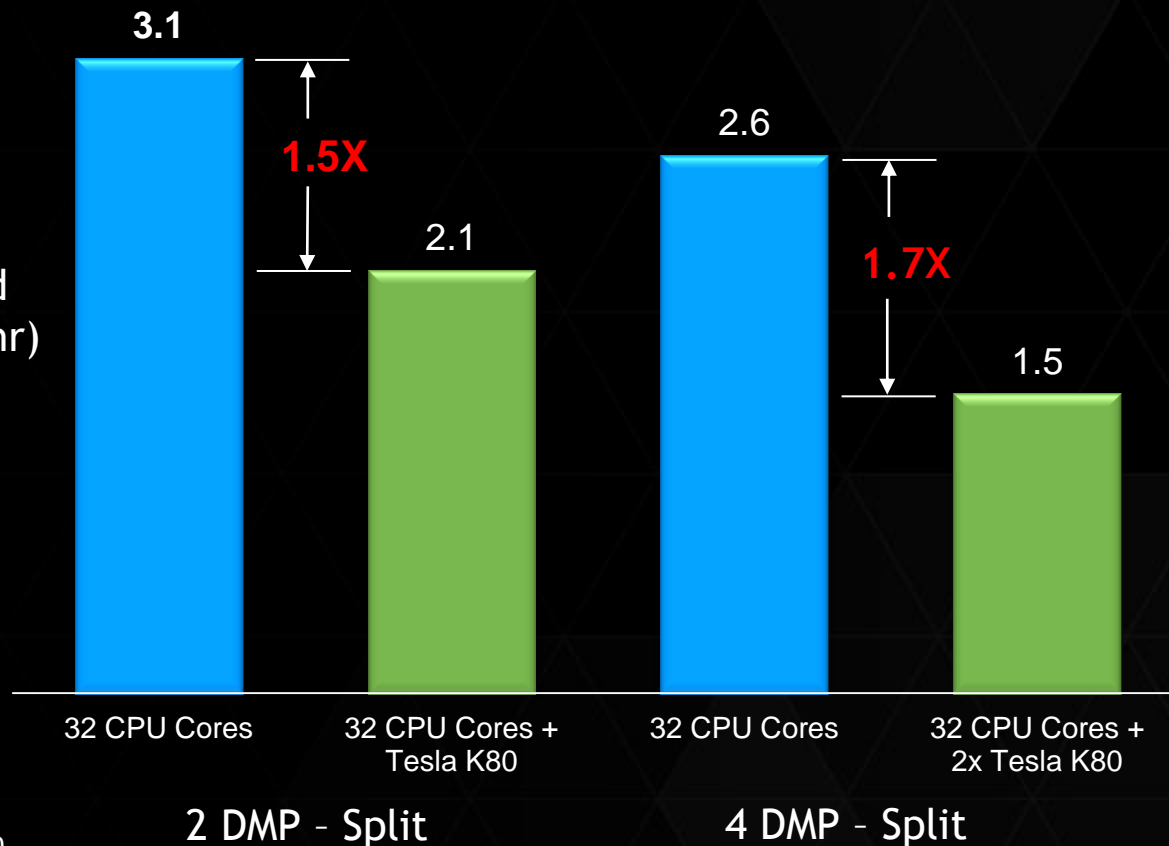
Large Model (~77 TFLOPs), 4.71M DOF, Nonlinear Static, Direct Sparse Solver
 Abaqus 6.14-2 with Intel Xeon E5-2697v2, 2.70 GHz CPU, 128 GB memory; Tesla GPUs

ABAQUS PERFORMANCE WITH GPU

1.7X FASTER WITH NVIDIA K80 GPU's



Elapsed Time (hr)




Model: s4e; 16.7 MDOF, Nonlinear Static, Direct Sparse Solver; Abaqus 6.14 with Intel Xeon Haswell E5-2698v3 (16-core), 2.3 GHz CPU, 256 GB memory; Tesla K80

