NVIDIA Experiences with Porting Large-Scale Engineering Codes to GPUs

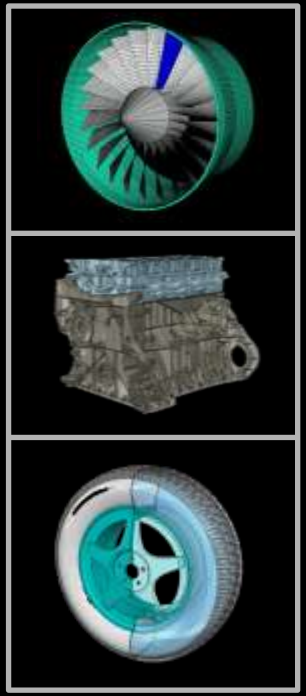

Stan Posey, HPC Industry Development
NVIDIA, Santa Clara, CA, USA
sposey@nvidia.com

NVIDIA Introduction and HPC Evolution of GPUs

NVIDIA

- Public, based in Santa Clara, CA | ~$4B revenue | ~5,500 employees
- Founded in 1999 with primary business in semiconductor industry
  - Products for graphics in workstations, notebooks, mobile devices, etc.
  - Began R&D of GPUs for HPC in 2004, released first Tesla and CUDA in 2007
- Development of GPUs as a co-processing accelerator for x86 CPUs

HPC Evolution of GPUs

- **2004:** Began strategic investments in GPU as HPC co-processor
- **2006:** G80 first GPU with built-in compute features, 128 cores; CUDA SDK Beta
- **2007:** Tesla 8-series based on G80, 128 cores – CUDA 1.0, 1.1
- **2008:** Tesla 10-series based on GT 200, 240 cores – CUDA 2.0, 2.3
- **2009:** Tesla 20-series, code named “Fermi” up to 512 cores – CUDA SDK 3.0

3 Generations of Tesla in 3 Years

How NVIDIA Tesla GPUs are Deployed in Systems

	Data Center Products		Workstation
	Tesla M205 / M2070 Adapter	Tesla S2050 1U System	Tesla C2050 / C2070 Workstation Board
GPUs	**1 Tesla GPU**	**4 Tesla GPUs**	**1 Tesla GPU**
Single Precision	1030 Gigaflops	4120 Gigaflops	1030 Gigaflops
Double Precision	515 Gigaflops	2060 Gigaflops	515 Gigaflops
Memory	3 GB / 6 GB	12 GB (3 GB / GPU)	3 GB / 6 GB
Memory B/W	148 GB/s	148 GB/s	144 GB/s

Engineering Disciplines and Related Software

- Computational Structural Mechanics (CSM) implicit for strength (stress) and vibration
  - Structural strength at minimum weight, low-frequency oscillatory loading, fatigue
    - ANSYS; ABAQUS/Standard; MSC.Nastran; NX Nastran; Marc

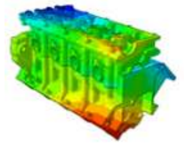

- Computational Structural Mechanics (CSM) explicit for impact loads; structural failure
  - Impact over short duration; contacts – crashworthiness, jet engine blade failure, bird-strike
    - LS-DYNA; ABAQUS/Explicit; PAM-CRASH; RADIOSS

- Computational Fluid Dynamics (CFD) for flow of liquids (~water) and gas (~air)
  - Aerodynamics; propulsion; reacting flows; multiphase; cooling/heat transfer
    - ANSYS FLUENT; STAR-CD; STAR-CCM+; CFD++; ANSYS CFX; AcuSolve; PowerFLOW

- Computational Electromagnetics (CEM) for EM compatibility, interference, radar
  - EMC for sensors, controls, antennas; low observable signatures; radar-cross-section
    - ANSYS HFSS; ANSYS Maxwell; ANSYS SIwave; XFdtd; FEKO; Xpatch; SIGLBC; CARLOS; MM3D

Motivation for CPU Acceleration with GPUs

IDC's Top 10 HPC Market Predictions for 2010

February 17, 2010

6. x86 Processors Will Dominate, But GPGPUs Will Gain Traction as x86 Hits the Wall

- x86 processors went from near-zero to hero in HPC in the past decade, largely replacing RISC.
- x86 will continue to dominate, but GPGPUs will start making their presence felt more in 2010.
- Multiple Large HPC procurements have substantial GPGPU content.
  - GPGPUs play a crucial role in ORNL's planned exascale system.
- GPGPUs provide more peak/Linpack flops per dollar for politics and will inevitably provide more sustained flops for suitable applications.
- In 2010, some ISVs will announce plans to redesign their apps with GPGPUs in mind.

GPU Progress Status for Engineering Codes

GPU Status	Structural Mechanics	Fluid Dynamics	Electromagnetics
Available Today	ANSYS Mechanical AFEA Abaqus/Standard (beta)	AcuSolve Moldflow Culises (OpenFOAM) Particleworks	Nexxim EMPro CST MS XFdtd SEMCAD X
Release Coming in 2011	LS-DYNA *implicit* Marc	CFD++ LS-DYNA CFD	Xpatch
Product Evaluation	RADIOSS *implicit* PAM-CRASH *implicit* MD Nastran NX Nastran	CFD-ACE+ FloEFD Abaqus/CFD	
Research Evaluation	LS-DYNA Abaqus/Explicit	FLUENT/CFX STAR-CCM+	HFSS

GPU Considerations for Engineering Codes

- Initial efforts are linear solvers on GPU, but it's not enough
  - Linear solvers ~50% of profile time -- only 2x speed-up is possible
  - More of application will be moved to GPUs in progressive stages

- Most codes use a parallel domain decomposition method
  - This fits GPU model very well and preserves costly MPI investment

- All codes are parallel and scale across multiple CPU cores
  - Fair GPU vs. CPU comparisons should be CPU-socket-to-GPU-socket
  - Comparisons presented here are made against 4-core Nehalem

Leading ISVs Who Develop Engineering Codes

NVIDIA

ISV	Application
ANSYS	ANSYS CFD (FLUENT and CFX); ANSYS Mechanical; HFSS
SIMULIA	Abaqus/Standard; Abaqus/Explicit
LSTC	LS-DYNA
MSC.Software	MD Nastran; Marc; Adams
CD-adapco	STAR-CD; STAR-CCM+
Altair	RADIOSS
Siemens	NX Nastran
ESI Group	PAM-CRASH; PAM-STAMP
Metacomp	CFD++
ACUSIM	AcuSolve
Autodesk	Moldflow

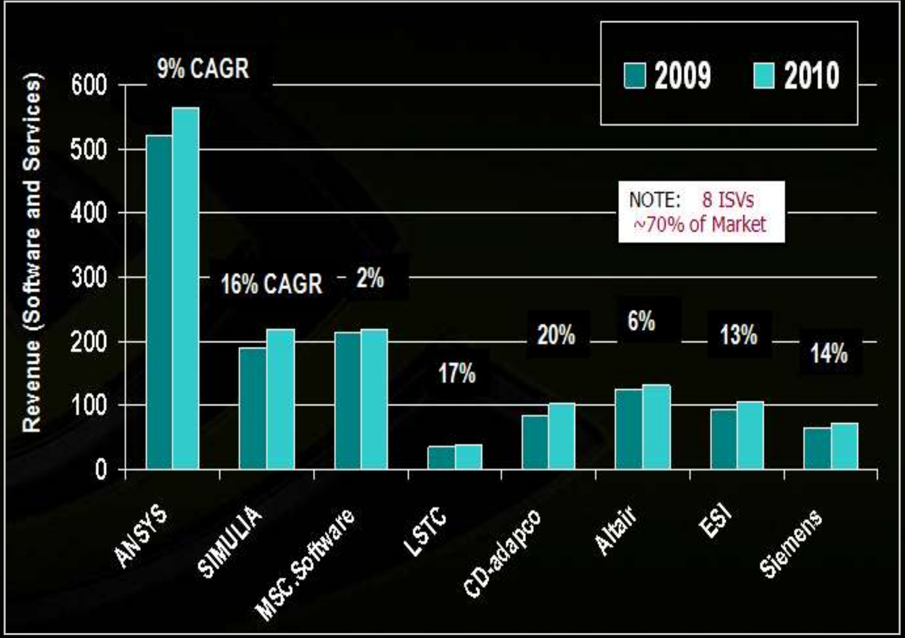

GPU Priority by ISV Market Opportunity and "Fit"

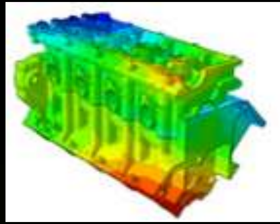

#1 Computational Structural Mechanics (CSM) implicit for strength (stress) and vibration

ANSYS | ABAQUS/Standard | MSC.Nastran; Marc | NX Nastran | LS-DYNA | RADIOSS

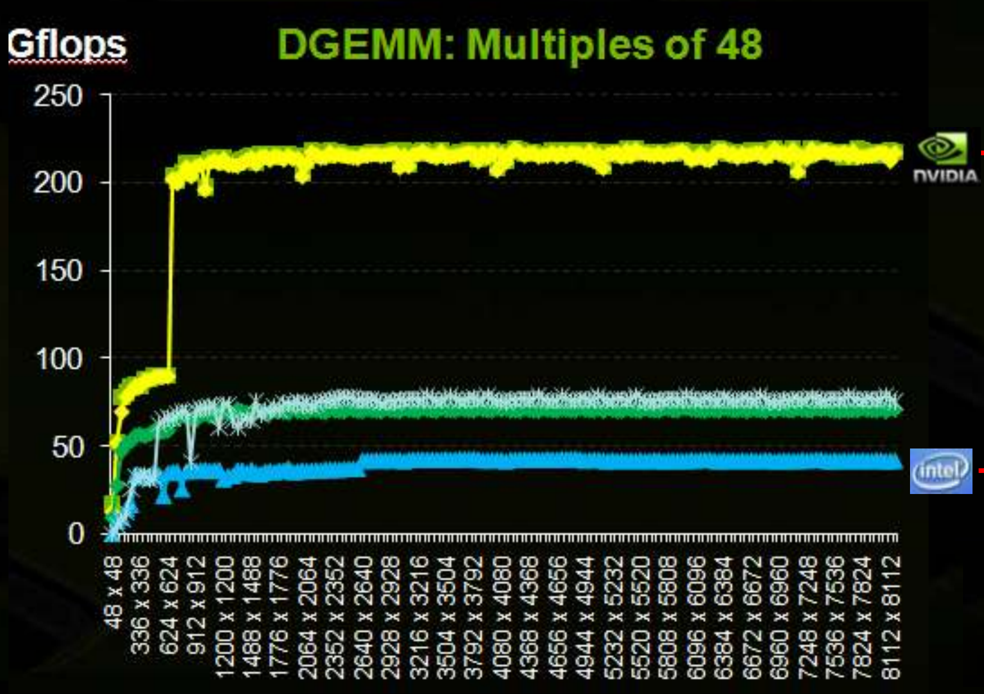

Typical Computational Profiles of CSM Implicit

Direct linear equation solver ~75%

Tesla C2050 4x Faster DGEMM vs. QC Nehalem

Tesla C2050 (ECC off) | Tesla C2050 (ECC on) | Tesla C1060 | MKL 4 Threads | MKL 8 Threads

cuBLAS 3.1: NVIDIA Tesla C1060, Tesla C2050 (Fermi)

MKL 10.2.4.32: Quad-Core Intel Xeon 5550, 2.67 GHz

DGEMM Improved 36% With CUDA 3.2 (Nov 10)

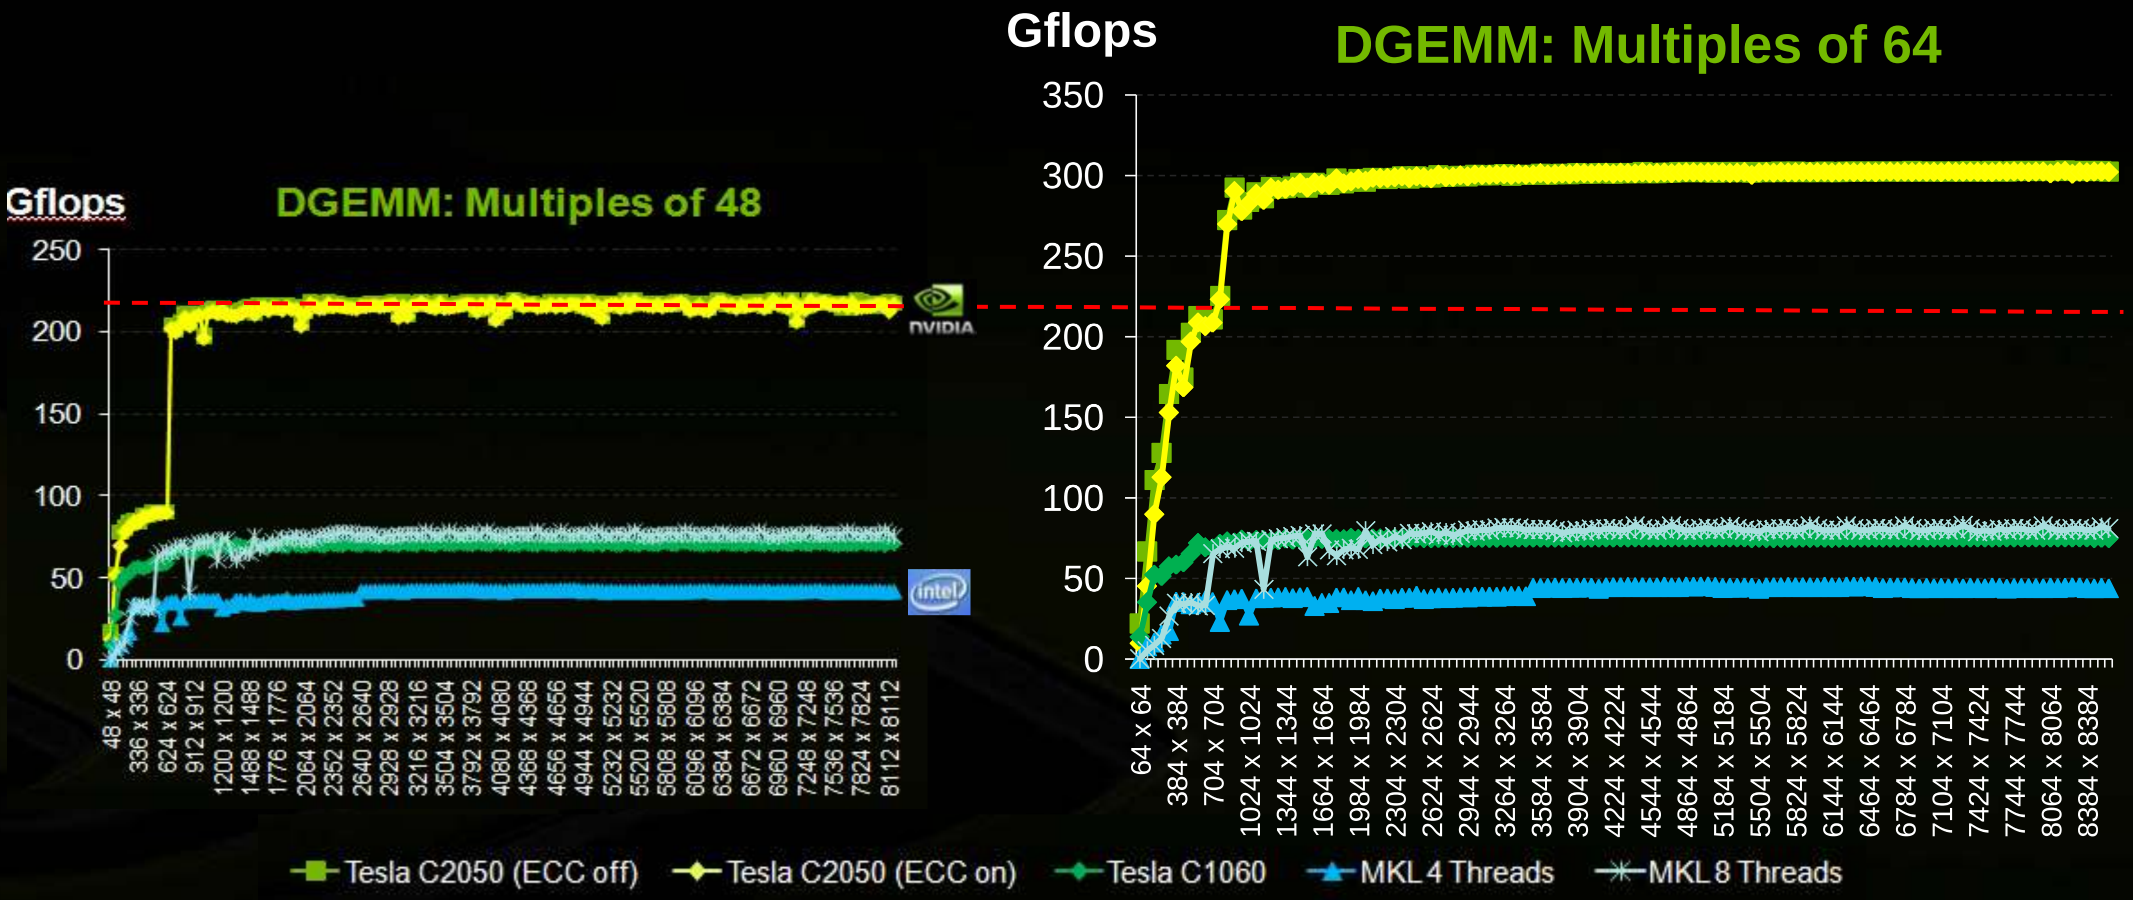

cuBLAS 3.1: NVIDIA Tesla C1060, Tesla C2050 (Fermi)
MKL 10.2.4.32: Quad-Core Intel Xeon 5550, 2.67 GHz

