GPU-Acceleration of CAE Simulations

Bhushan Desam
NVIDIA Corporation
bdesam@nvidia.com

AGENDA

- GPUs in Enterprise Computing
- Business Challenges in Product Development
- NVIDIA GPUs for CAE Applications
- Computational Fluid Dynamics (CFD) Applications
- Computational Structural Mechanics (CSM) Applications
- Computational Electro Magnetics (CEM) Applications
- NVIDIA Solutions
- Value of GPU-accelerated Simulations

NVIDIA Enterprise Group
Visualization, Accelerated Computing & Virtualization

QUADRO
Revolutionizing Design & Visualization

TESLA
Accelerating Momentum in HPC and Big Data Analytics

GRID
Enabling End-to-End Enterprise Virtualization

TESLA

Accelerating Computing

GPUs enable tremendous breakthroughs by simply enabling us to do more, faster

ANSYS, SIMULIA, and other major ISVs leverage GPUs to accelerate engineering simulation for better design

The world's top10 energy-efficient supercomputers use NVIDIA GPUs according to TheGreen500 list

AGENDA

- GPUs in Enterprise Computing
- Business Challenges in Product Development
- NVIDIA GPUs for CAE Applications
- Computational Fluid Dynamics (CFD) Applications
- Computational Structural Mechanics (CSM) Applications
- Computational Electro Magnetics (CEM) Applications
- NVIDIA Solutions
- Value of GPU-accelerated Simulations

Business Challenges in Product Development

Simulation is playing an important role in meeting these challenges, thus add value to many product development companies

Changing Role of Simulation in Product Development

From insight to product innovation

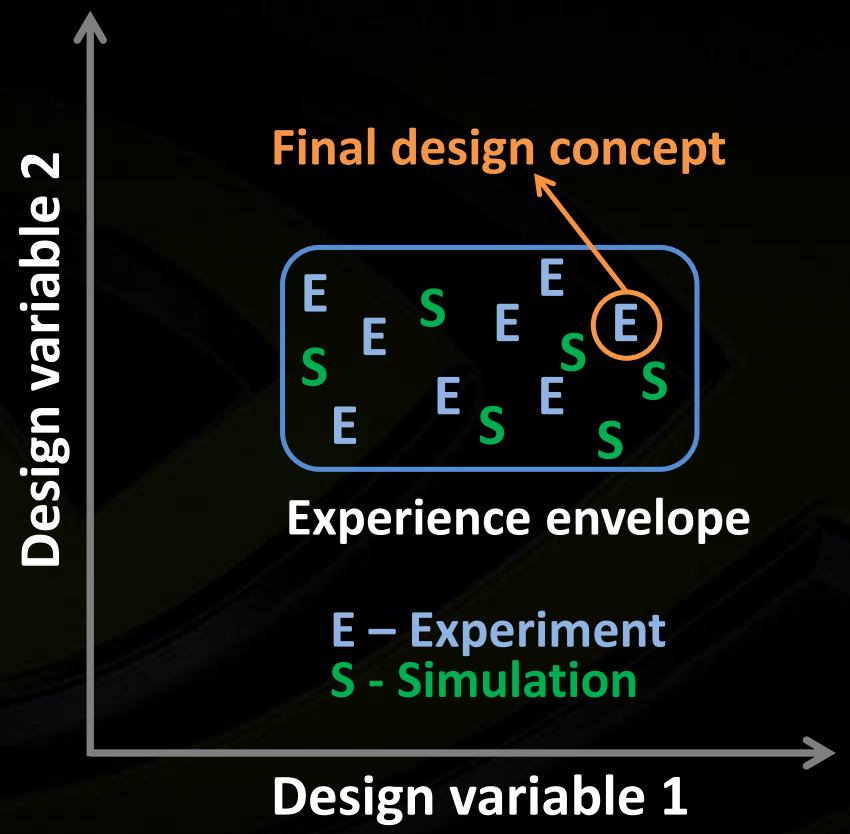

Building insight

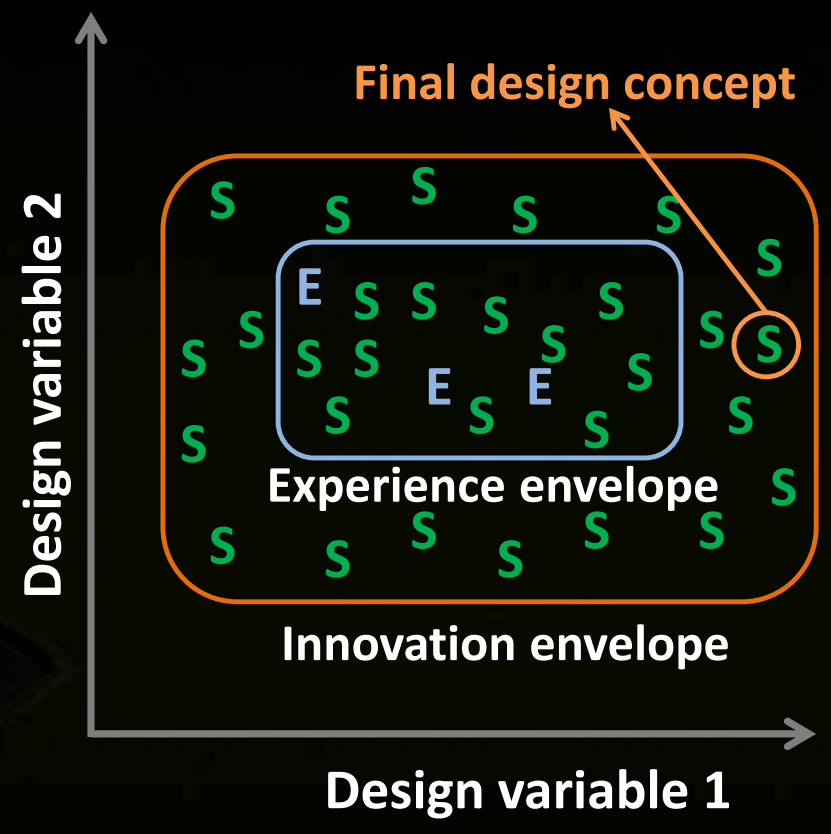

Simulation-driven product innovation

Computing Capacity is Still a Major Challenge

Frequency of limiting size/detail in simulation models due to compute infrastructure or turnaround time limitations