ABAQUS PERFORMANCE WITH GPU

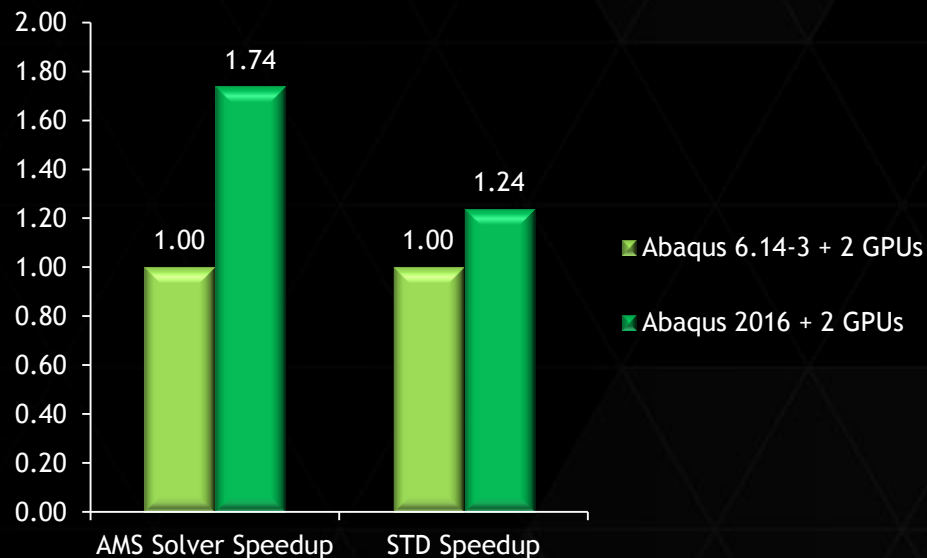
Abaqus/AMS 2016 solver

1.5x faster than v6.14



20M DOF and 12k modes

1.2x faster than v6.14




3M DOF and 5k modes

HP ProLiant SL250s Gen8, Intel Xeon Ivy Bridge E5-2680e (2x 10 cores), 2.8 GHz, 192 GB Memory and 2x Tesla K40m GPUs

ABAQUS PERFORMANCE WITH GPU

2 additional tokens for all GPUs



| Node 0 | Node 1 | Node 2 | Node 3 | Node 4 | Node 5 |
|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|
| 16 core Intel Sandybridge | 16 core Intel Sandybridge | 16 core Intel Sandybridge | 16 core Intel Sandybridge | 16 core Intel Sandybridge | 16 core Intel Sandybridge |
| 2 x NVIDIA K20m (Kepler) | 2 x NVIDIA K20m (Kepler) | 2 x NVIDIA K20m (Kepler) | 2 x NVIDIA K20m (Kepler) | 2 x NVIDIA K20m (Kepler) | 2 x NVIDIA K20m (Kepler) |
| 256 GB | 256 GB | 256 GB | 256 GB | 256 GB | 256 GB |

- ▶ 6 compute nodes
- ▶ 2 MPI processes per compute node
- ▶ Accelerated DMP execution mode (an optional feature in 6.14)