cuBLAS 3.2: NVIDIA Tesla C1060, Tesla C2050 (Fermi)
MKL 10.2.4.32: Quad-Core Intel Xeon 5550, 2.67 GHz

Basics of Implicit CSM Implementations

Implicit CSM – deployment of a multi-frontal direct sparse solver

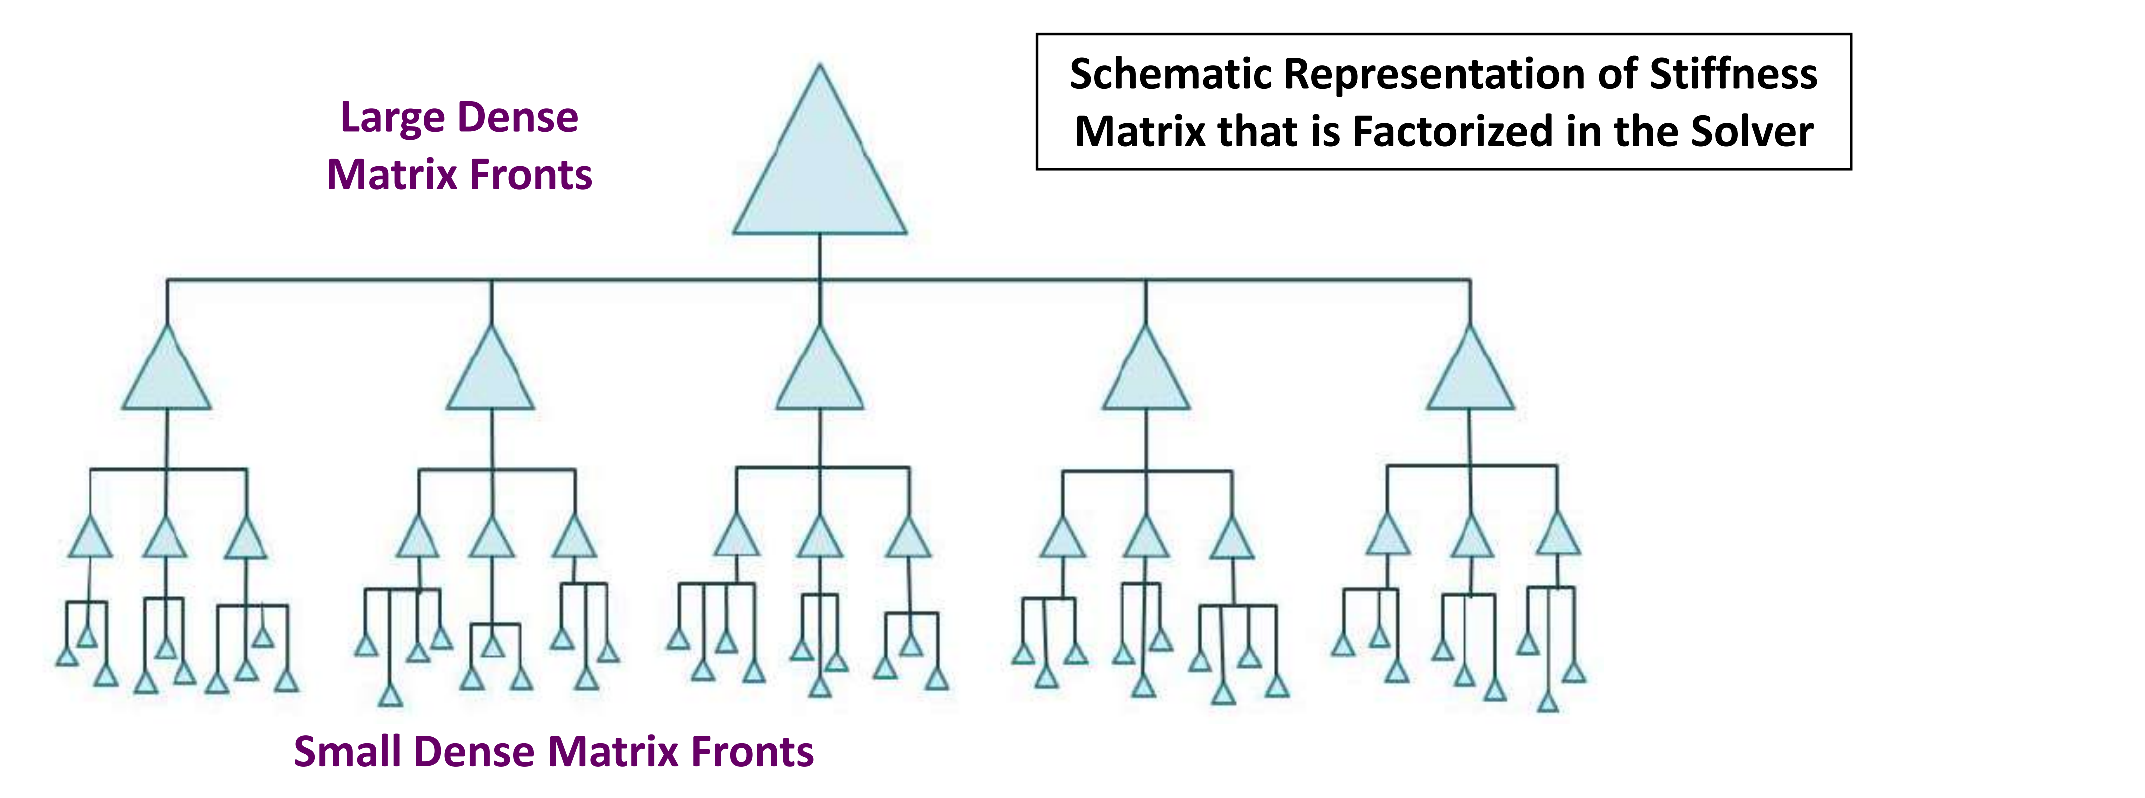

Basics of Implicit CSM Implementations

Implicit CSM – deployment of a multi-frontal direct sparse solver

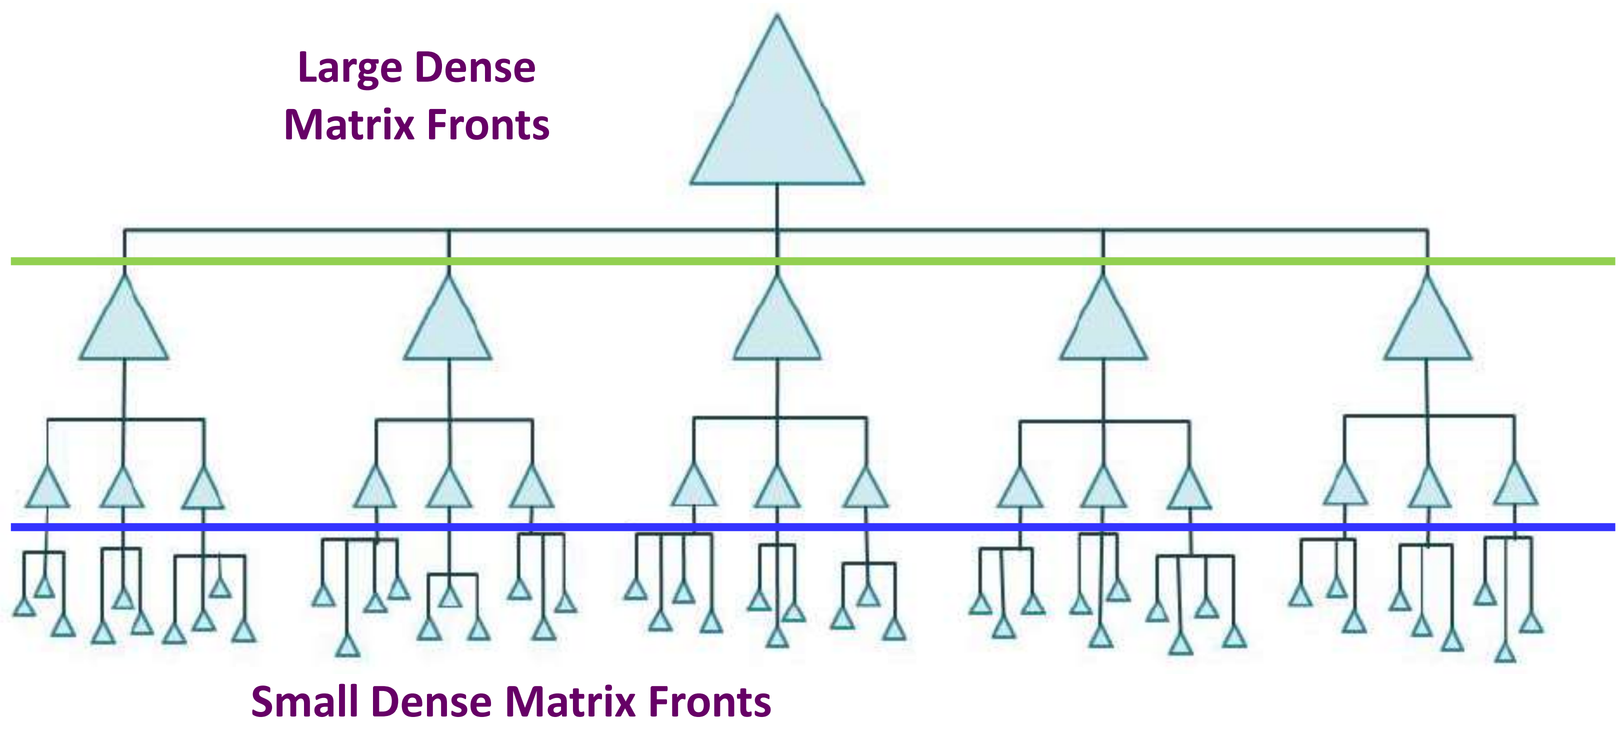

Upper threshold:
Fronts too large for single GPU memory need multiple GPUs

Lower threshold:
Fronts too small to overcome PCIe data transfer costs stay on CPU cores

ANSYS Performance Study by HP and NVIDIA

HP ProLiant SL390 Server Configuration

- Single server node – 12 total CPU cores, 1 GPU
- 2 x Xeon X5650 HC 2.67 GHz CPUs (Westmere)
- 48 GB memory – 12 x 4GB 1333 MHz DIMMs
- NVIDIA Tesla M2050 GPU with 3 GB memory
- RHEL5.4, MKL 10.25, NVIDIA CUDA 3.1 – 256.40
- *Study conducted at HP by Domain Engineering*

HP SL390 Server

+

NVIDIA Tesla M2050 GPU

HP Z800 Workstation Configuration

- 2 x Xeon X5570 QC 2.8 GHz CPUs (Nehalem)
- 48 GB memory
- NVIDIA Tesla C2050 with 3 GB memory
- RHEL5.4, Intel MKL 10.25, NVIDIA CUDA 3.1
- *Study conducted at NVIDIA by Performance Lab*

HP Z800 Workstation

+

NVIDIA Tesla C2050 GPU

ANSYS Mechanical Model – V12sp-5

- Turbine geometry, 2,100 K DOF and SOLID187 FE’s
- Single load step, static, large deflection nonlinear
- ANSYS Mechanical 13.0 direct sparse solver