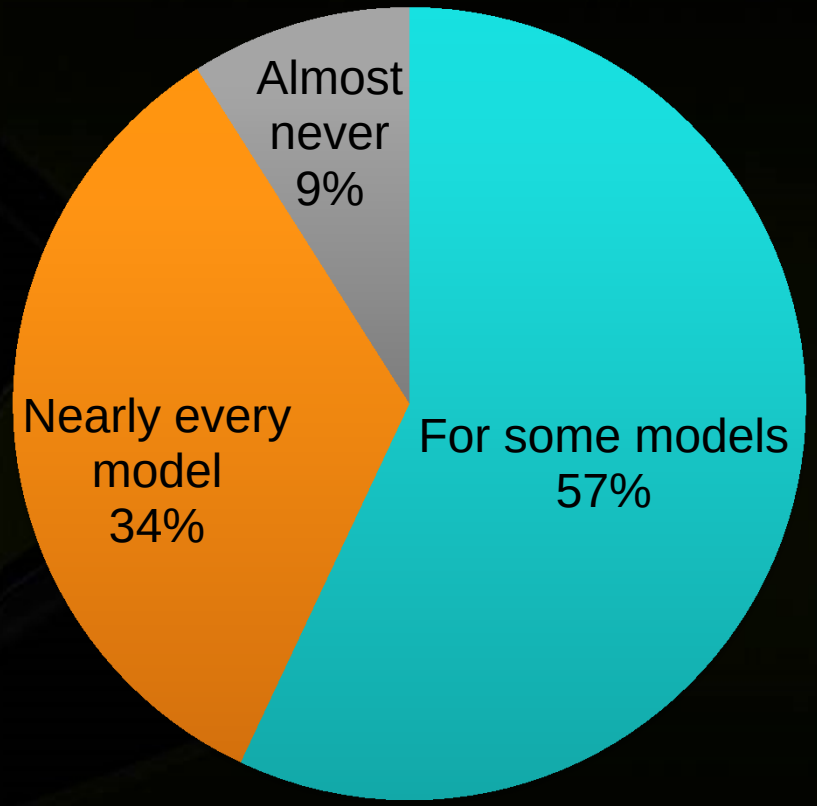

Source: Survey by ANSYS with over 1,800 respondents

Increasing GPU Performance & Memory Bandwidth

Peak Double Precision FP

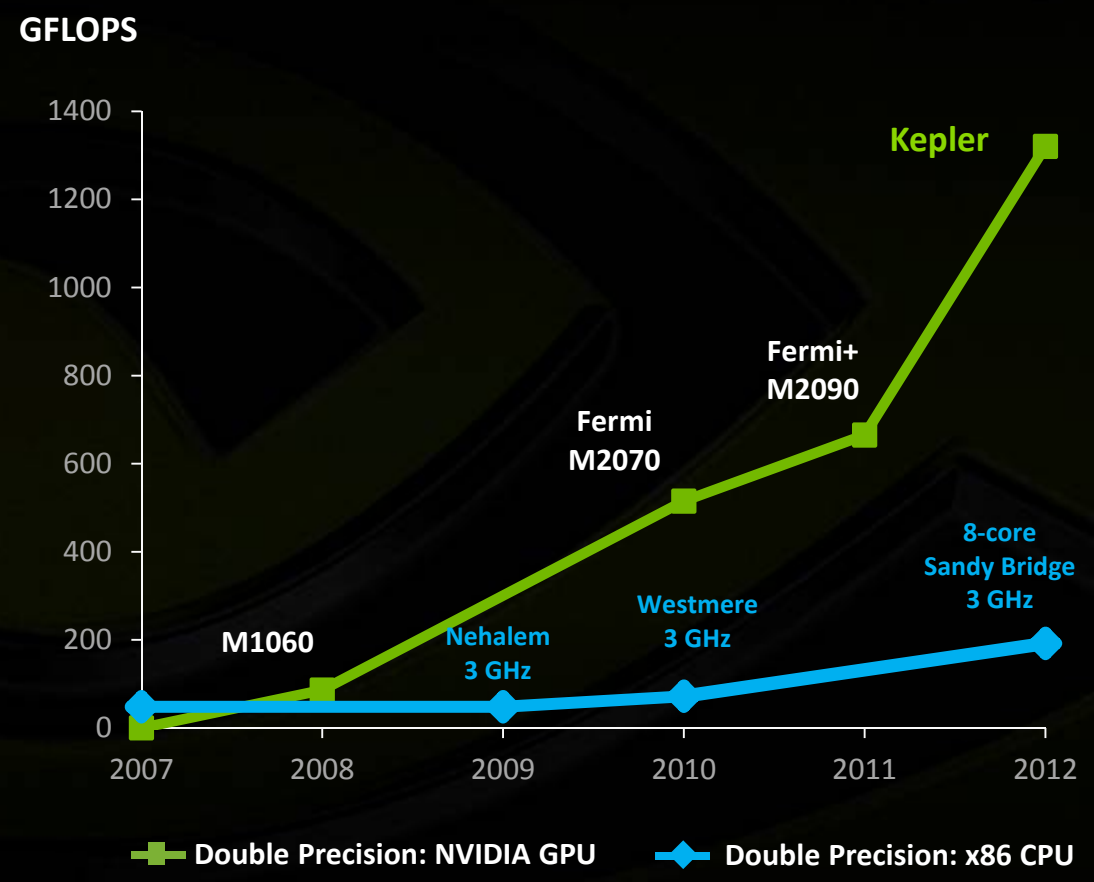

Peak Memory Bandwidth

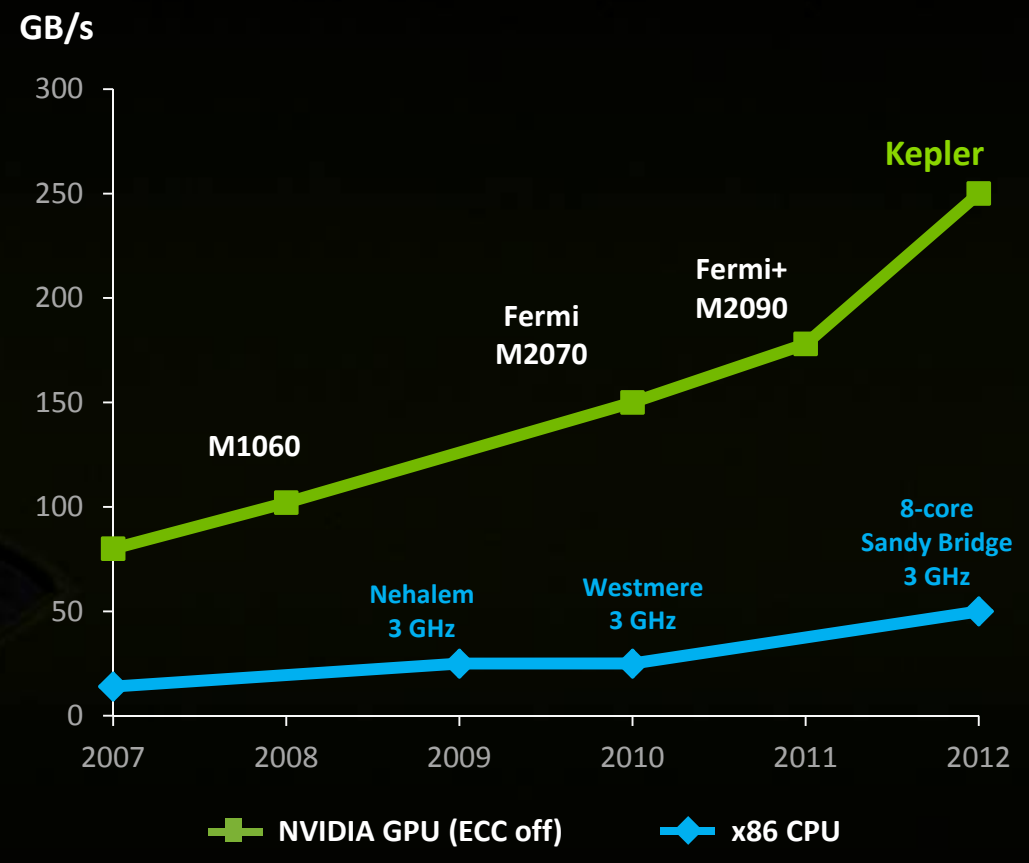

AGENDA

- GPUs in Enterprise Computing
- Business Challenges in Product Development
- NVIDIA GPUs for CAE Applications
- Computational Fluid Dynamics (CFD) Applications
- Computational Structural Mechanics (CSM) Applications
- Computational Electro Magnetics (CEM) Applications
- NVIDIA Solutions
- Value of GPU-accelerated Simulations

Basics of GPU Computing

- **GPU is an accelerator attached to an x86 CPU**
- **GPU acceleration is user-transparent**
  - Jobs launch and complete without additional user steps
- **Schematic of a CPU with an attached GPU accelerator**
  - CPU begins/ends job, GPU manages heavy computations