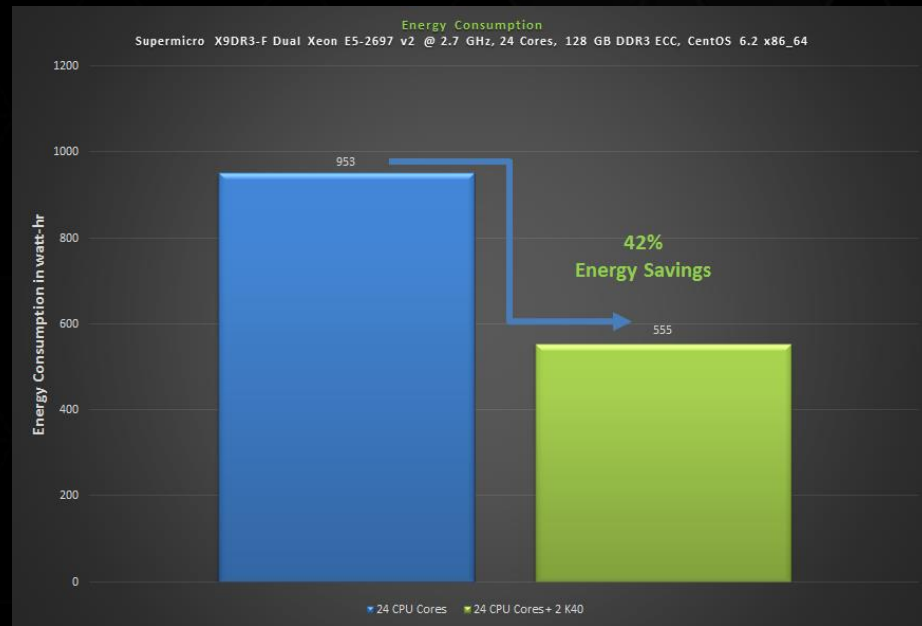
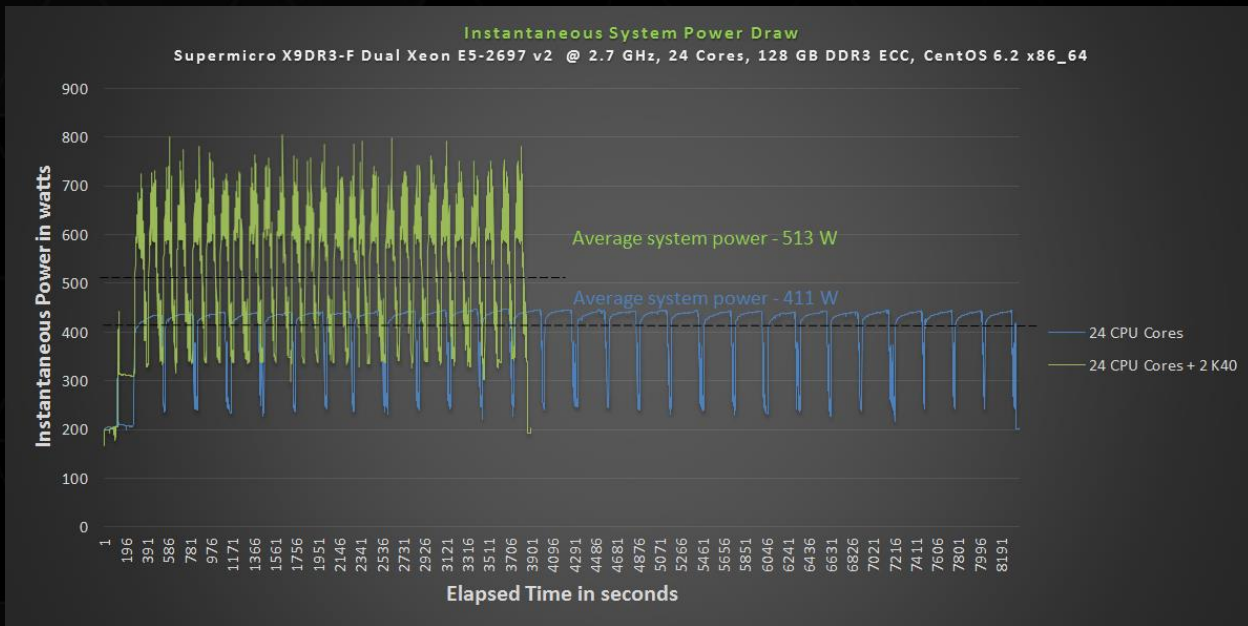
AGENDA

- Engineering & IT Challenges/Trends
- NVIDIA GPU Solutions
- Abaqus GPU Computing
- **Power Consumption**
- Which GPUs & Systems for CAE ?
- Remarks
- Q & A