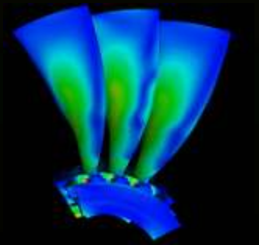

ANSYS Mechanical for Westmere GPU Server

NOTE: Results Based on ANSYS Mechanical R13 SMP Direct Solver Sep 2010

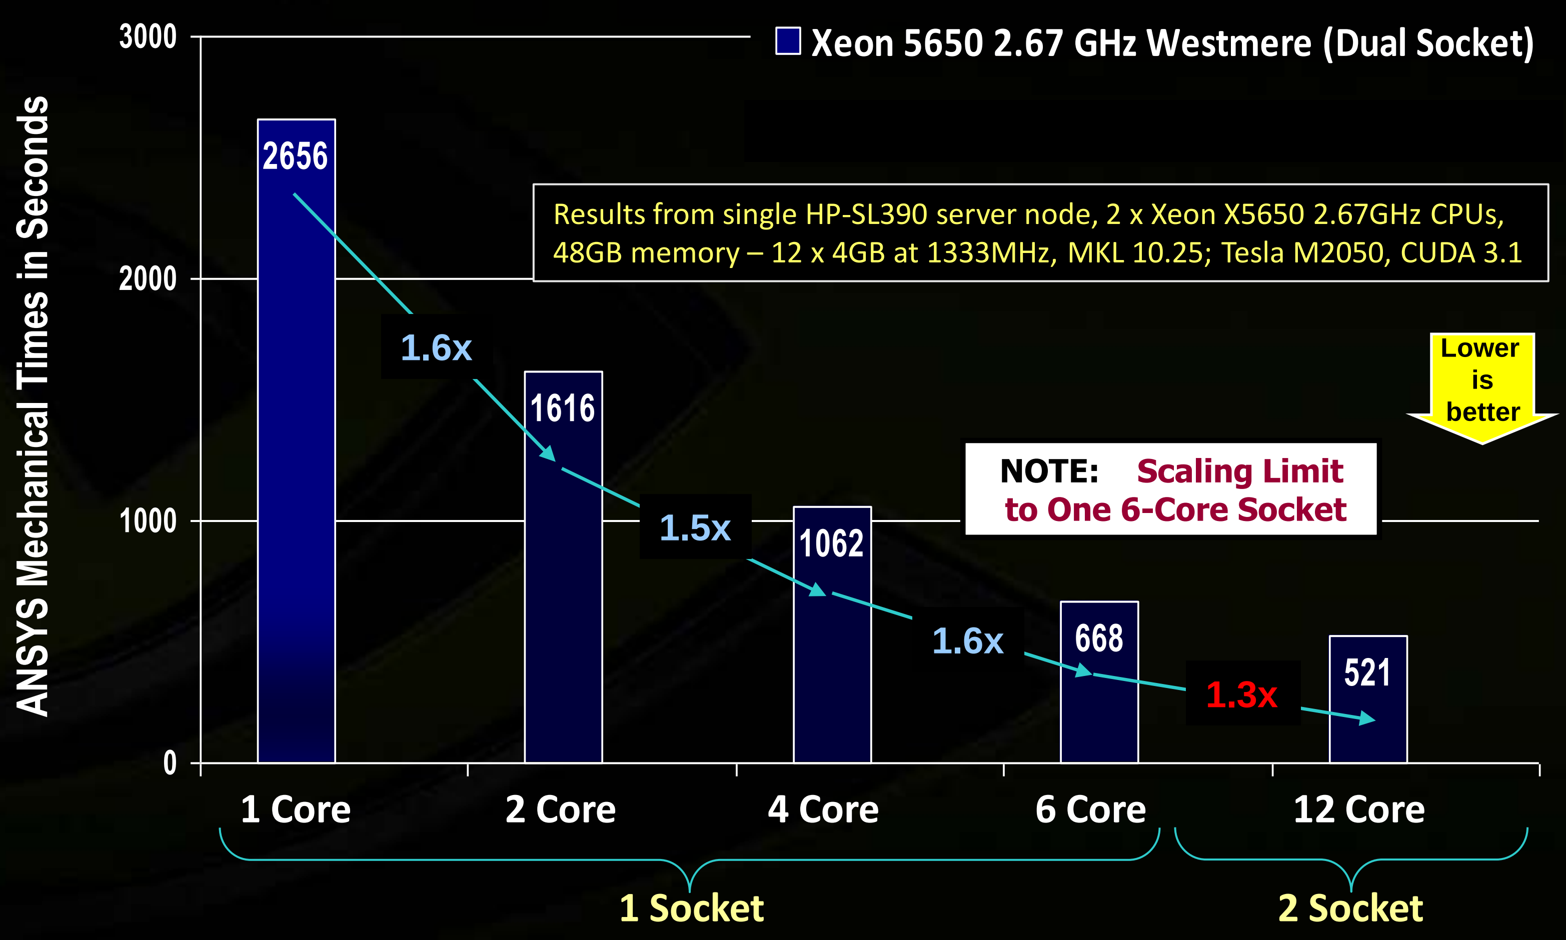

V12sp-5 Model

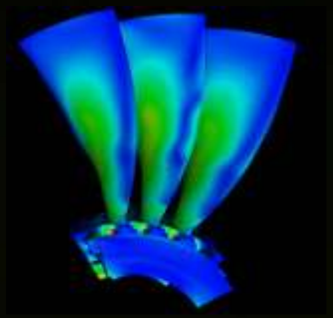

- Turbine geometry
- 2,100 K DOF
- SOLID187 FEs
- Static, nonlinear
- One load step
- Direct sparse

ANSYS Mechanical for Westmere GPU Server

NOTE: Results Based on ANSYS Mechanical R13 SMP Direct Solver Sep 2010

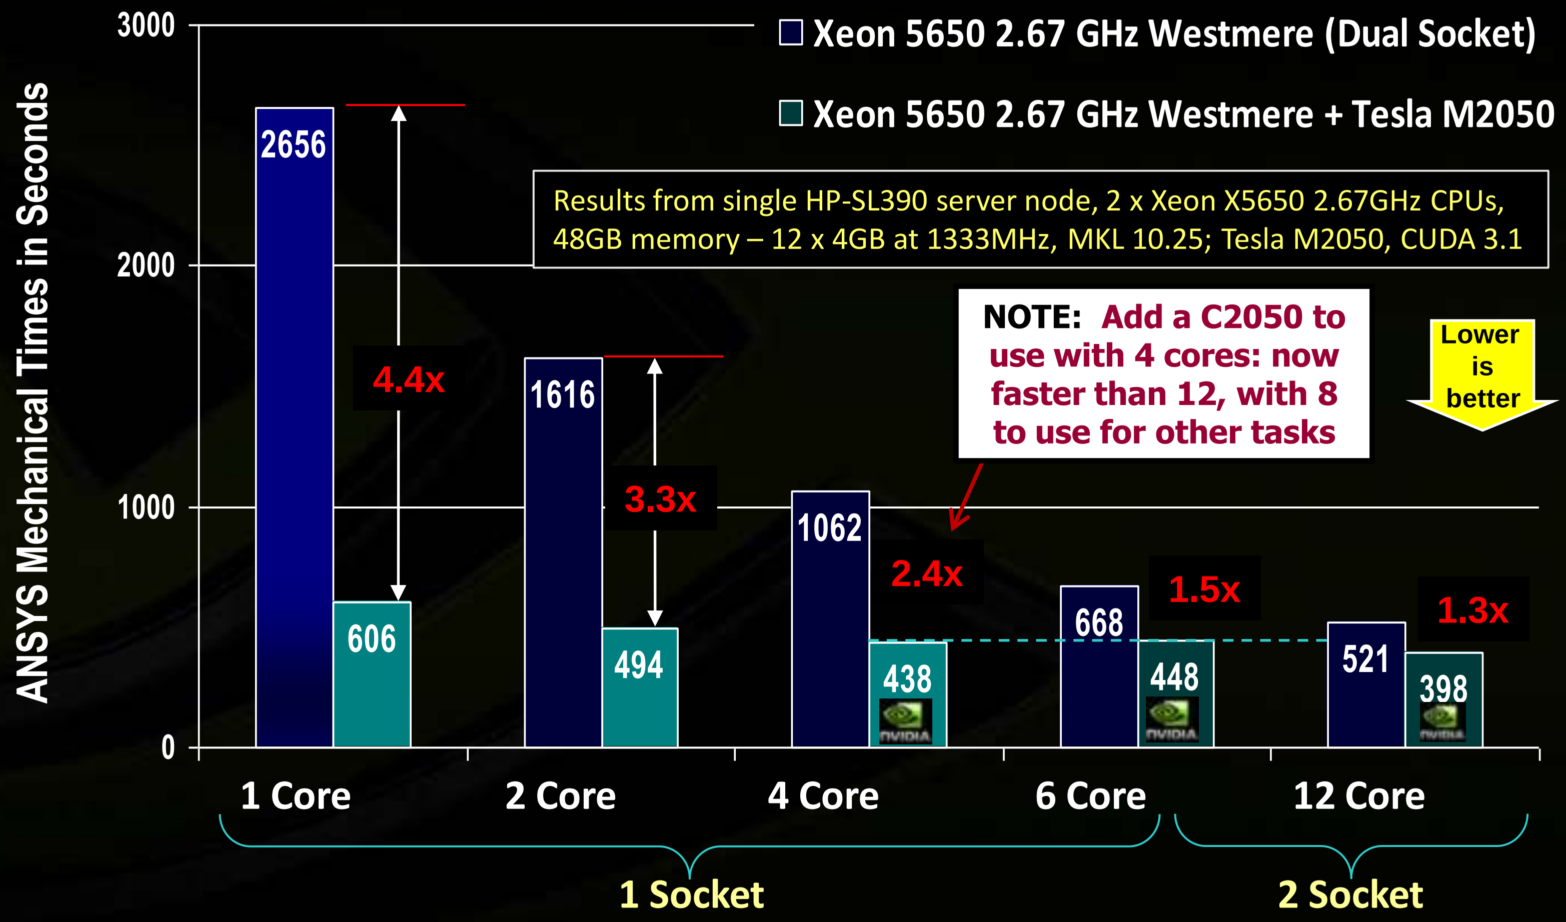

V12sp-5 Model

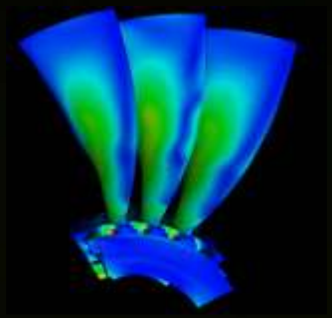

- Turbine geometry
- 2,100 K DOF
- SOLID187 FEs
- Static, nonlinear
- One load step
- Direct sparse

ANSYS Mechanical for Nehalem GPU Workstation

NOTE: Results Based on ANSYS Mechanical R13 Direct SMP Solver Sep 2010

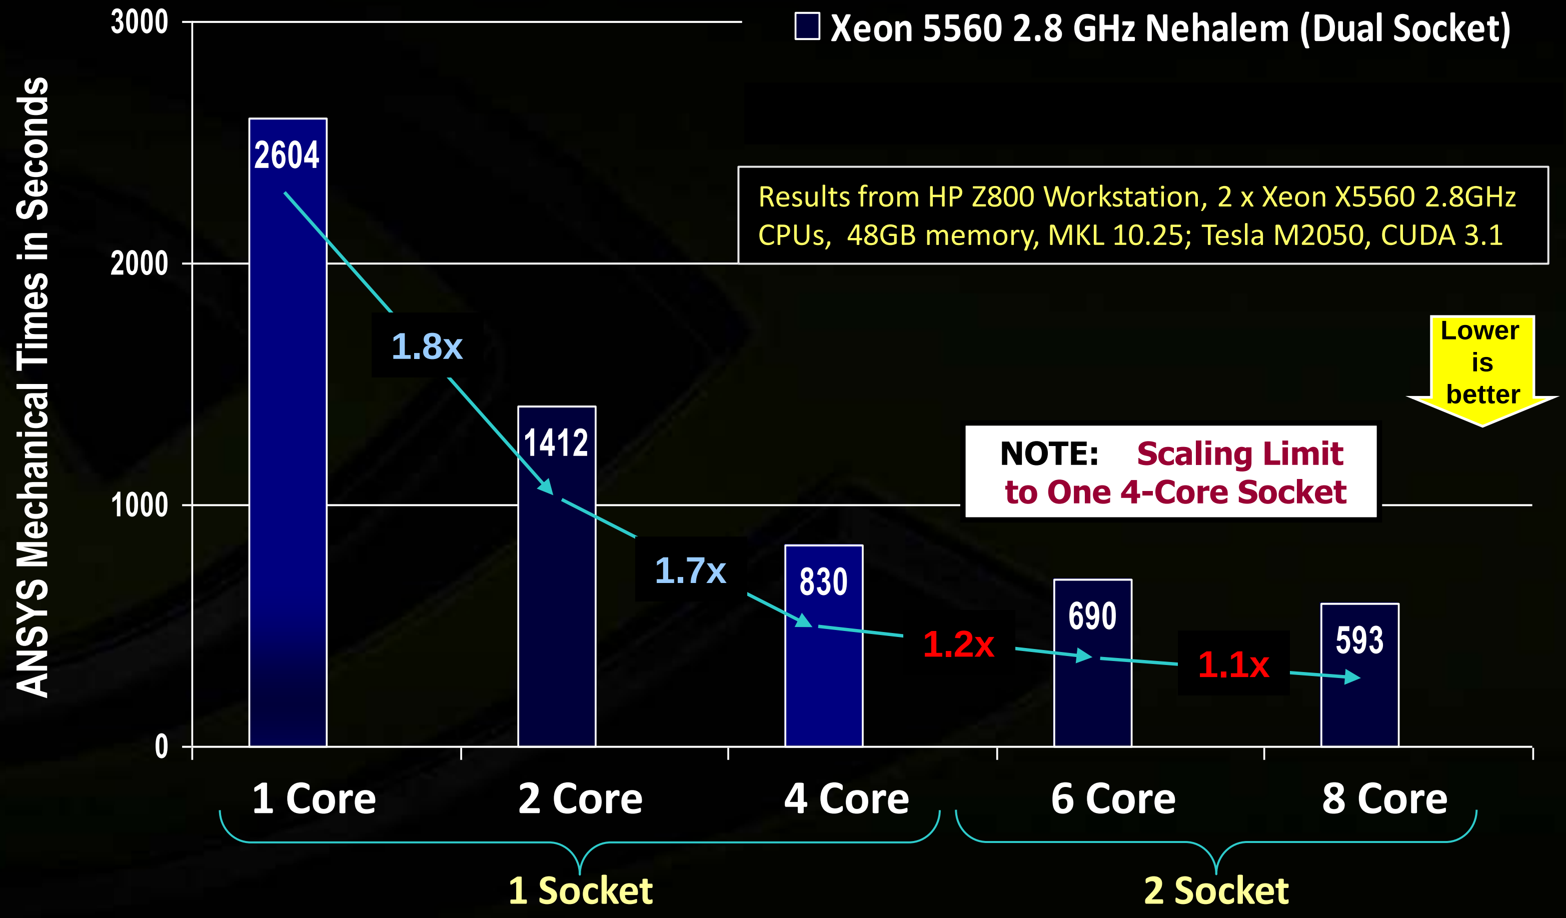

V12sp-5 Model

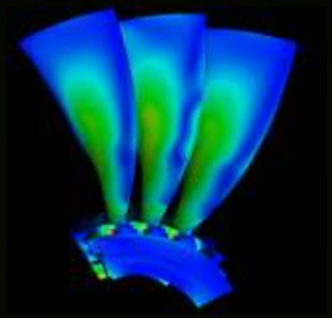

- Turbine geometry
- 2,100 K DOF
- SOLID187 FEs
- Static, nonlinear
- One load step
- Direct sparse