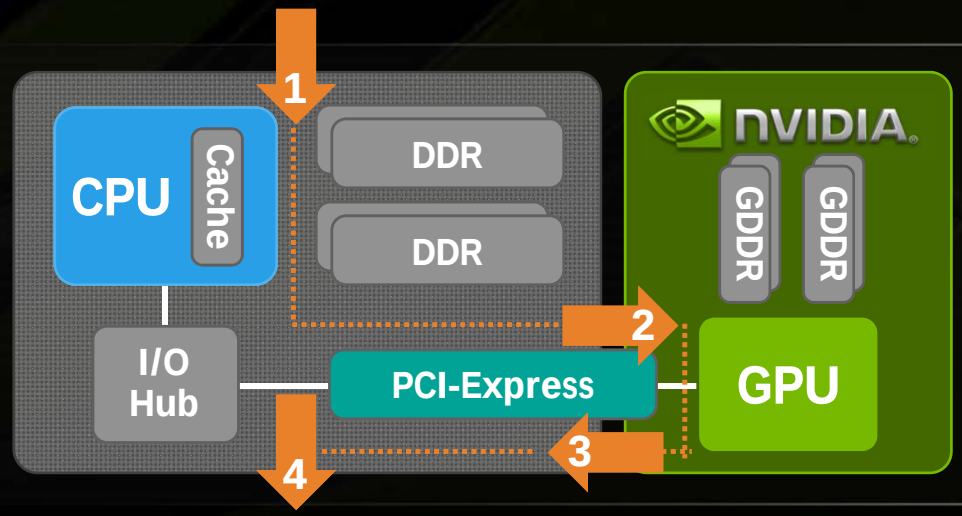

Schematic of an x86 CPU with a GPU accelerator

1. Job launched on CPU
2. Solver operations sent to GPU
3. GPU sends results back to CPU
4. Job completes on CPU

GPU Acceleration of a CAE Application

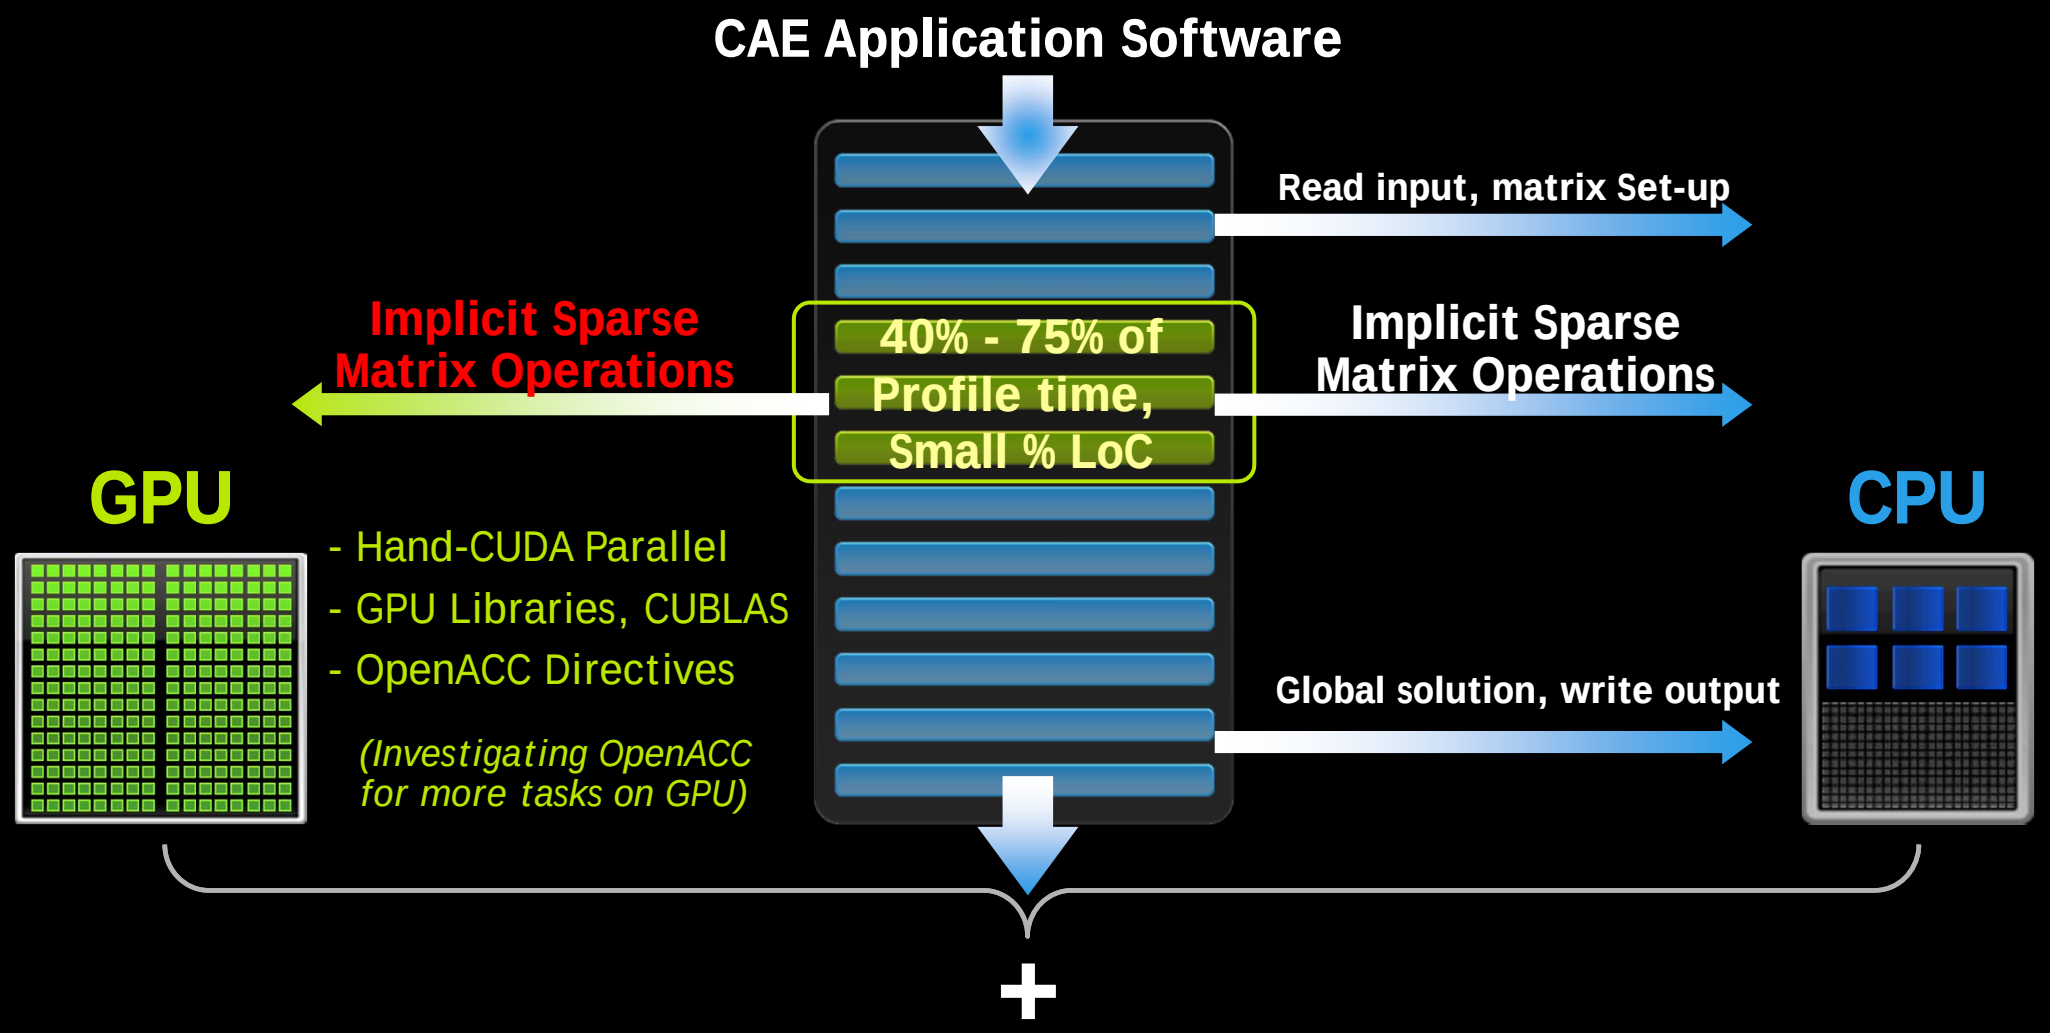

AGENDA

- GPUs in Enterprise Computing
- Business Challenges in Product Development
- NVIDIA GPUs for CAE Applications
- Computational Fluid Dynamics (CFD) Applications
- Computational Structural Mechanics (CSM) Applications
- Computational Electro Magnetics (CEM) Applications
- NVIDIA Solutions
- Value of GPU-accelerated Simulations