POWER CONSUMPTION STUDY

Abaqus 6.14

No additional tokens for 1 or 2 GPUs



- ▶ Adding GPUs to a CPU-only node resulted in a **2.2x** speed-up while reducing energy consumption by **42%**

AGENDA

- Engineering & IT Challenges/Trends
- NVIDIA GPU Solutions
- Abaqus GPU Computing
- Power Consumption
- Which GPUs & Systems for CAE ?
- Remarks
- Q & A

NVIDIA GPU FOR CAE

| | Workstations | Clusters / Servers |
|----------------------|--------------------------------|---|
| Visualization | Quadro K-Series | Tesla K20X*, K40, GRID K2 |
| Computing | Tesla K20, K40 Quadro K6000 | Tesla K20, K20X, K40, K80 |
| Remote Visualization | Quadro K6000 | Tesla K20*, K20X*, K40, K80, GRID K2 |
| Virtualization | GRID K2, K6000 | GRID K2, K6000 |



* Passively cooled GPUs only; GOM(Graphics-Only Mode) needs to be enabled

HP SYSTEMS WITH NVIDIA GPU FOR COMPUTING



HP ProLiant Gen9 servers

HP Apollo 2000

Scalable Multi-node



HP ProLiant XL190r
2x NVIDIA K40

HP Apollo 6000

Rack-scale Efficiency



HP ProLiant XL250a
2P + 2x NVIDIA Tesla K40 or K80

HP Apollo 8000

Performance Density



HP ProLiant XL750f
2P + 2 NVIDIA Tesla K40 XL (K40d)

HP COMPUTERS WITH NVIDIA GPU FOR GRAPHICS



HP workstation-class graphics

HP Desktop Workstations


High-end graphics & Computing



HP Z840
Up to 2x NVIDIA Tesla K40

HP Graphics Server Blade

For Client Virtualization



HP ProLiant WS460c Gen9
up to 6x NVIDIA Quadro K3100M

HP Graphics Server Blade with Expansion

For Client Virtualization






HP ProLiant WS460c Gen9
NVIDIA Quadro K6000/K5000/K4000,
GRID K2/K1

AGENDA

- Engineering & IT Challenges/Trends
- NVIDIA GPU Solutions
- Abaqus GPU Computing
- Power Consumption
- Which GPUs & Systems for CAE ?
- Remarks
- Q & A

ABAQUS GPU COMPUTING

- ▶ GPU Support for Abaqus/Standard since 2011 (v6.11)
- ▶ Current supported version: Abaqus 6.14 Refresh3
- ▶ Broad range of analysis types
- ▶ Multiple and selective GPU support
- ▶ Multi-GPU/node; multi-node DMP clusters
- ▶ Abaqus/AMS
- ▶ Abaqus GPU licensing based on tokens
 - ▶ Fewer consumption of tokens
- ▶ Performance gains vary
 - ▶ 2-3x speed-ups are common with large, solid models



NVIDIA BENEFITS FOR ABAQUS USERS




Increased
Throughput with
Faster Simulation
Runs



Fewer Simulation
Runs for Solution
Convergence



Improved
Team/Supplier
Collaboration



Move Simulation
Early in Design
Cycle

ABAQUS TEST DRIVE

www.accelerateabaqusongpu.com



- ▶ HP ProLiant SL250s Gen8 Server & NVIDIA Tesla GPUs
- ▶ 16 cores (2x 8-core E5-2600 Sandy Bridge), 128GB, 2x NVIDIA K20

Register

After registering, you will receive login instructions to access the cluster.

First Name * Last Name *

Company * Phone

Email * Confirm Email Address *

Abaqus site ID *


Workload input file¹ runtime on your current system: (Max. 48:00 Hours:minutes)
Hours * Minutes *

How many cores in your current system? *

What system and processor platform are you currently running? * If other, please specify.

J S S 2

1. Customer data files are maintained solely by SIMULIA and deleted immediately after test results are produced.



Thank you
Q & A

brajagopalan@nvidia.com
sharon.shaw@hp.com