ANSYS Mechanical for Nehalem GPU Workstation

NOTE: Results Based on ANSYS Mechanical R13 Sparse Direct Solver Sep 2010

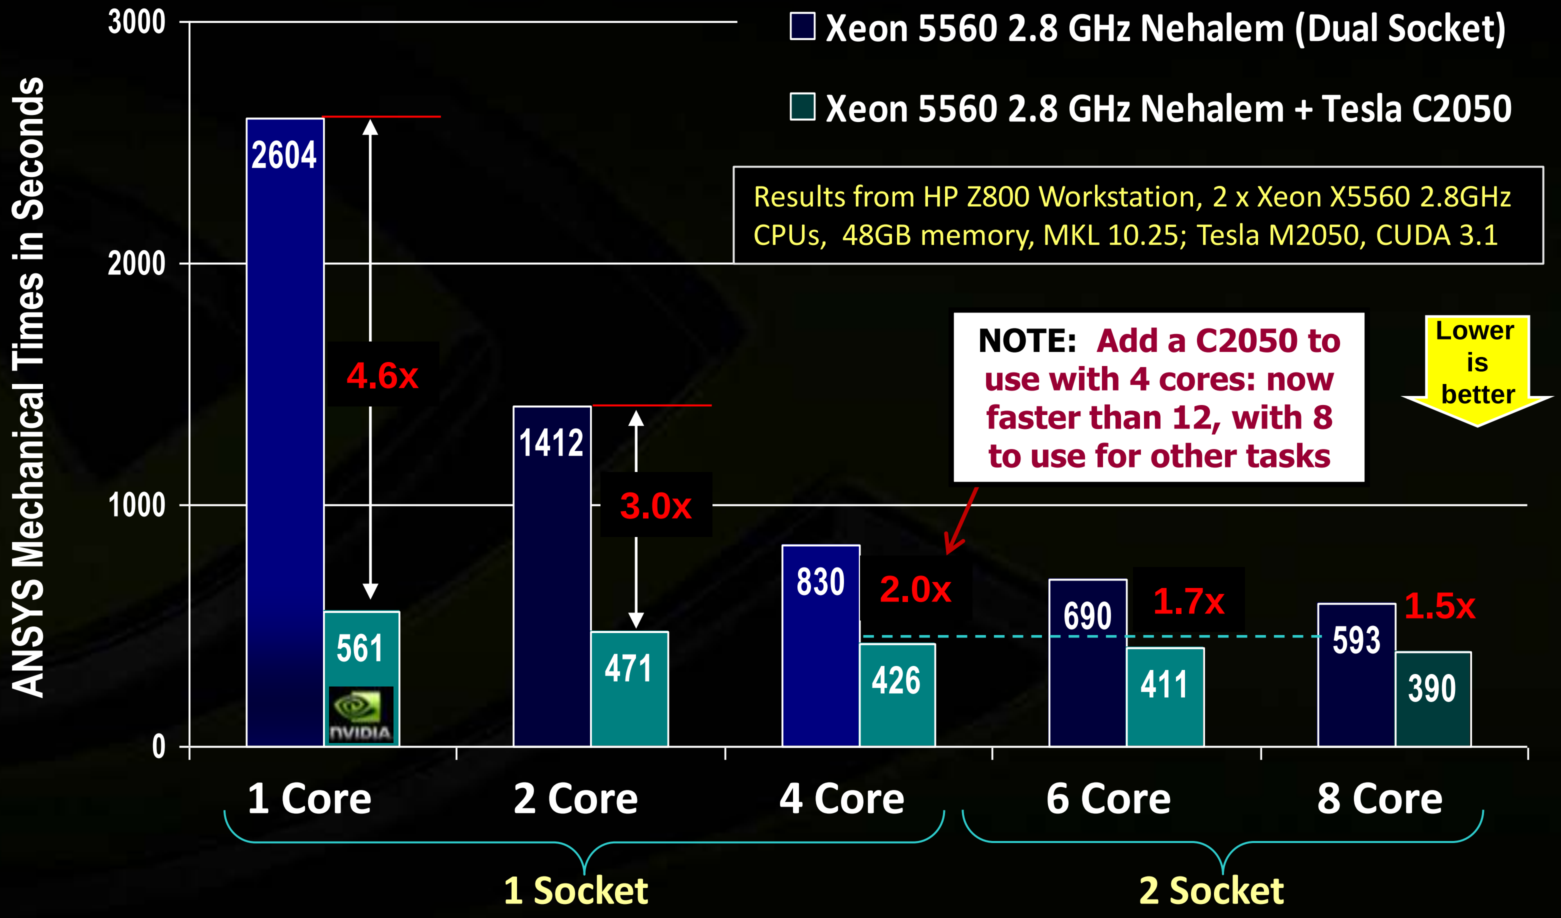

V12sp-5 Model

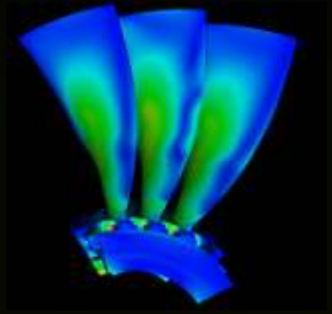

- Turbine geometry
- 2,100 K DOF
- SOLID187 FEs
- Static, nonlinear
- One load step
- Direct sparse

Effects of System CPU Memory for V12sp-5 Model

NOTE: Results Based on ANSYS Mechanical R13 SMP Direct Solver Sep 2010

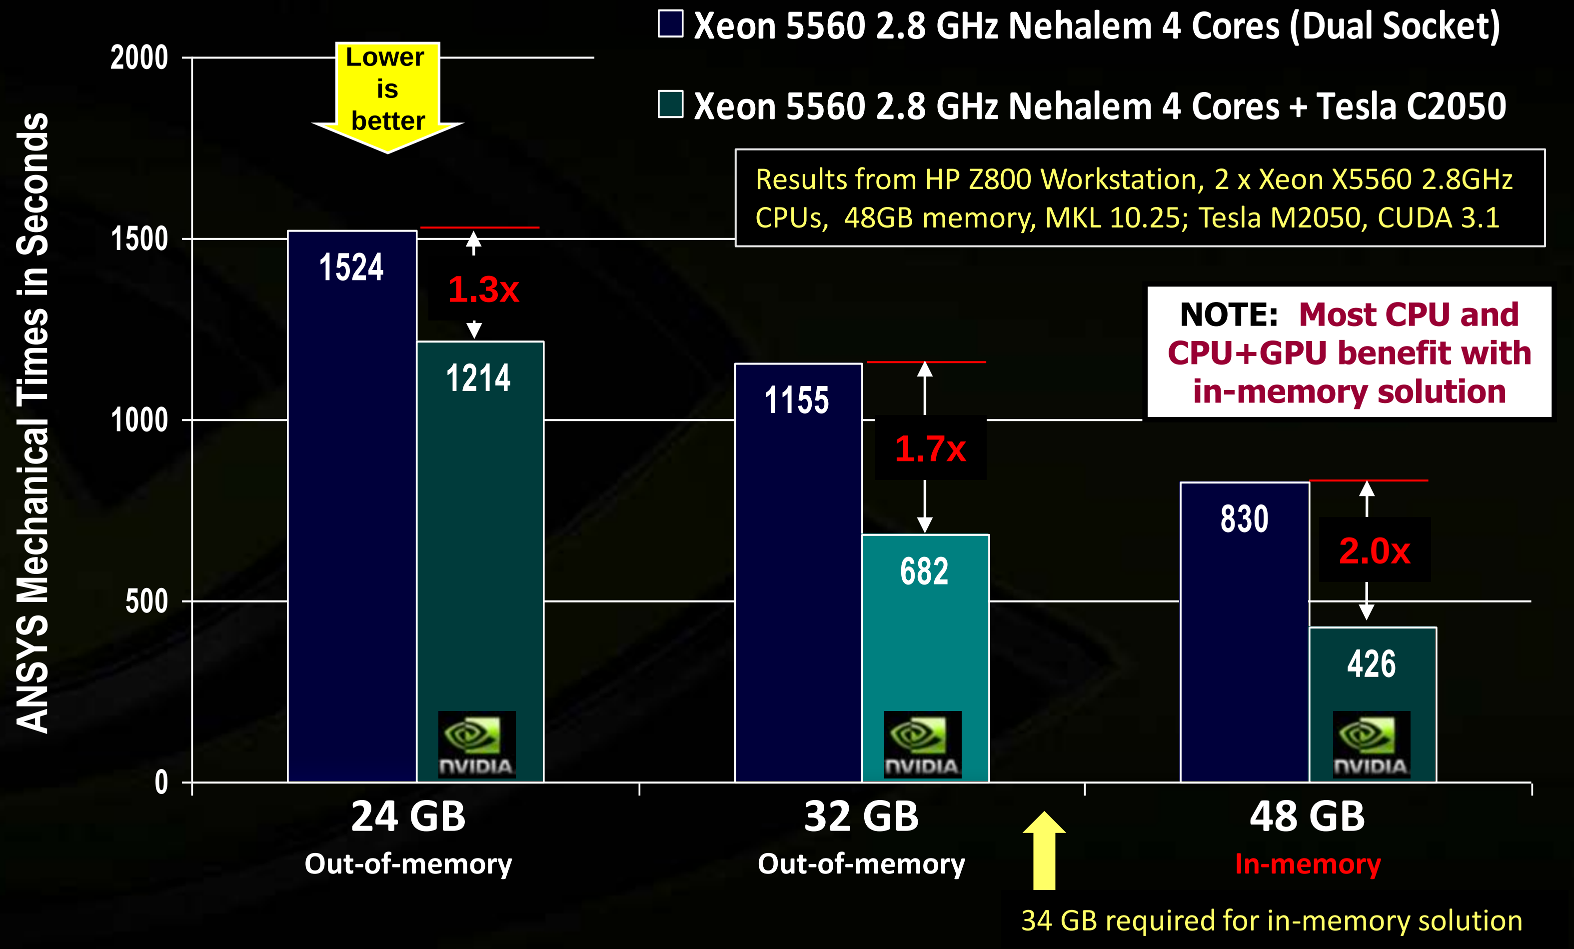

V12sp-5 Model

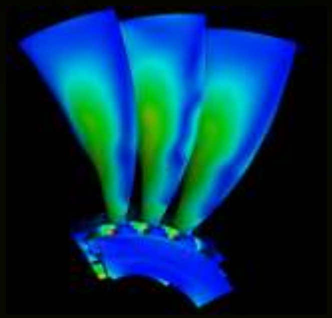

- Turbine geometry
- 2,100 K DOF
- SOLID187 FEs
- Static, nonlinear
- One load step
- Direct sparse

Effects of System CPU Memory for V12sp-5 Model

NOTE: Results Based on ANSYS Mechanical R13 SMP Direct Solver Sep 2010

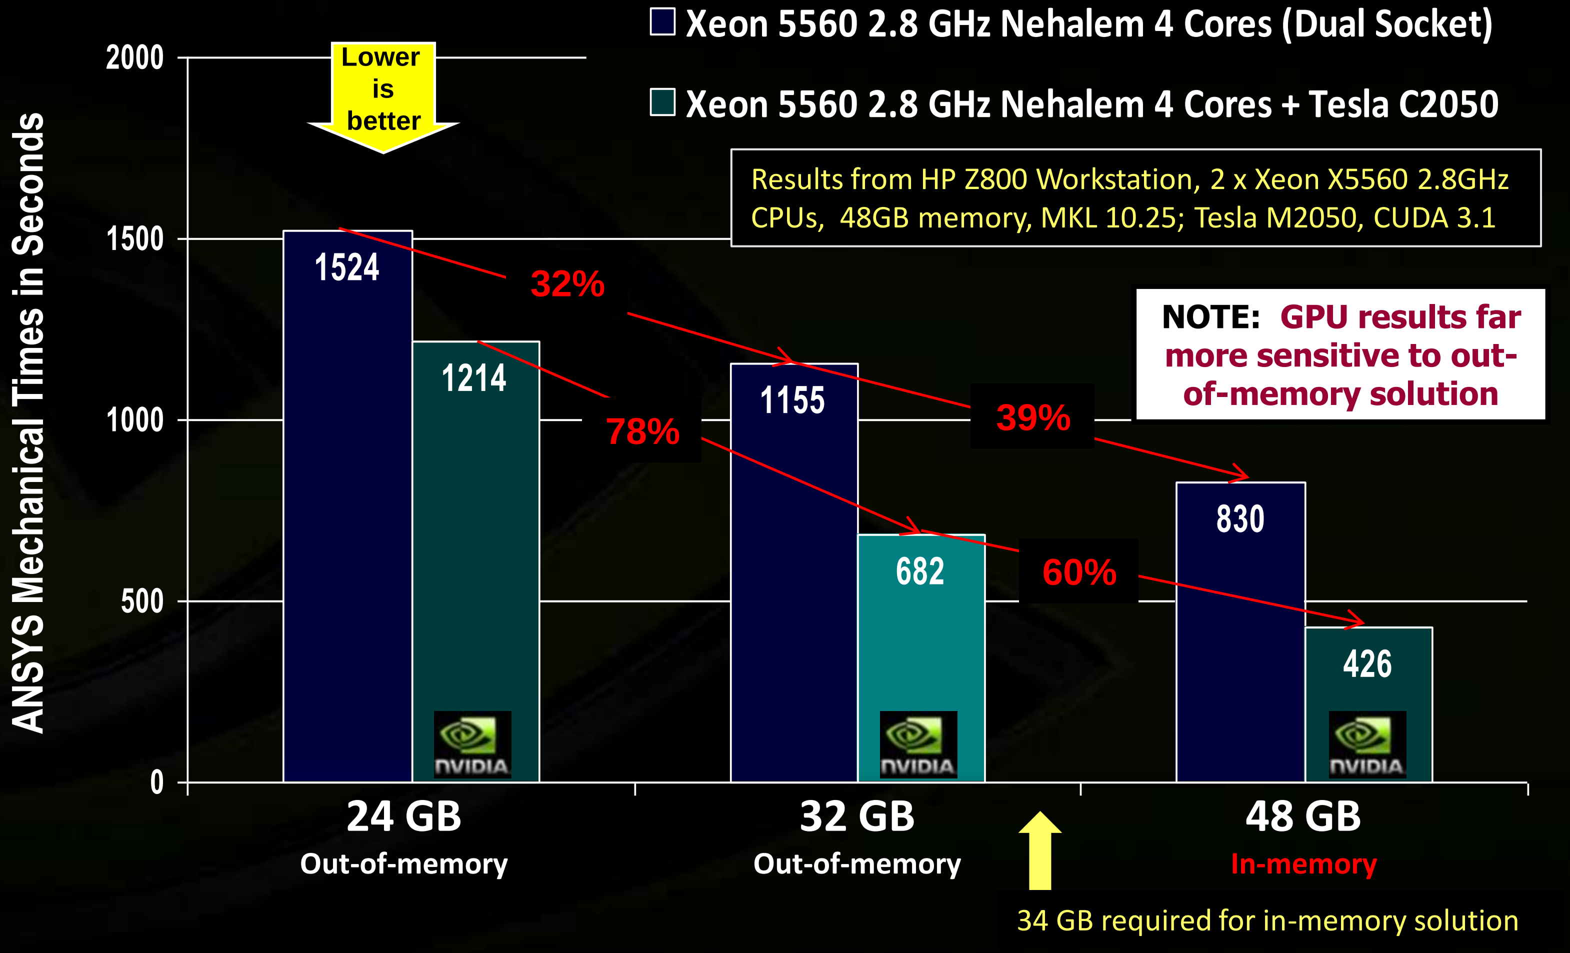

V12sp-5 Model

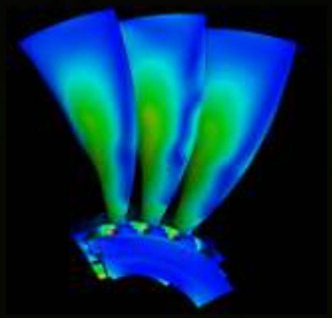

- Turbine geometry
- 2,100 K DOF
- SOLID187 FEs
- Static, nonlinear
- One load step
- Direct sparse

Economics of Engineering Codes in Practice

Cost Trends in CAE Deployment: **Costs in People and Software Continue to Increase**