GPU-accelerated CFD Applications

APPLICATION	DESCRIPTION	SUPPORTED FEATURES	MULTI-GPU SUPPORT
Altair AcuSolve	General purpose CFD software	Linear equation solver	Yes
ANSYS Fluent	General purpose CFD software	Radiation heat transfer model, linear equation solver	Yes
Autodesk Moldflow	Plastic mold injection software	Linear equation solver	Single only
CPFD Barracuda-VR and Barracuda	Fluidized bed modeling software	Linear equation solver, particle calculations	Single only
FluiDyna Culises for OpenFOAM	Solver library for general purpose CFD software	Linear equation solvers	Yes
FluiDyna LBultra	General purpose CFD software	Lattice-Boltzmann solver	Yes
Prometech Particleworks	Particle-based CFD software	Implicit and explicit solvers	Single only
Vratis ARAEL	General purpose CFD software based on FVM with OpenFOAM compatibility	Linear equation solver	Yes
Turbostream Ltd.	CFD software for turbomachinery flows	Explicit solver	Yes
Vratis SpeedIT extreme for OpenFOAM	Solver library for general purpose CFD software	Linear equation solvers	Yes

ANSYS® Fluent 15.0

GPU Acceleration in Fluent 15.0

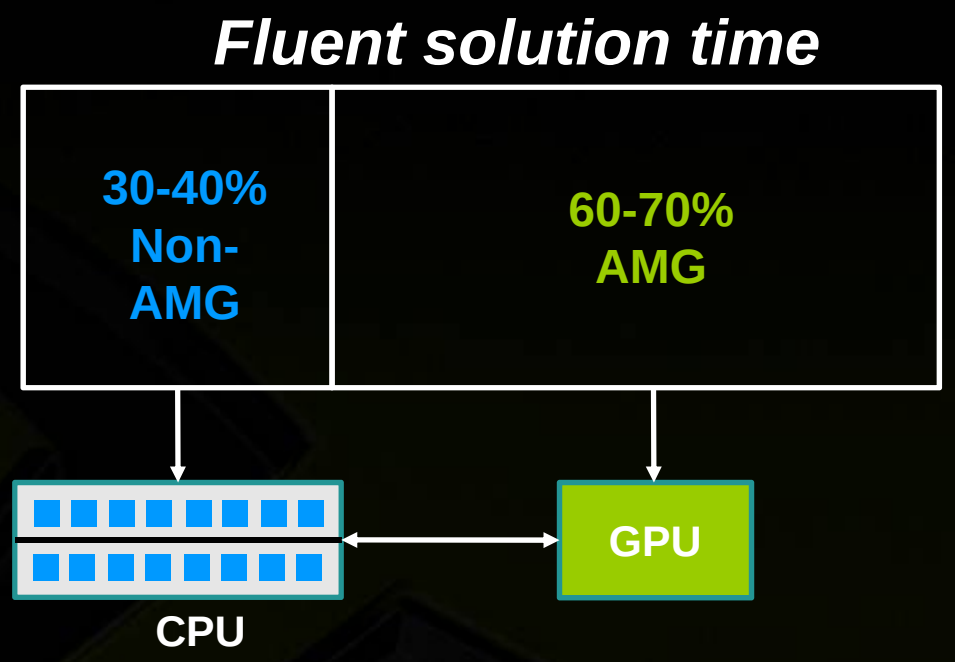

- In flow problems, pressure-based coupled solver typically spends 60-70% time in AMG where as segregated solver only spends 30-40% of the time in AMG
- Higher AMG times are ideal for GPU acceleration, thus coupled-problems benefit from GPUs

Fluent Speed-up from GPU acceleration

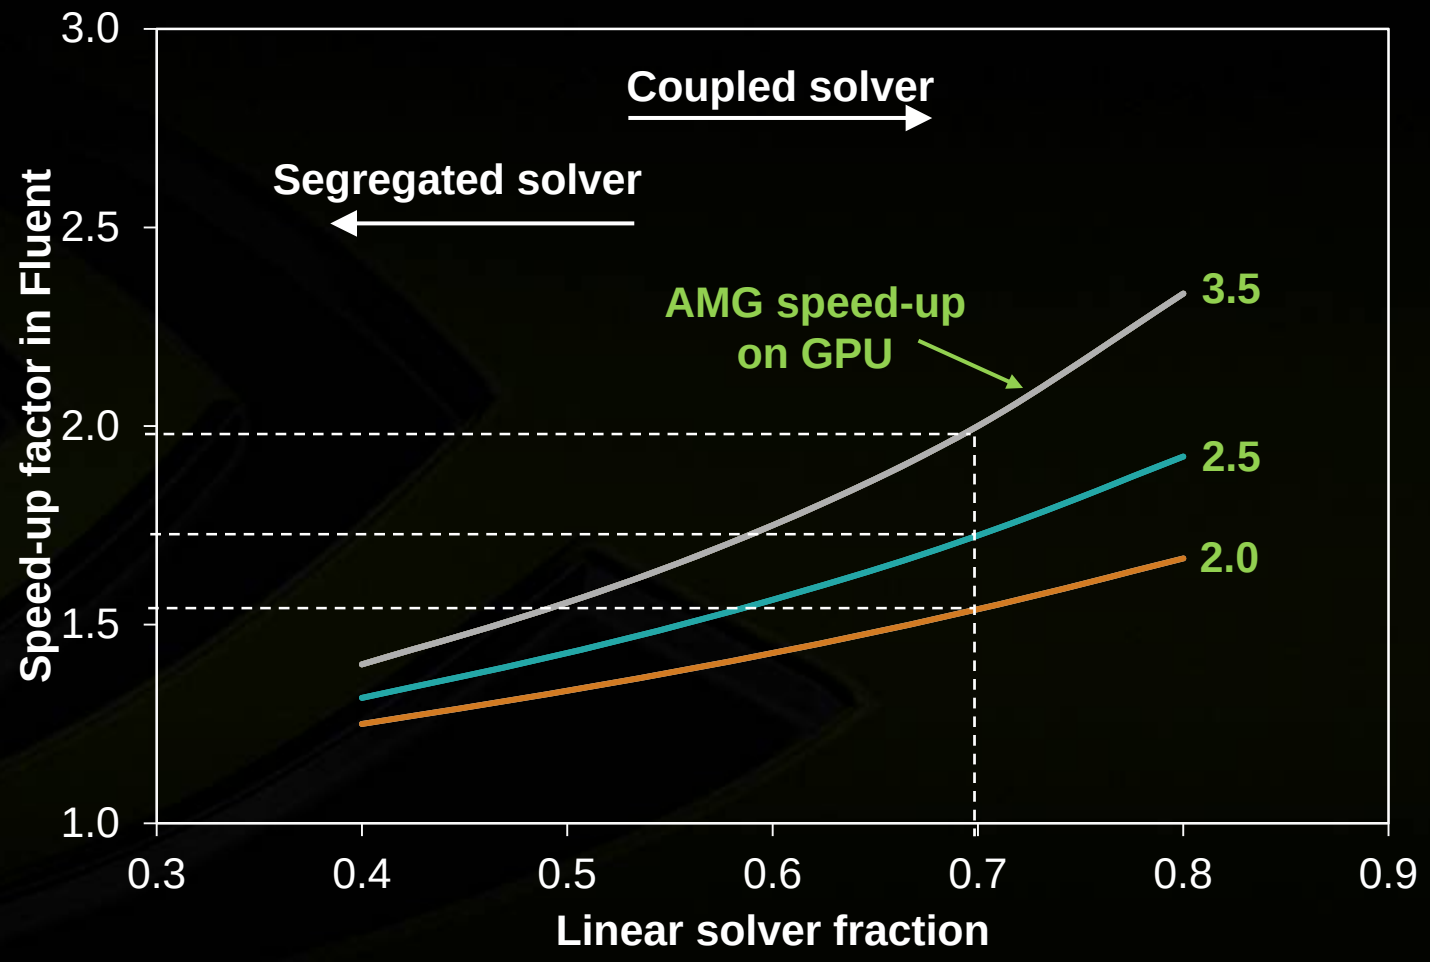

ANSYS 15.0 HPC licenses (new)

- Treats each GPU socket as a CPU core, which significantly increases simulation productivity from HPC licenses
- All ANSYS HPC products unlock GPUs in 15.0, including HPC, HPC Pack, HPC Workgroup, and HPC Enterprise products.