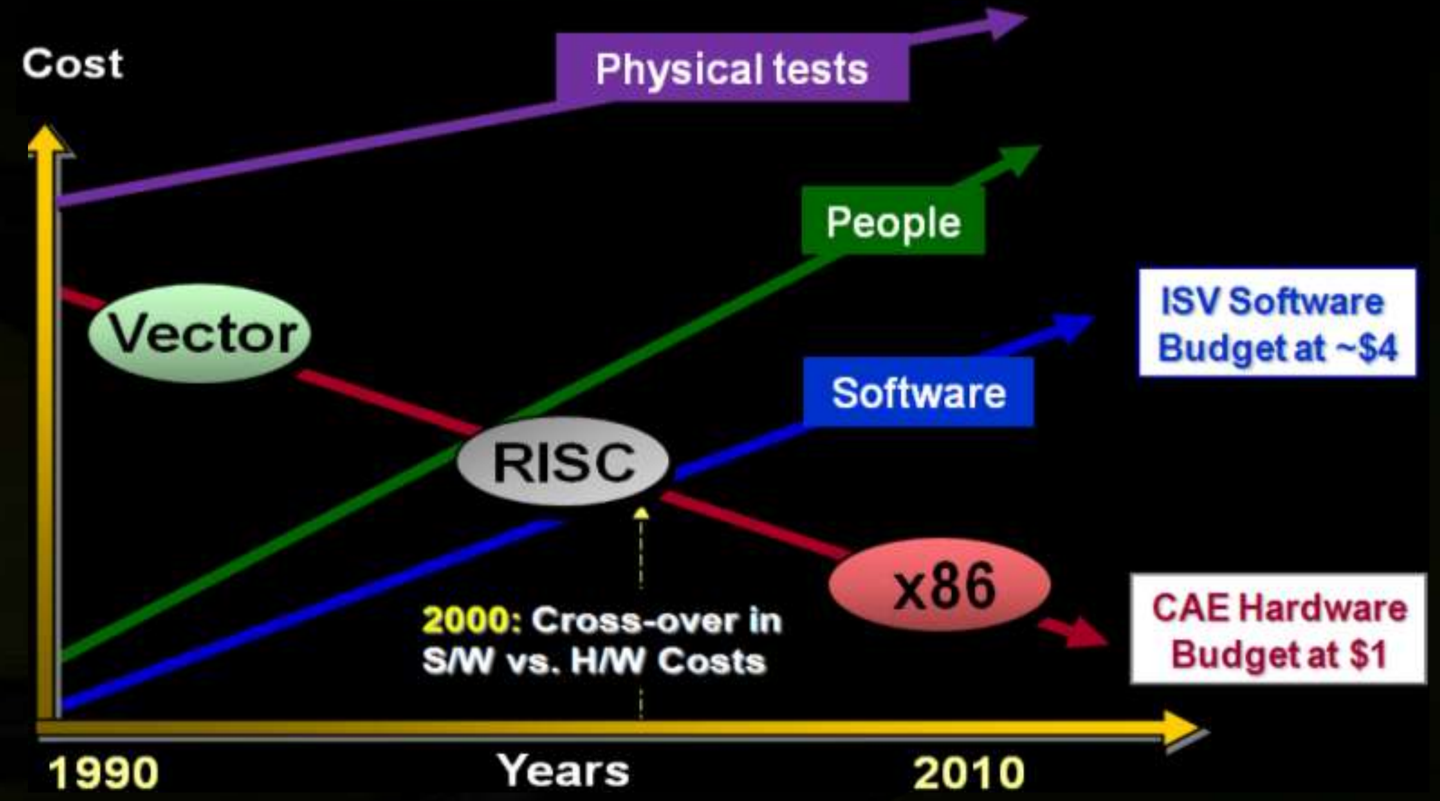

- Historically hardware very expensive vs. ISV software and people
- Software budgets are now 4x vs. hardware
- Increasingly important that hardware choices drive cost efficiency in people and software

Abaqus/Standard for Nehalem GPU Workstation

Abaqus/Standard: Based on v6.10-EF Direct Solver – Tesla C2050, CUDA 3.1 vs. 4-core Nehalem

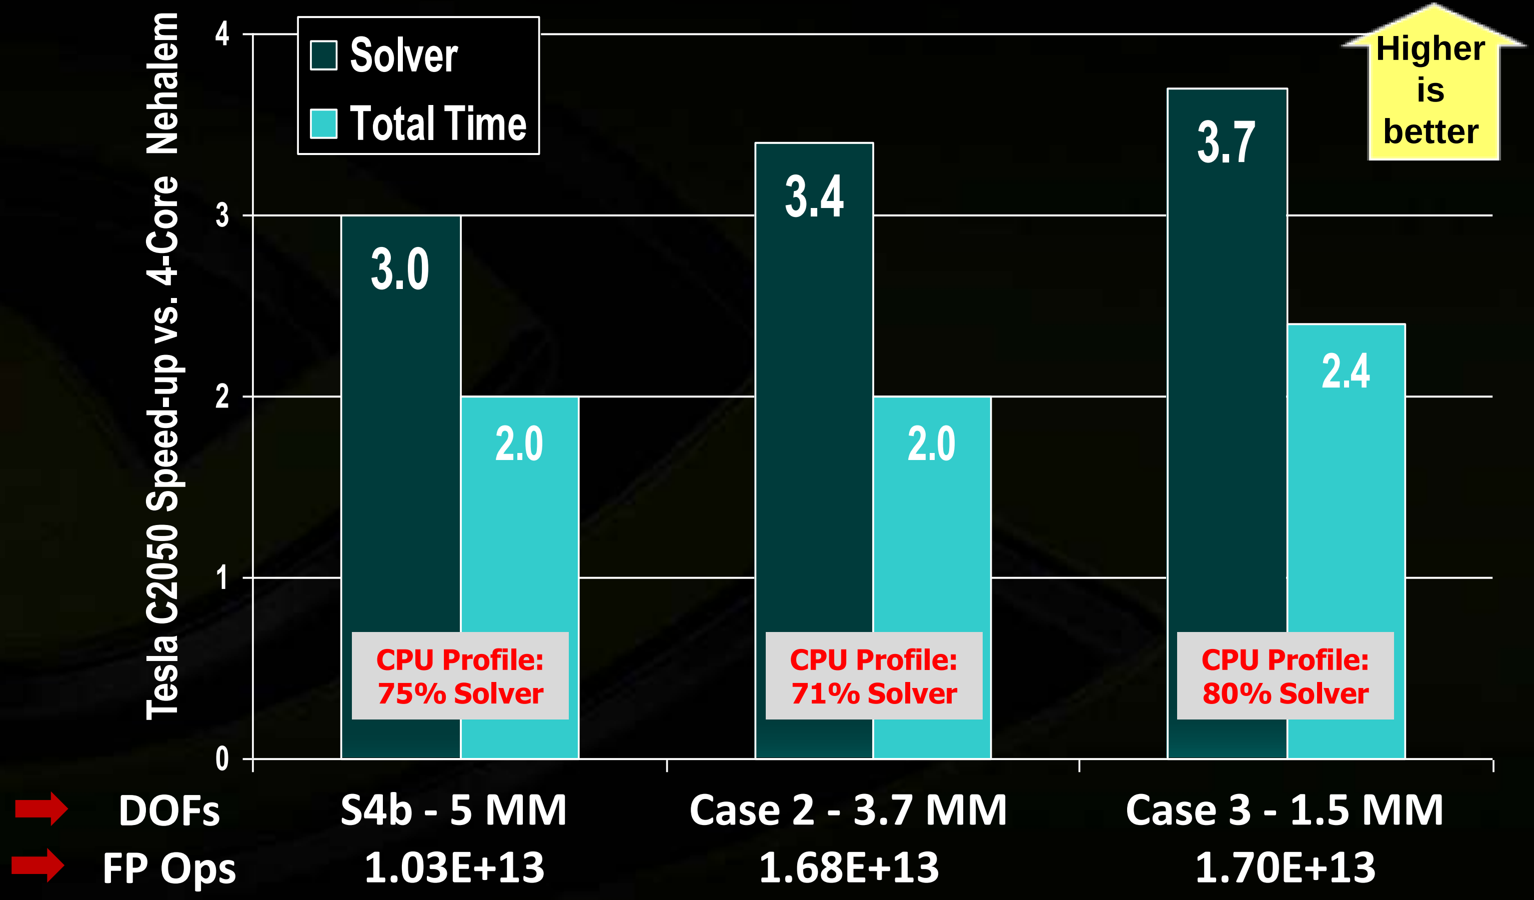

Source: SIMULIA Customer Conference, 27 May 2010:

"Current and Future Trends of High Performance Computing with Abaqus"

Presentation by Matt Dunbar

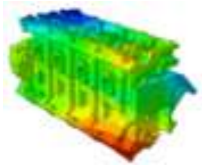

S4b: Engine Block Model of 5 MM DOF

NOTE: Solver Performance Increases with FP Operations

Results Based on 4-core CPU

Abaqus and NVIDIA Automotive Case Study

NOTE: Preliminary Results Based on Abaqus/Standard v6.10-EF Direct Solver

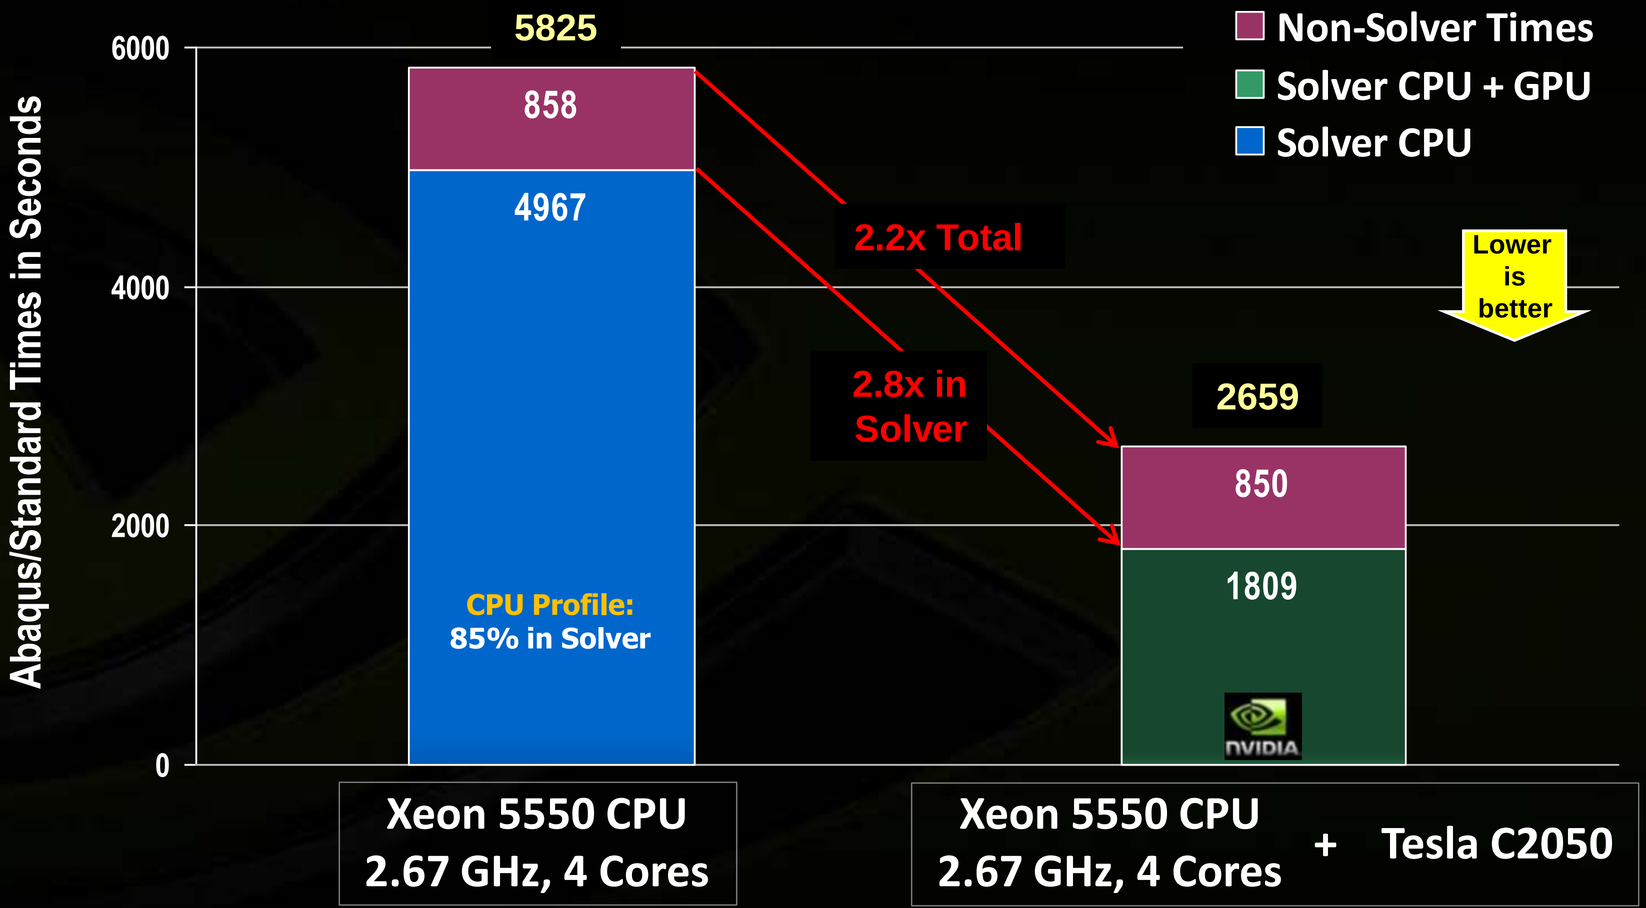

Results from HP Z800 Workstation, 2 x Xeon X5550 2.67 GHz CPUs, 48GB memory, MKL 10.25; Tesla C2050 with CUDA 3.1

Engine Model

- 1.5M DOF
- 2 Iterations
- 5.8e12 Ops per Iteration

Abaqus and NVIDIA Automotive Case Study

NOTE: Preliminary Results Based on Abaqus/Standard v6.10-EF Direct Solver

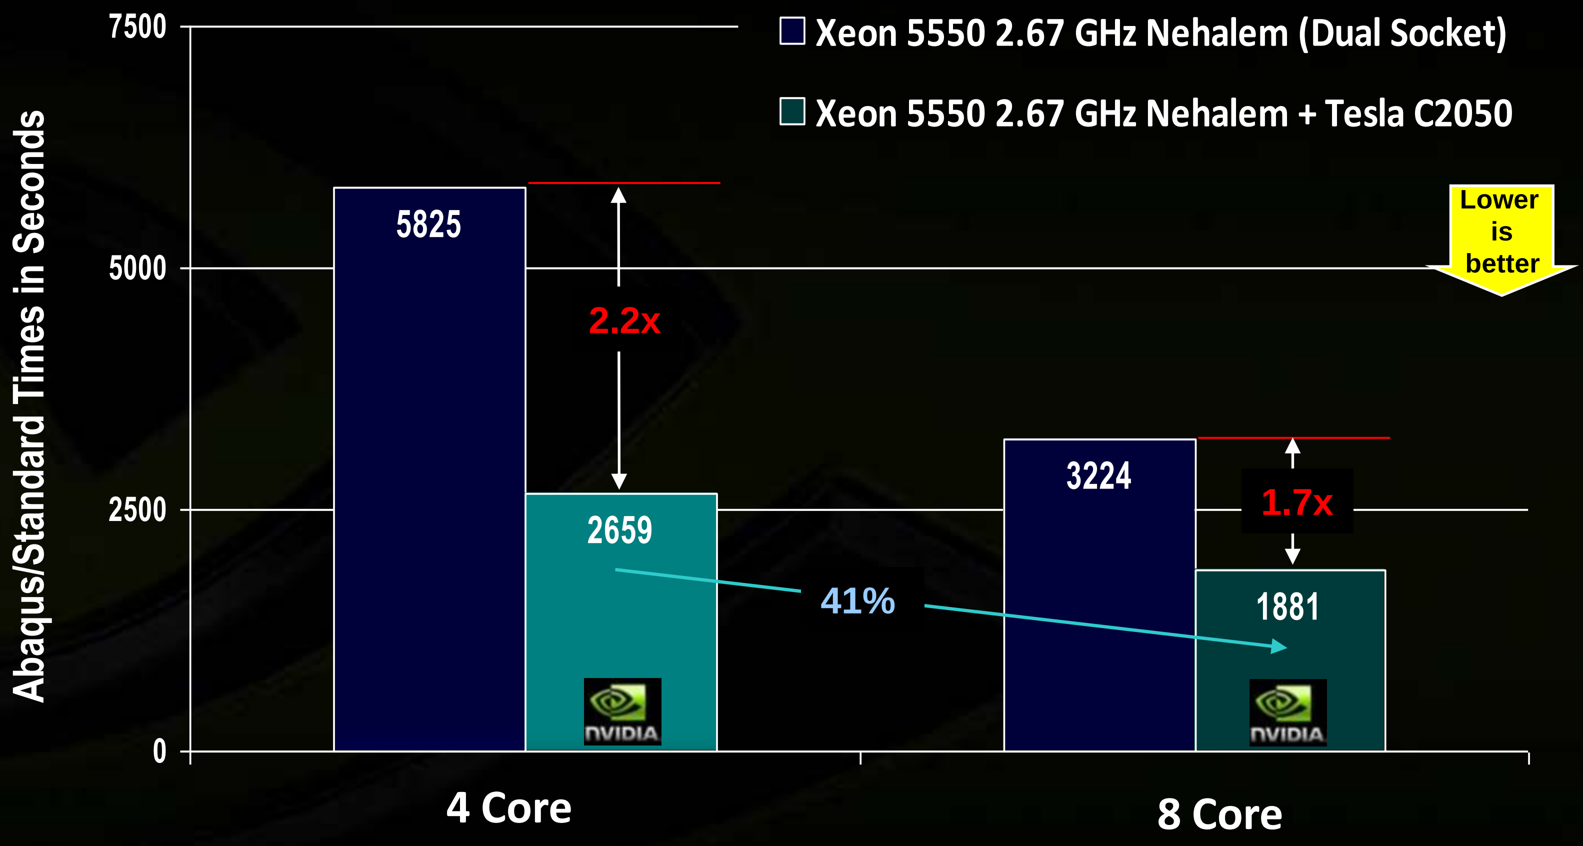

Engine Model

- 1.5M DOF
- 2 Iterations
- 5.8e12 Ops per Iteration

Results from HP Z800 Workstation, 2 x Xeon X5550 2.67 GHz CPUs, 48GB memory, MKL 10.25; Tesla C2050 with CUDA 3.1