GPU support in ANSYS – Licensing comparison

License type	ANSYS 14.5	ANSYS 15.0
HPC per core licenses	None	1 HPC license – 1 GPU 2 HPC licenses – 2 GPUs 'N' HPC licenses – 'N' GPUs
HPC pack	1 HPC pack – 1 GPU 2 HPC packs – 4 GPUs	1 HPC pack – 4 GPUs 2 HPC packs – 16 GPUs

GPU Acceleration of Water Jacket Analysis

ANSYS Fluent 15.0 performance on pressure-based coupled Solver

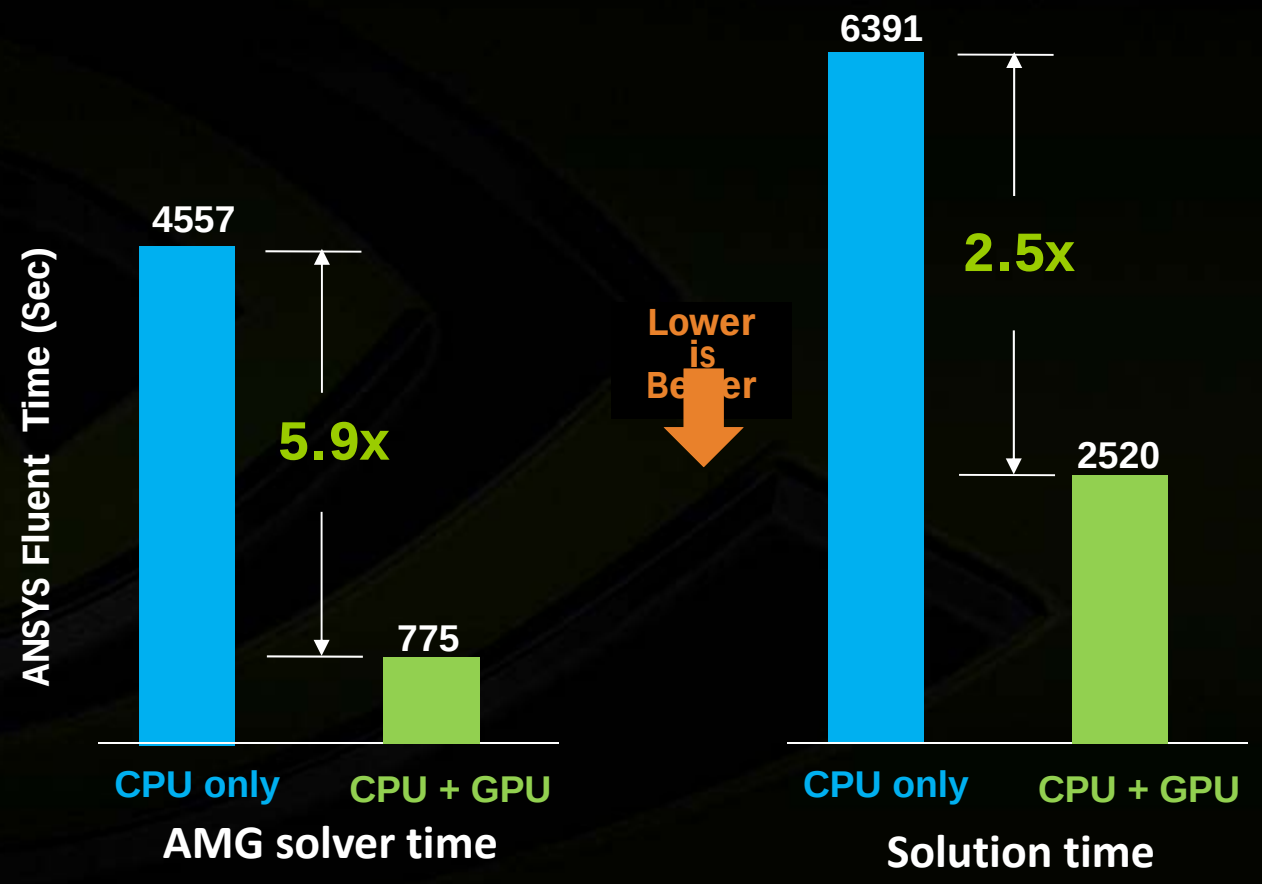

Water jacket model

- Unsteady RANS model
- Fluid: water
- Internal flow
- CPU: Intel Xeon E5-2680; 8 cores
- GPU: 2 X Tesla K40

NOTE: Times for 20 time steps

GPU value proposition for Fluent 15.0

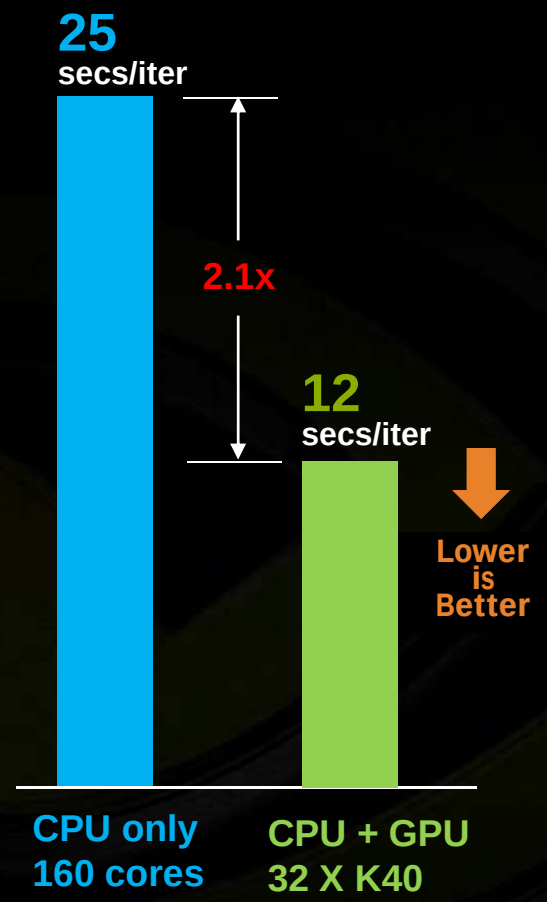

All results are based on turbulent flow over an F1 case (144million cells) over 1000 iterations; steady-state, pressure-based coupled solver with single-precision; CPU: 8 Ivy Bridge nodes with 20 cores each, F-cycle, size 8; GPU: 32 X Tesla K40, V-cycle, and size 2.

Formula 1 aerodynamic study (144 million cells)

2X

Additional productivity from HPC licenses and GPUs

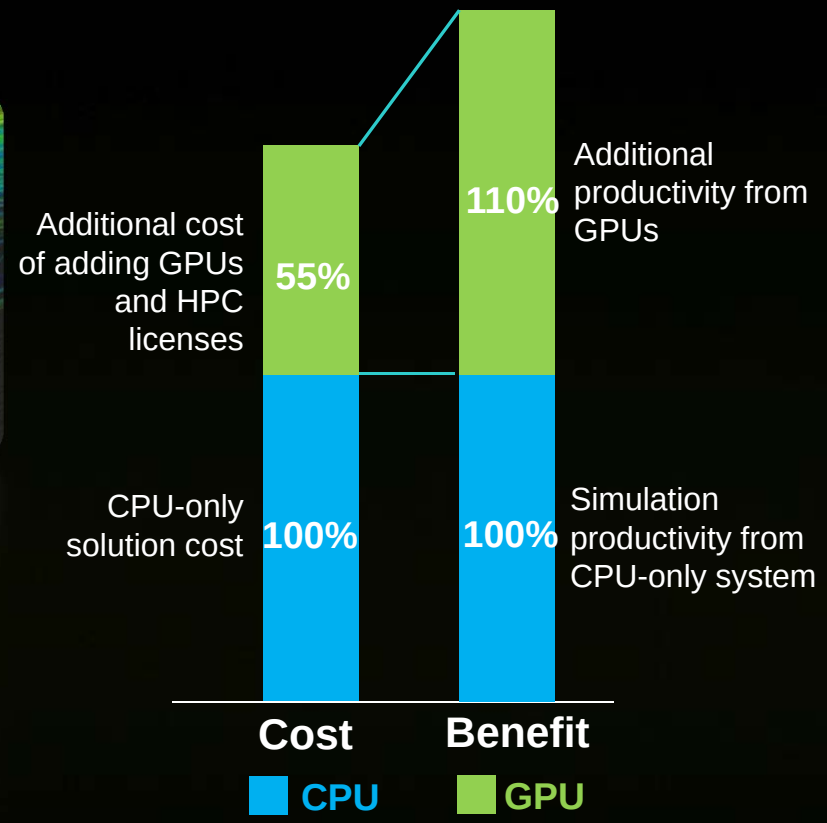

CPU-only solution cost is approximated and includes both hardware and paid-up software license costs. Benefit/productivity is based on the number of completed Fluent jobs/day.

GPU Scaling in Fluent 15.0 (F1 case)

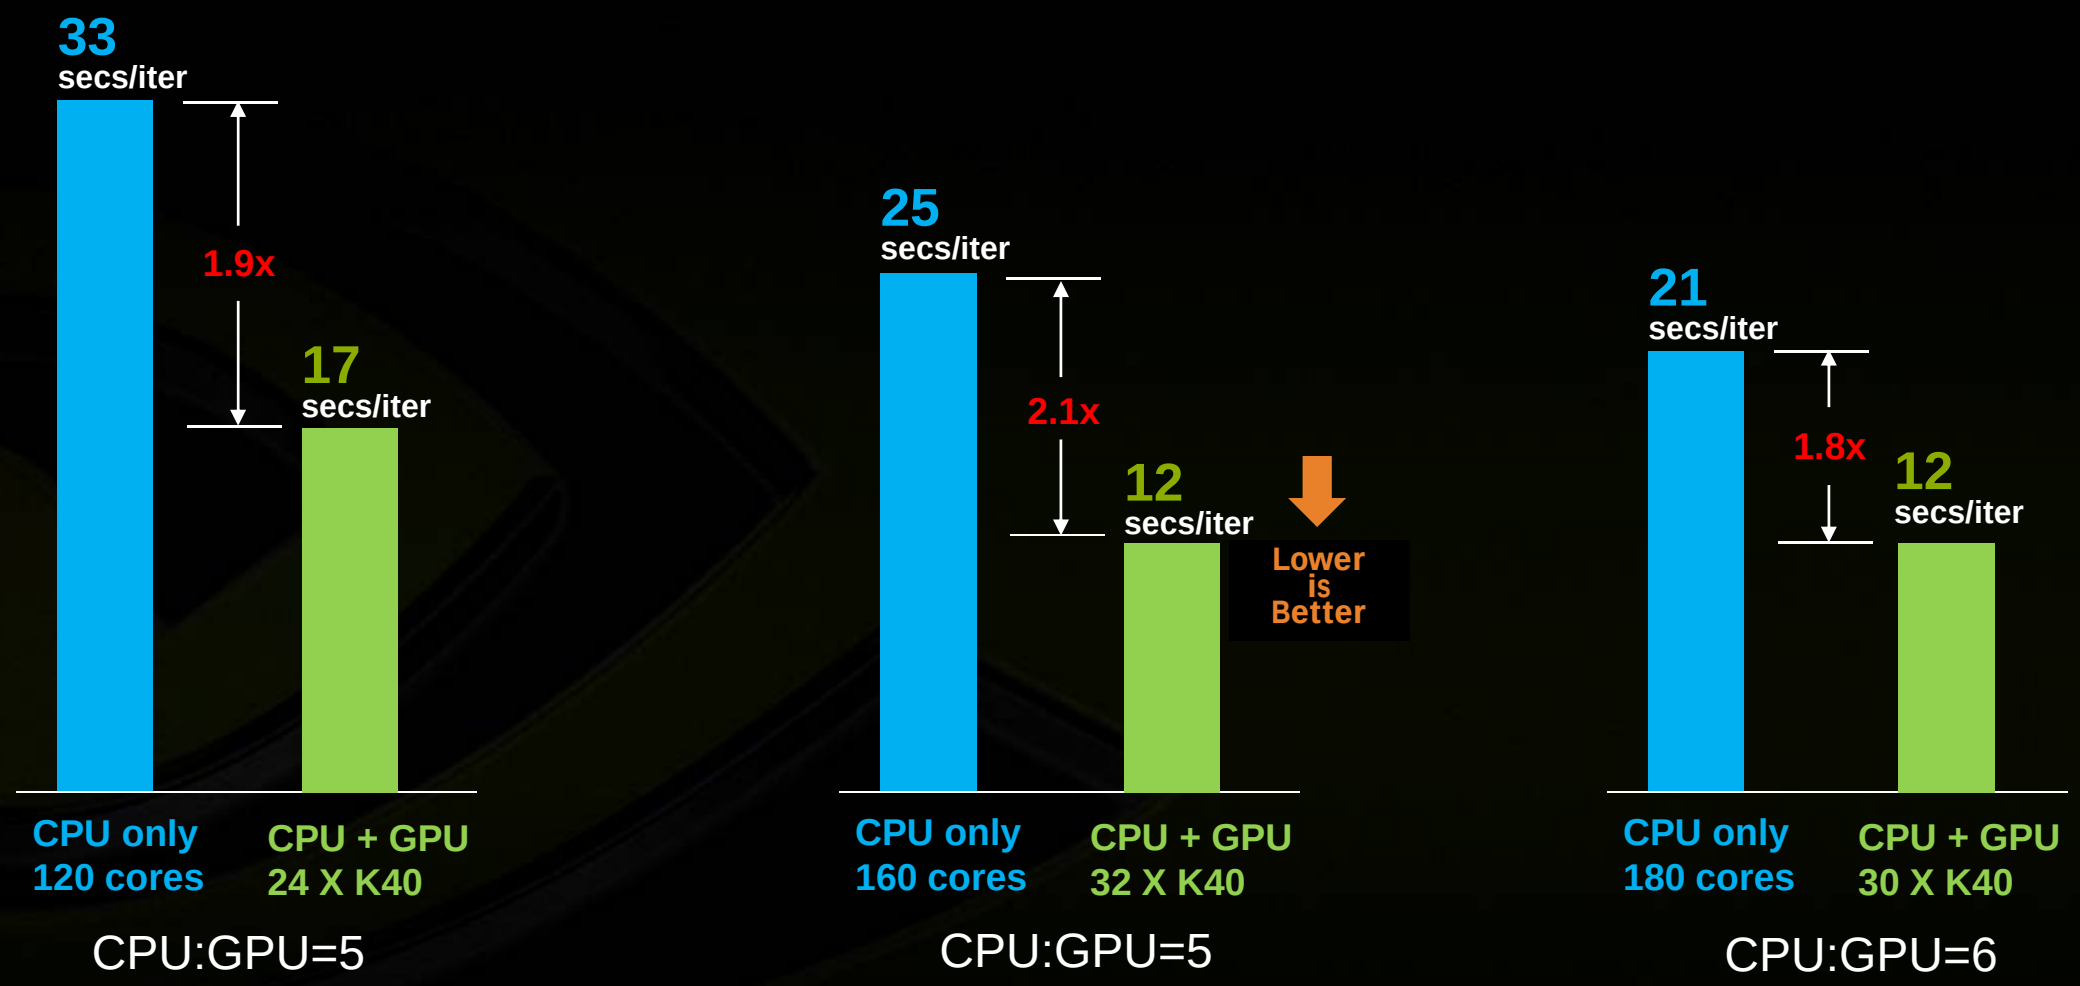

All results are based on turbulent flow over an F1 case (144million cells) over 1000 iterations; steady-state, pressure-based coupled solver with single-precision; CPU: Ivy Bridge with 10 cores per socket, F-cycle, size 8; GPU: Tesla K40, V-cycle, and size 2 .

Shorter Time to Solution with GPUs at PSI Inc.

A customer success story

Objective

Meeting engineering services schedule & budget, and technical excellence are imperative for success.