LS-DYNA 971 Performance for GPU Acceleration

NOTE: Results of LS-DYNA Total Time for 300K DOF Implicit Model

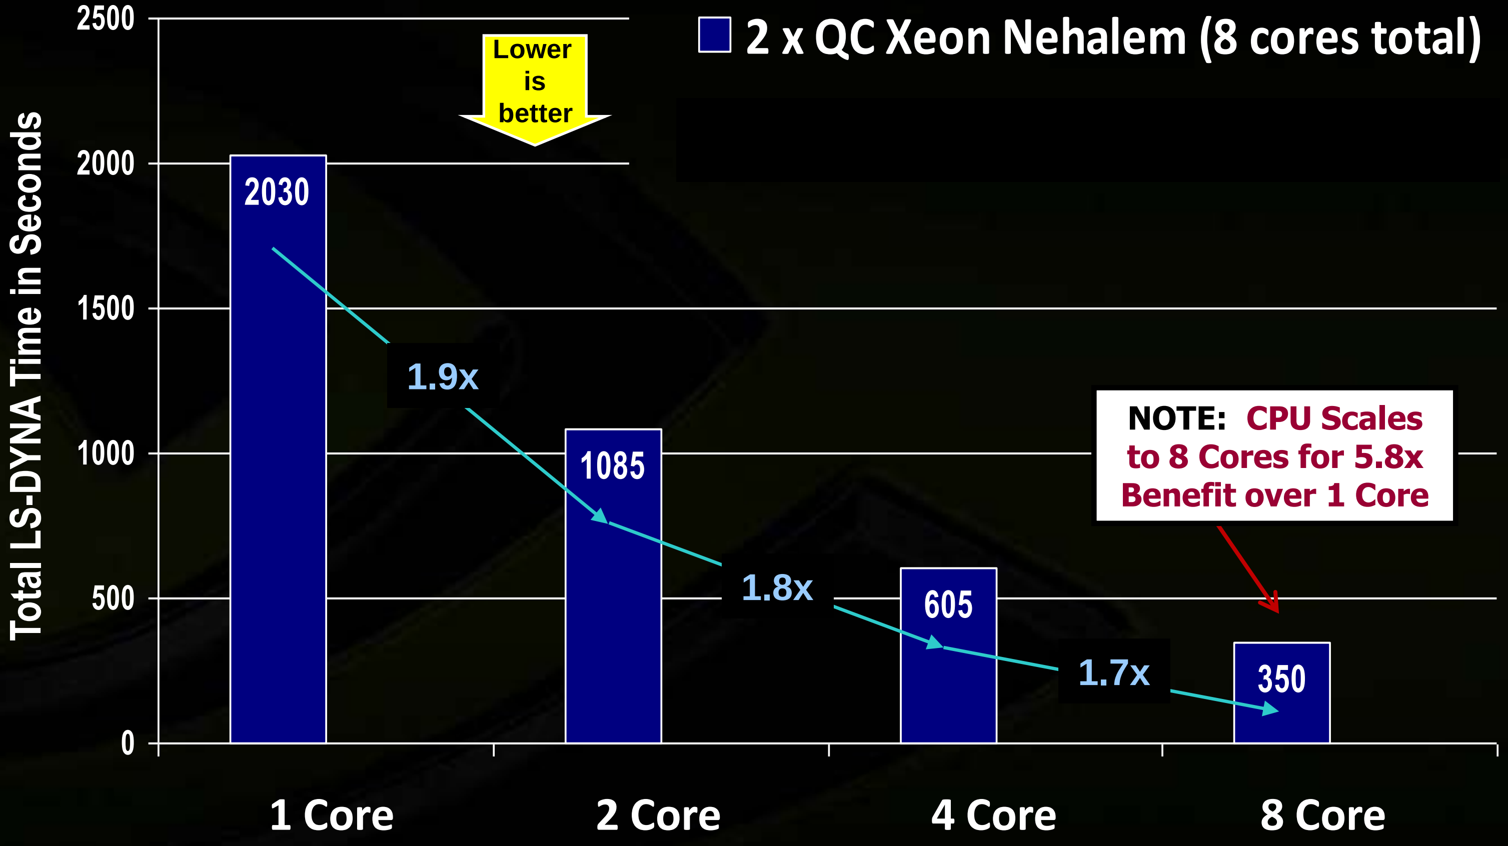

Results for CPU-only

LS-DYNA 971 Performance for GPU Acceleration

NOTE: Results of LS-DYNA Total Time for 300K DOF Implicit Model

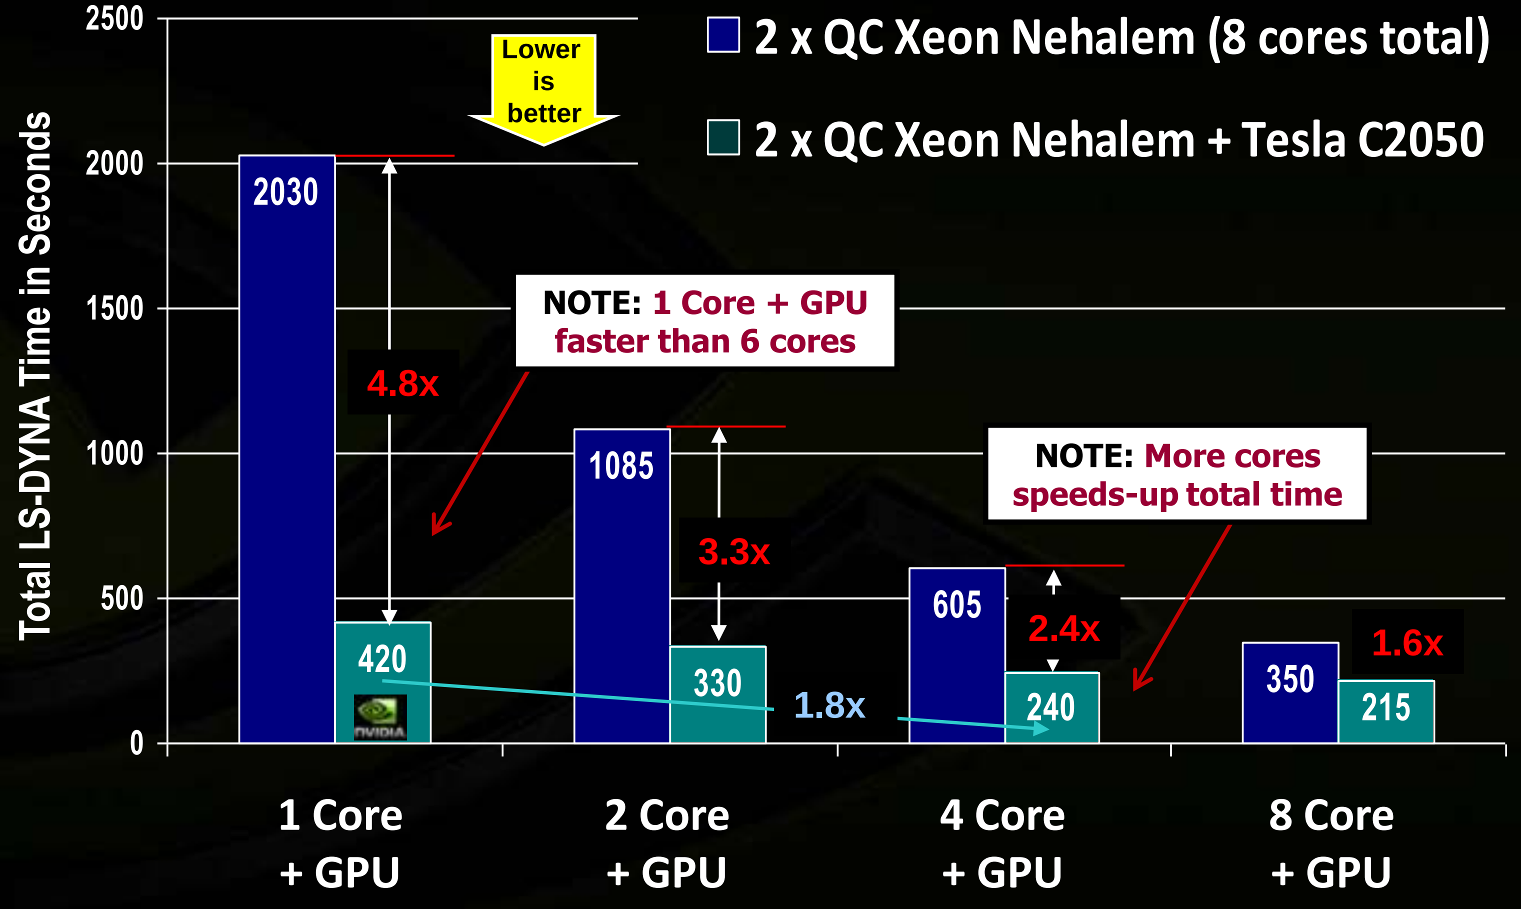

Add GPU Acceleration

OUTER3 Model

~300K DOF, 1 RHS

Distributed CSM and NVIDIA GPU Clusters

NOTE: Illustration Based on a Simple Example of 4 Partitions and 4 Compute Nodes

Model geometry is decomposed; partitions are sent to independent compute nodes on a cluster

Compute nodes operate distributed parallel using MPI communication to complete a solution per time step

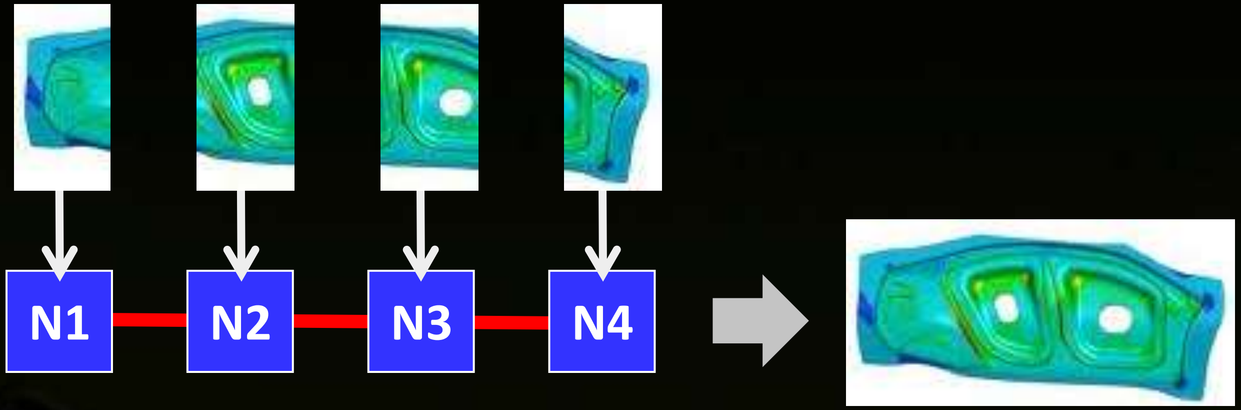

A global solution is developed at the completed time duration

Distributed CSM and NVIDIA GPU Clusters

NOTE: Illustration Based on a Simple Example of 4 Partitions and 4 Compute Nodes

Model geometry is decomposed; partitions are sent to independent compute nodes on a cluster

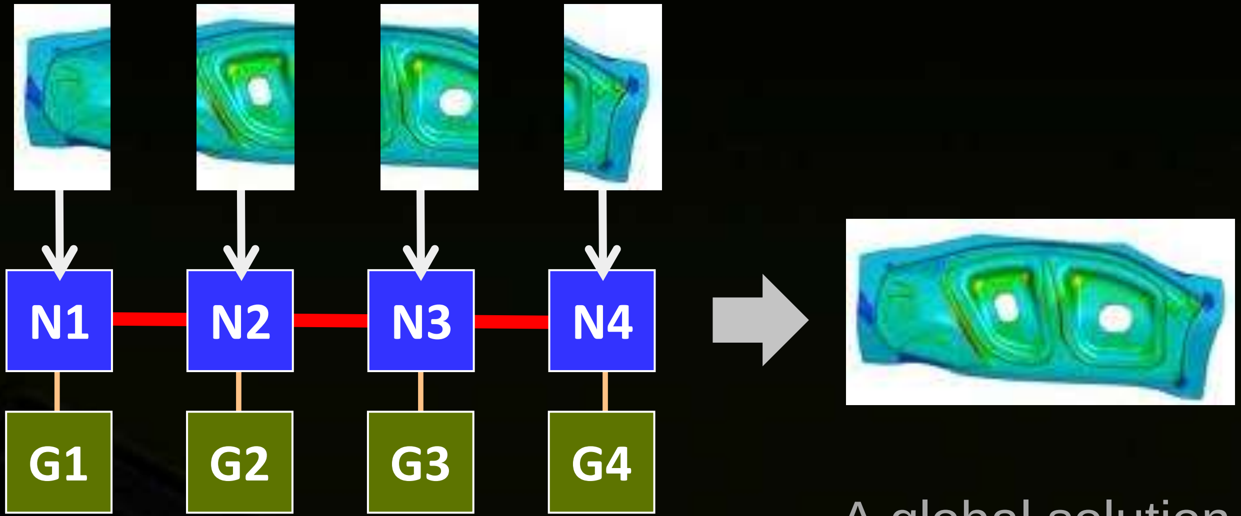

Compute nodes operate distributed parallel using MPI communication to complete a solution per time step

A global solution is developed at the completed time duration

A partition would be mapped to a GPU and provide shared memory OpenMP parallel – a 2nd level of parallelism in a hybrid model

GPU Priority by ISV Market Opportunity and "Fit"

#2 Computational Fluid Dynamics (CFD)

ANSYS CFD (FLUENT/CFX) | STAR-CCM+ | AcuSolve| CFD++| Particleworks | *OpenFOAM*

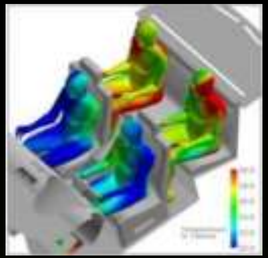

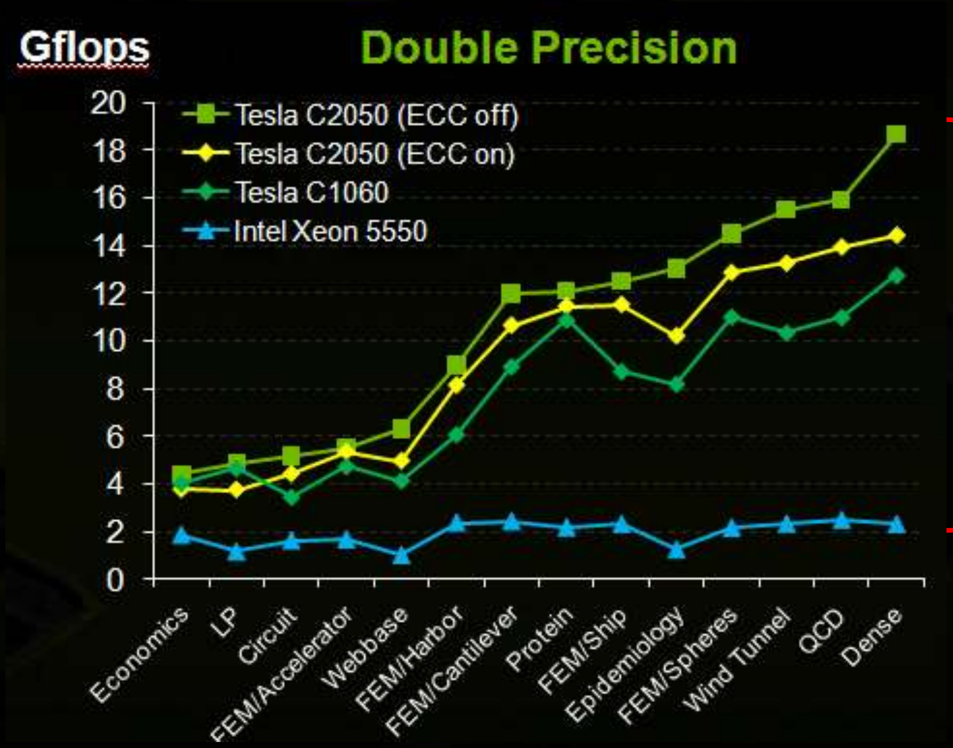

Typical Computational Profile of CFD (implicit)

Iterative linear equation solver ~50%

NOTE: Tesla C2050 9x Faster SpMV vs. QC Nehalem

SpMv: CUDA 3.0, Tesla C1060 and Tesla C2050
MKL 10.2: Intel Xeon 5550, 2.67 GHz

Performance of AcuSolve 1.8 on Tesla

AcuSolve: Profile is SpMV Dominant but Substantial Portion Still on CPU

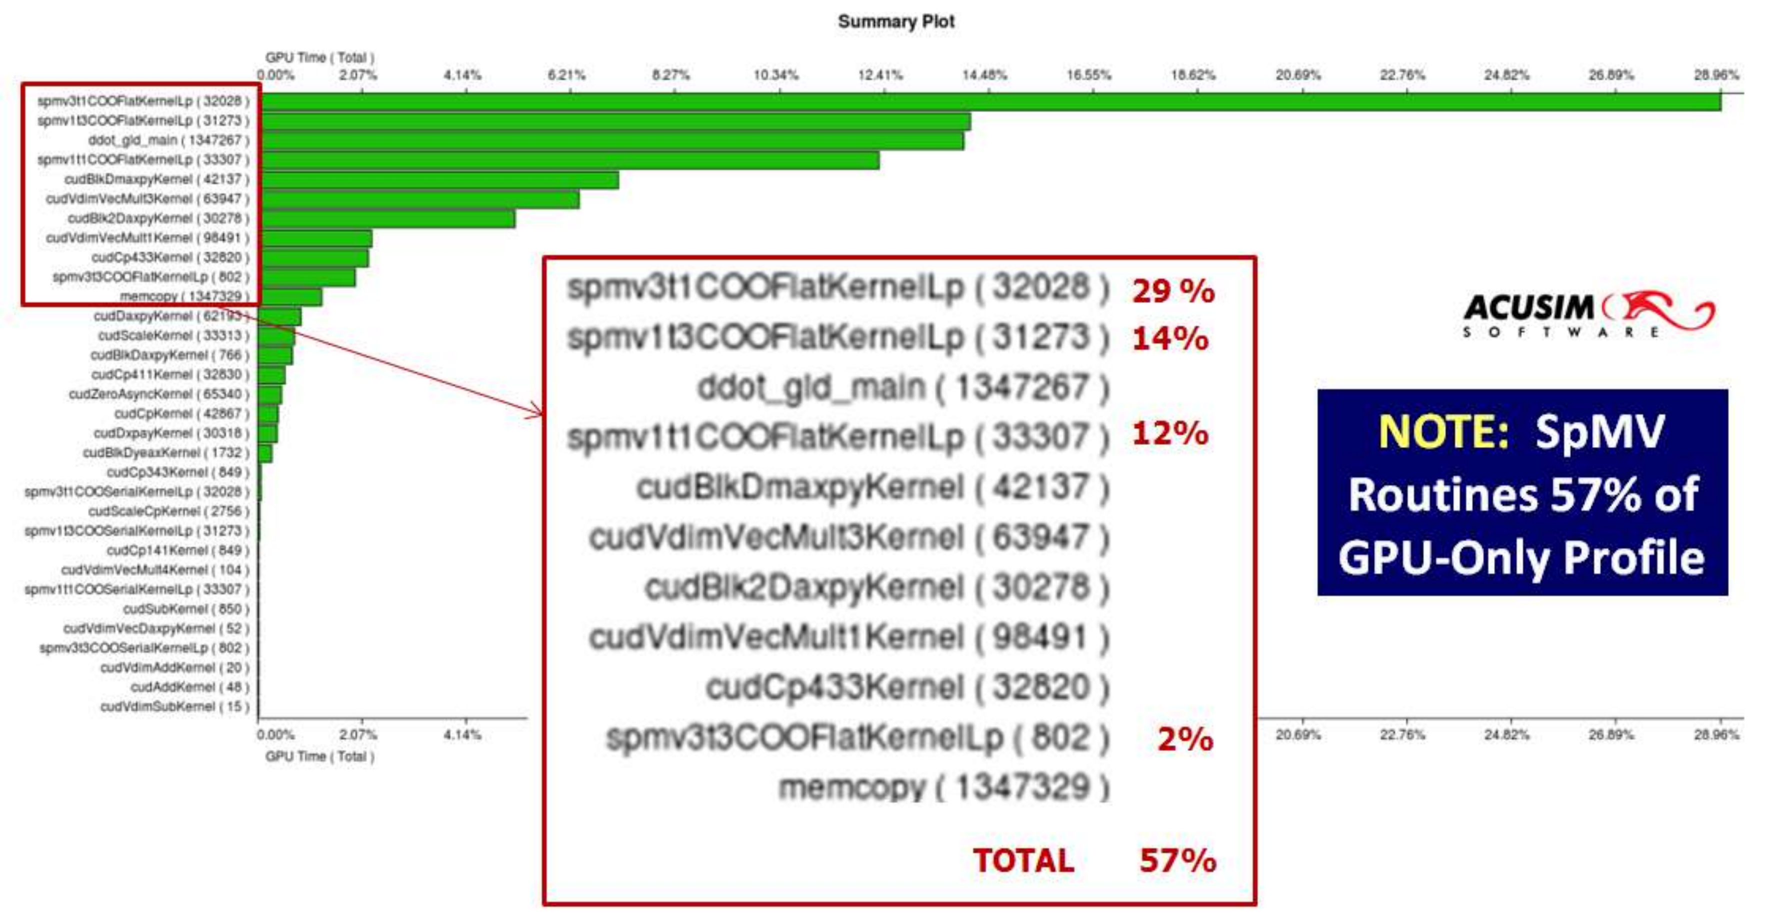

Performance of AcuSolve 1.8 on Tesla

AcuSolve: Comparison of Multi-Core Xeon CPU vs. Xeon CPU + Tesla GPU

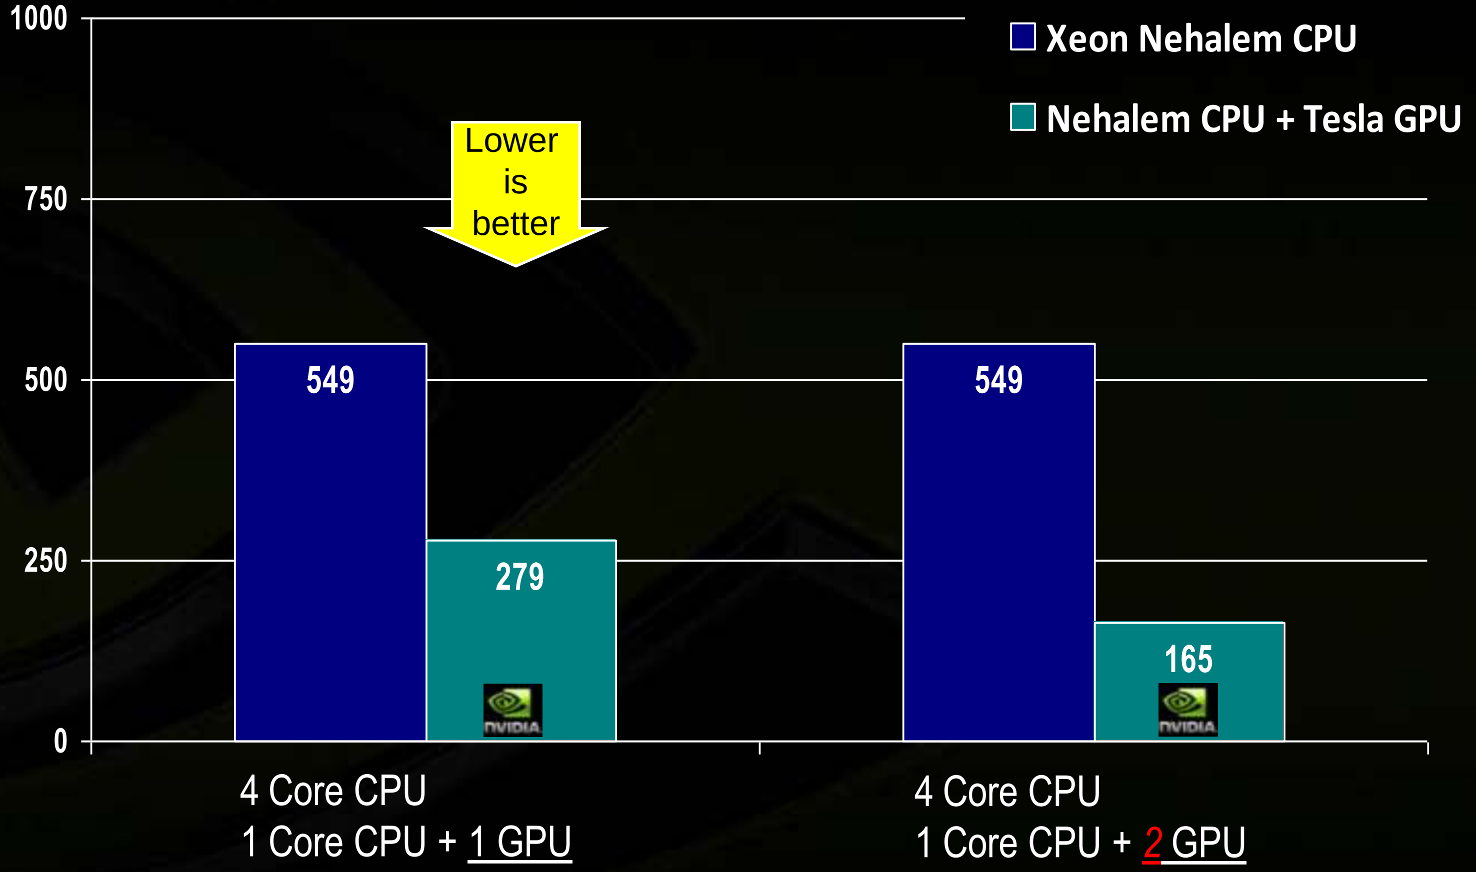

S-duct with 80K DOF

Hybrid MPI/Open MP for Multi-GPU test

CFD Developments and Publications on GPUs

48th AIAA Aerospace Sciences Meeting | Jan 2010 | Orlando, FL, USA

FEFLO: **Porting of an Edge-Based CFD Solver to GPUs**
[AIAA-2010-0523] Andrew Corrigan, Ph.D., Naval Research Lab; Rainald Lohner, Ph.D., GMU

FAST3D: **Using GPU on HPC Applications to Satisfy Low Power Computational Requirement**
[AIAA-2010-0524] Gopal Patnaik, Ph.D., US Naval Research Lab

OVERFLOW: **Rotor Wake Modeling with a Coupled Eulerian and Vortex Particle Method**
[AIAA-2010-0312] Chris Stone, Ph.D., Intelligent Light

CFD on Future Architectures | Oct 2009 | DLR Braunschweig, DE

Veloxi: **Unstructured CFD Solver on GPUs**
Jamil Appa, Ph.D., BAE Systems Advanced Technology Centre

elsA: **Recent Results with elsA on Many-Cores**
Michel Gazaix and Steve Champagneux, ONERA / Airbus France

Turbostream: **Turbostream: A CFD Solver for Many-Core Processors**
Tobias Brandvik, Ph.D. , Whittle Lab, University of Cambridge

Parallel CFD 2009 | May 2009 | NASA Ames, Moffett Field, CA, USA

OVERFLOW: **Acceleration of a CFD Code with a GPU**
Dennis Jespersen, NASA Ames Research Center

GPU Results for Grid-Based Continuum CFD

NVIDIA

Success Demonstrated in Full Range of Time and Spatial Schemes

	Structured Grid	Unstructured
Explicit [usually compressible]	UNIVERSITY OF CAMBRIDGE **TurboStream** Sandia National Laboratories, Oak Ridge National Laboratory **S3D** ~15x	BAE SYSTEMS **Veloxi** George Mason University ~8x **FEFLO**
Implicit [usually incompressible]	~4x **U.S. Engine Co.** ***Internal flows*** NCSA **DNS**	~2x *ISVs* ACUSIM SOFTWARE **AcuSolve** Autodesk **Moldflow**

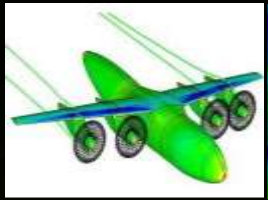

Aircraft aero

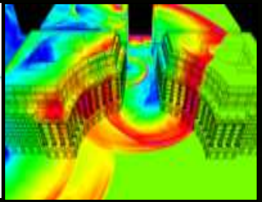

Bldg air blast

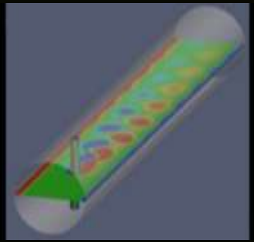

Chem mixer

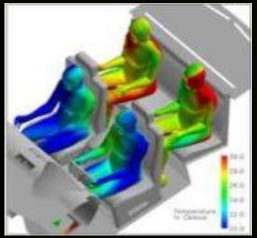

Auto climate

Speed-ups based on use of 4-core Xeon X5550 2.67 GHz

Culises: New CFD Solver Library for OpenFOAM

GPU-based HPC for Fluid Dynamics

FluiDyna INNOVATIVE ENGINEERING AND BEYOND

Culises

Aim : Acceleration of CFD simulation

A CUDA library for iterative solution of equation systems on GPUs

Features

- State-of-the art iterative solvers (Precond. CGs, Multigrid)
- Support of unstructured comput. meshes for efficient description of complex geon
- Support of single (4 byte) and double (8 byte) precision floating point numbers
- Interfaces to customer specific software packages (*OpenFOAM*,...)