HPC Solution

- PSI evaluates and implements the new technology in software (ANSYS 15.0) and hardware (NVIDIA GPU) as soon as possible.
- GPU produces a 43% reduction in Fluent solution time on an Intel Xeon E5-2687 (8 core, 64GB) workstation equipped with an NVIDIA K40 GPU

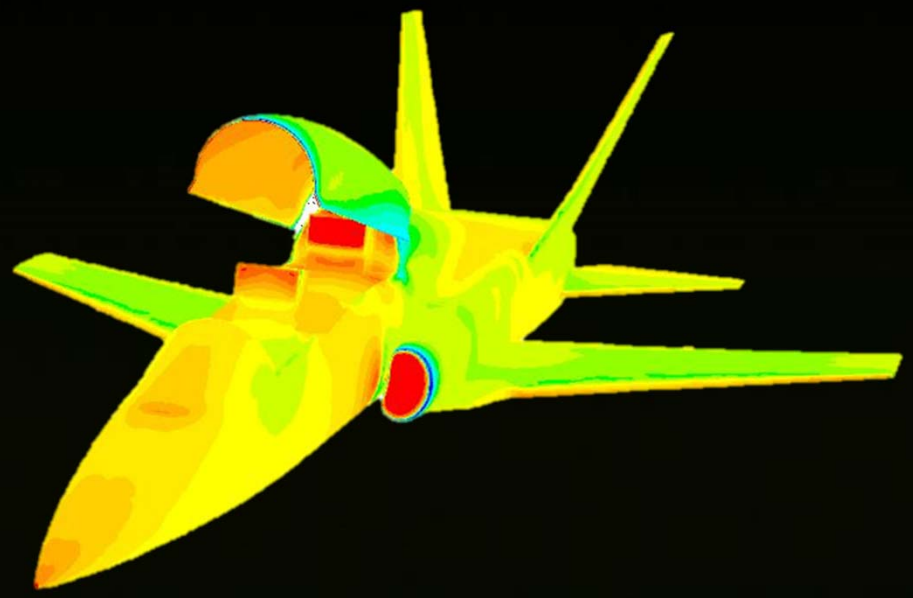

Design Impact

Increased simulation throughput allows meeting delivery-time requirements for engineering services.

Images courtesy of Parametric Solutions, Inc.

ANSYS Fluent 15.0 Resources

- ANSYS solution web page at NVIDIA – http://www.nvidia.com/ansys
- GPU user guide for ANSYS Fluent 15.0 available at - http://www.ansys.com/Resource+Library/Technical+Briefs/Accelerating+ANSYS+Fluent+15.0+Using+NVIDIA+GPUs
- Previously recorded ANSYS IT webcast series webinar titled “How to Speed Up ANSYS 15.0 with NVIDIA GPUs”, which is available at http://www.ansys.com

FluiDyna®
INNOVATIVE ENGINEERING AND BEYOND

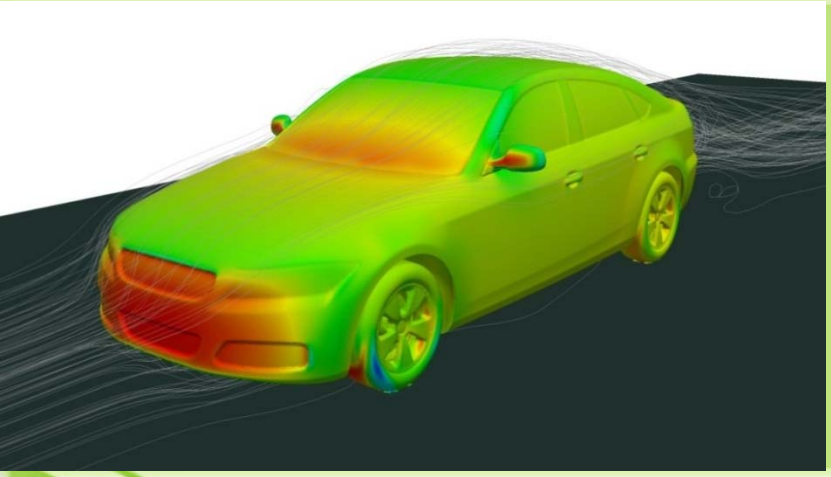

GPU Acceleration of Computational Fluid Dynamics (CFD) in Industrial Applications using Culises and aeroFluidX

Library Culises
Concept and Features

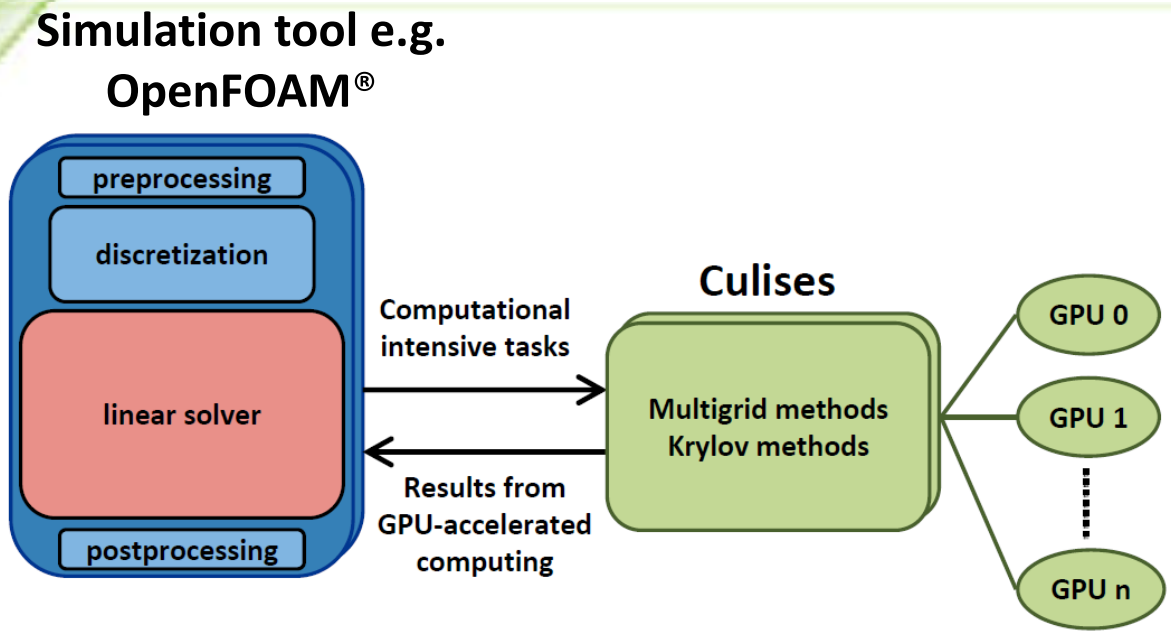

Culises = **Cu**da **Li**brary for **S**olving Linear **E**quation **S**ystems

See also www.culises.com

- State-of-the-art solvers for solution of linear systems
  - Multi-GPU and multi-node capable
  - Single precision or double precision available
- Krylov subspace methods
  - CG, BiCGStab, GMRES for symmetric /non-symmetric matrices
  - Preconditioning options
    - Jacobi (Diagonal)
    - Incomplete Cholesky (IC)
    - Incomplete LU (ILU)
    - Algebraic Multigrid (AMG), see below
- Stand-alone multigrid method
  - Algebraic aggregation and classical coarsening
  - Multitude of smoothers (Jacobi, Gauss-Seidel, ILU etc.)
- Flexible interfaces for arbitrary applications e.g.: established coupling with OpenFOAM®

Culises: Auto OEM Model

Multi-GPU runs

- **Automotive industrial setup (Japanese OEM)**
- CPU linear solver for pressure: geometric-algebraic multigrid (GAMG) of OpenFoam
  GPU linear solver for pressure: AMG preconditioned CG (AMGPCG) of Culises
- 200 SIMPLE iterations

Grid Cells	CPU cores Intel E5-2650	GPUs Nvidia K40	Linear solve time [s]	Total simulation time [s]	Speedup linear solver	Speedup total simulation
18M	8	1	1779	8407	3.83	1.60
18M	8	2	1238	7846	5.50	1.71
18M	16	2	1194	4564	2.50	1.39
62M	16	2	4170	16337	2.62	1.42
62M	32	4	2488	7905	1.90	1.29