Benefits

- Acceleration of comput. expensive algorithms of existing customer software
- Significant reduction of computing times
- Increased resolution for improved detailling of the computer model of the real sys
- Porting complex software packages is avoided, but only the most expensive part
- Repeated validation of complete software packages is avoided

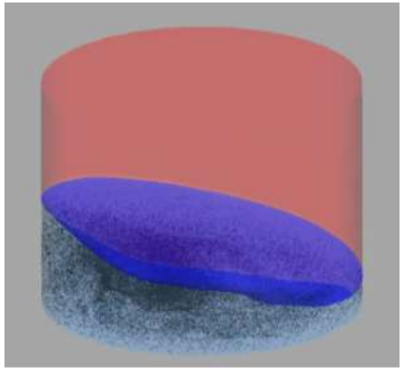

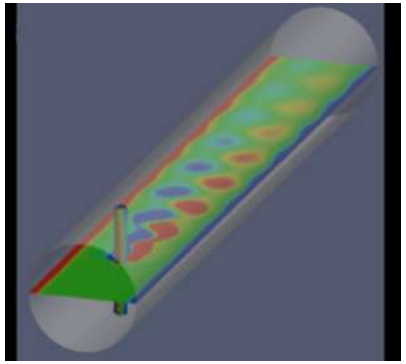

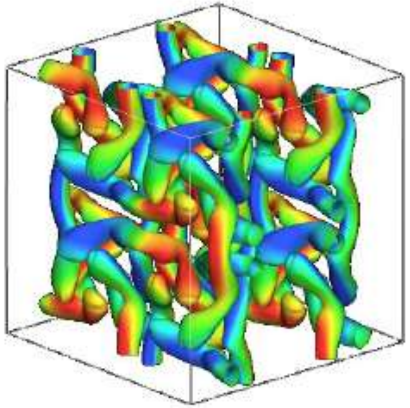

Fig 1: Taylor Green Vortex (iso-contours of Q-criterion colored by velocity)

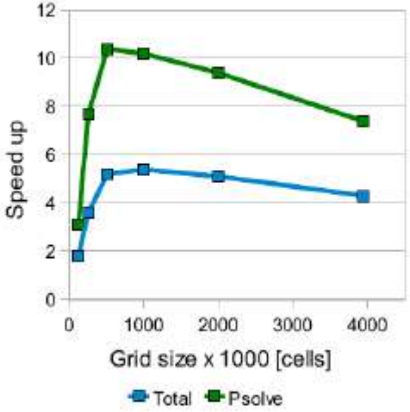

Fig 2: Acceleration of pressure solver 'Psolve' for hybrid GPU-CPU code compared to CPU-only code

Prometech and Particle-Based CFD for Multi-GPUs

- **Particleworks from Prometech Software**

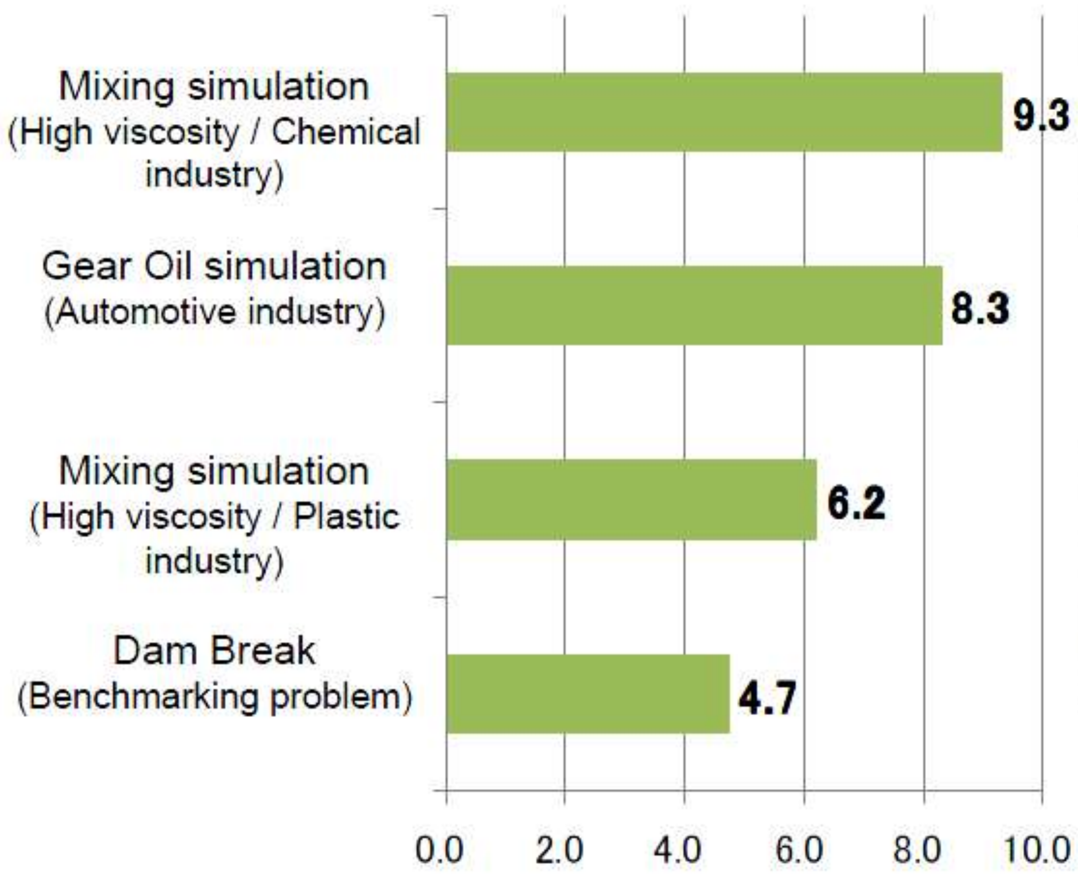

Reference 1.0 equals Intel Core i7 4cores

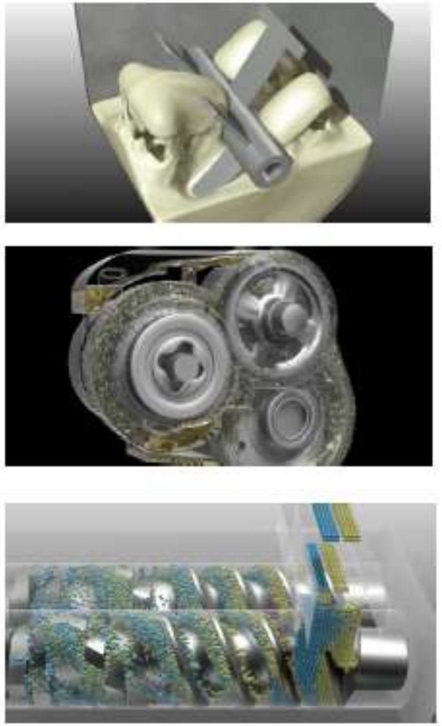

MPS-based method developed at the University of Tokyo [Prof. Koshizuka]

Preliminary results for Particleworks 2.5 with released planned for 2011

Performance is relative to 4 cores of Intel i7 CPU

Contact Prometech for release details

IMPETUS AFEA Results for GPU Computing

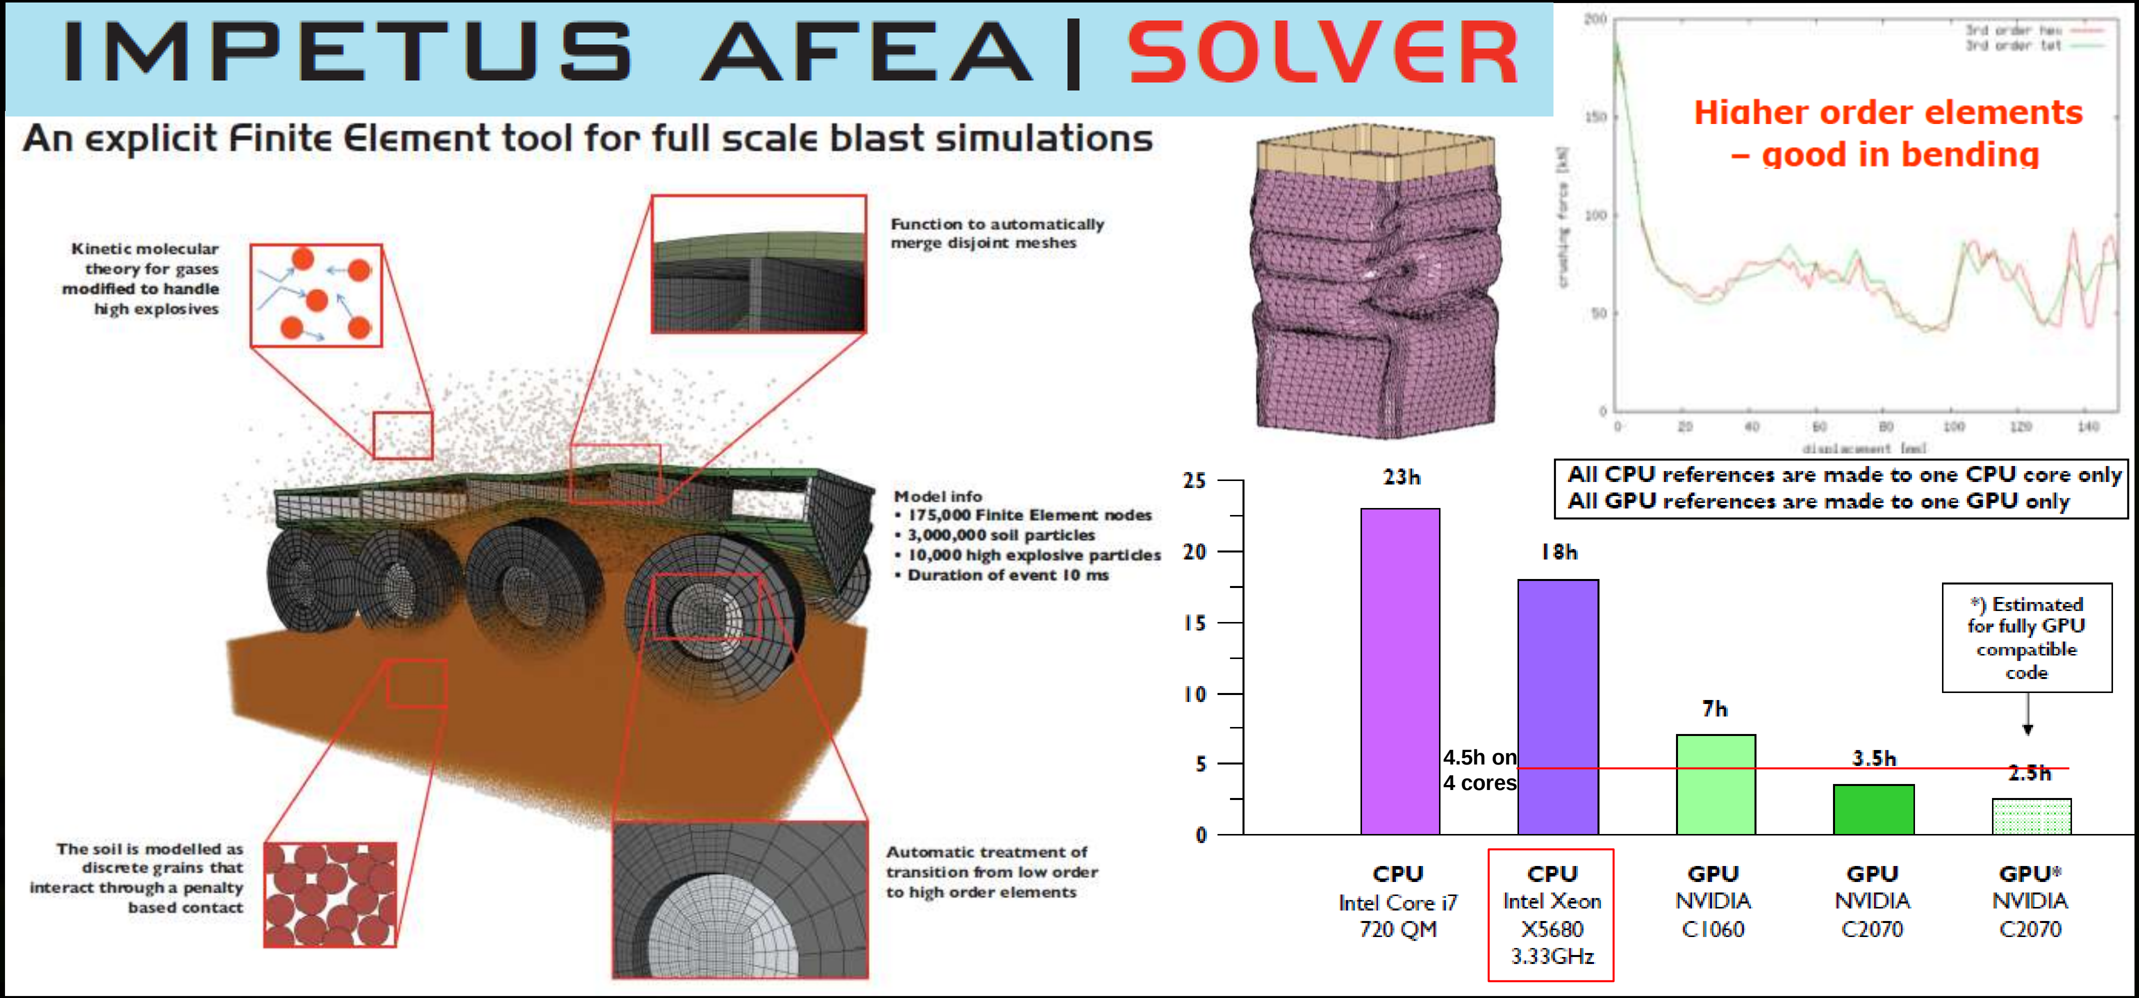

Summary of Engineering Code Progress for GPUs

- **GPUs are an Emerging HPC Technology for ISVs**
  - **Industry Leading ISV Software is GPU-Enabled Today**
- **Initial GPU Performance Gains are Encouraging**
  - **Just the beginning of more performance and more applications**
- **NVIDIA Continues to Invest in ISV Developments**
  - **Joint technical collaborations at most Engineering ISVs**

Contributors to the ISV Performance Studies

SIMULIA

- Mr. Matt Dunbar, Technical Staff, Parallel Solver Development
- Dr. Luis Crivelli, Technical Staff, Parallel Solver Development

ANSYS

- Mr. Jeff Beisheim, Technical Staff, Solver Development

USC Institute for Information Sciences

- Dr. Bob Lucas, Director of Numerical Methods

ACUSIM (Now a Division of Altair Engineering)

- Dr. Farzin Shakib, Founder and President

Thank You, Questions ?

Stan Posey | CAE Market Development

NVIDIA, Santa Clara, CA, USA