Culises

Potential speedup for hybrid approach

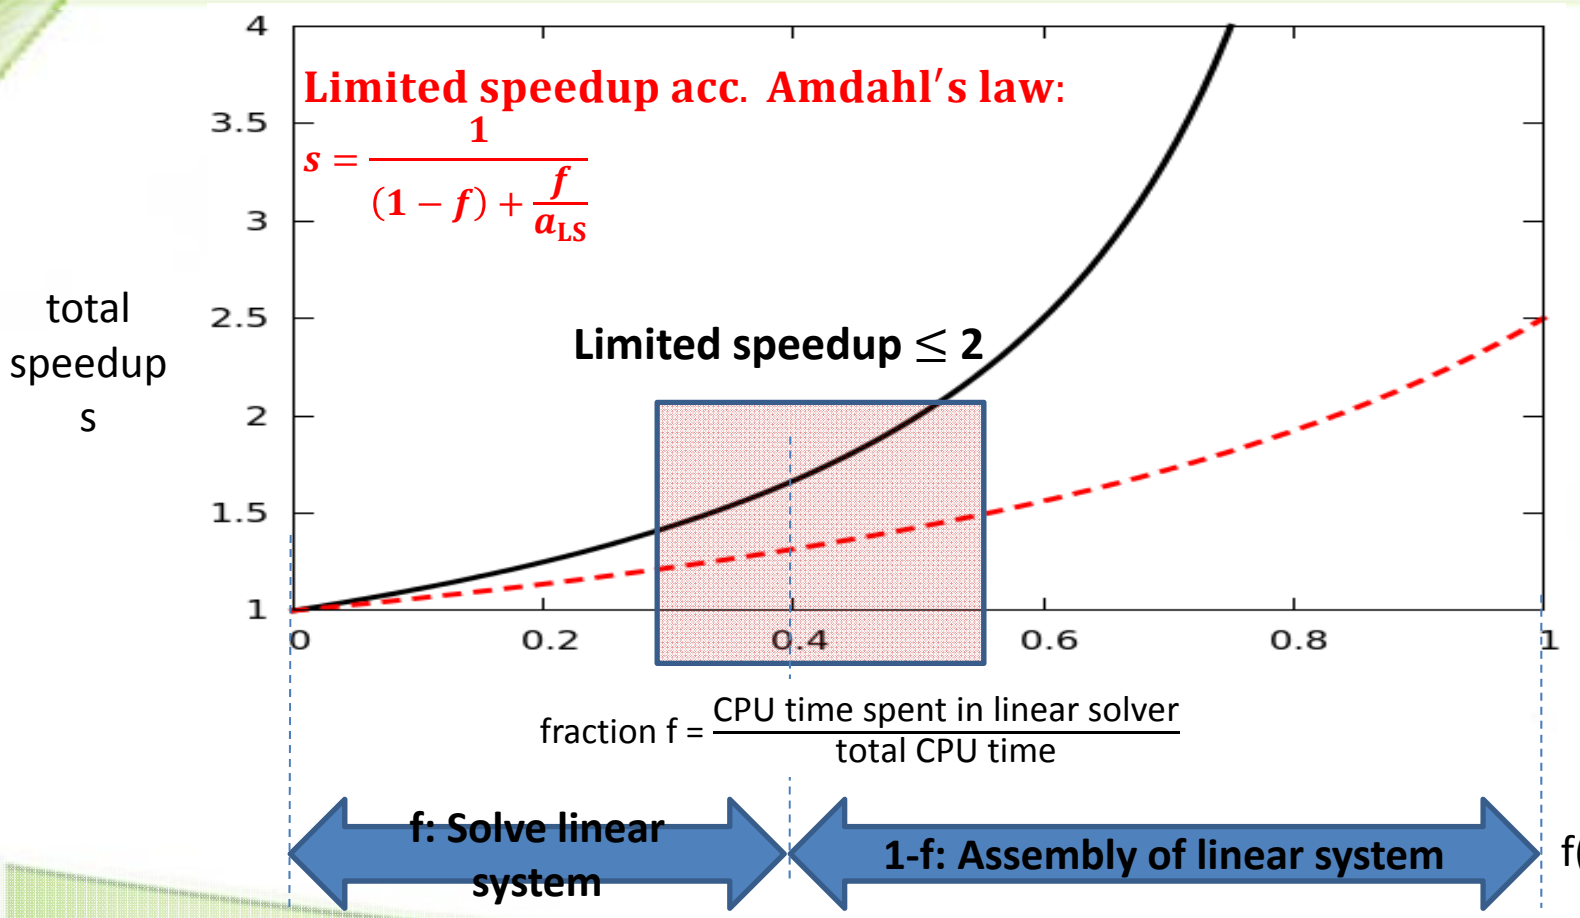

acceleration of linear solver on GPU:

$a_{LS} \rightarrow \infty$, a_{MA} = 1.0

a_{LS} = 2.5, a_{MA} = 1.0

a_{LS} :Speedup linear solver

a_{MA}: Speedup matrix assembly

f(steady-state run) << f(transient run)

aeroFluidX

an extension of the hybrid approach

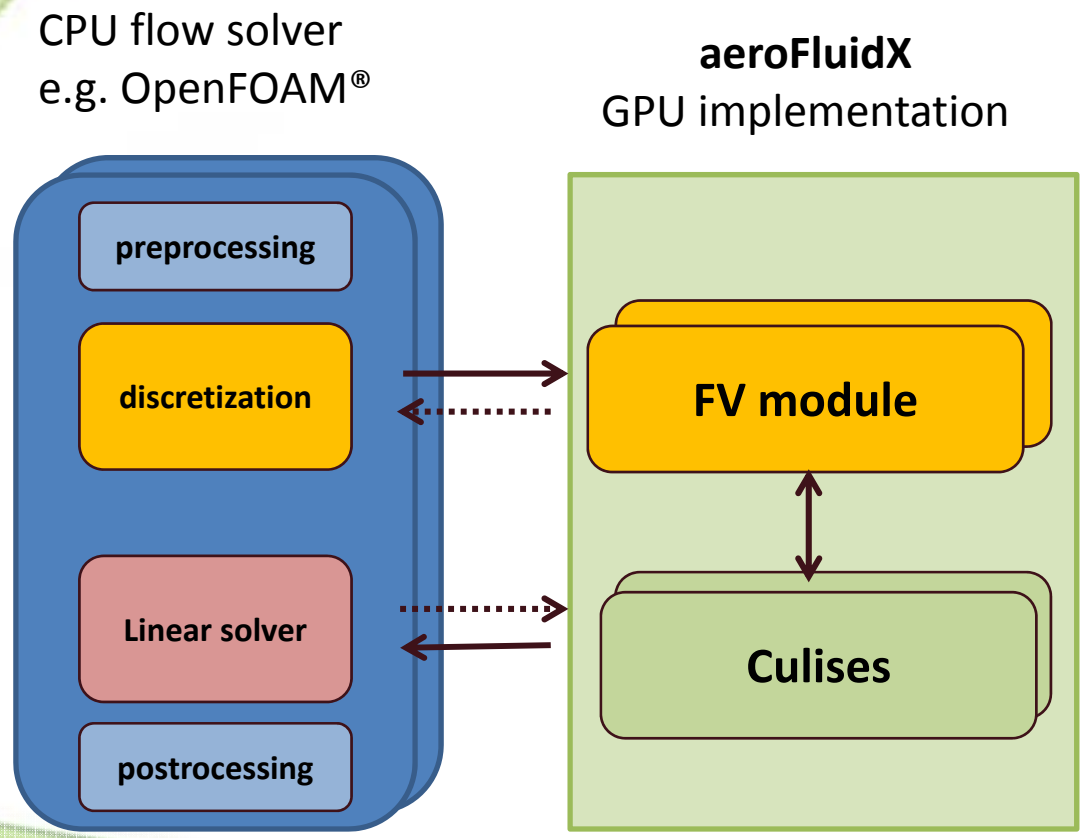

- Porting discretization of equations to GPU
  - ➔ discretization module (Finite Volume) running on GPU
  - ➔ Possibility of direct coupling to Culises
    - ➔ Zero overhead from CPU-GPU-CPU memory transfer and matrix format conversion
    - ➔ Solution of momentum equations also beneficial
- OpenFOAM® environment supported
  - ➔ Enables plug-in solution for OpenFOAM® customers
  - ➔ But communication with other input/output file formats possible

aeroFluidX

NACA0012 airfoil flow

- CPU: Intel E5-2650 (all 8 cores)
  GPU: Nvidia K40
- 4M grid cells (unstructured)
- Running 100 SIMPLE steps with:
  - OpenFOAM® (OF)
    - pressure: GAMG
    - Velocitiy: Gauss-Seidel
  - OpenFOAM® (OFC)
    - Pressure: Culises AMGPCG (1.5x)
    - Velocity: Gauss-Seidel
  - aeroFluidX (AFXC)
    - Pressure: Culises AMGPCG
    - Velocity: Culises Jacobi
- Total speedup:
  - OF (1x)
  - OFC 1.22x
  - AFXC 1.82x

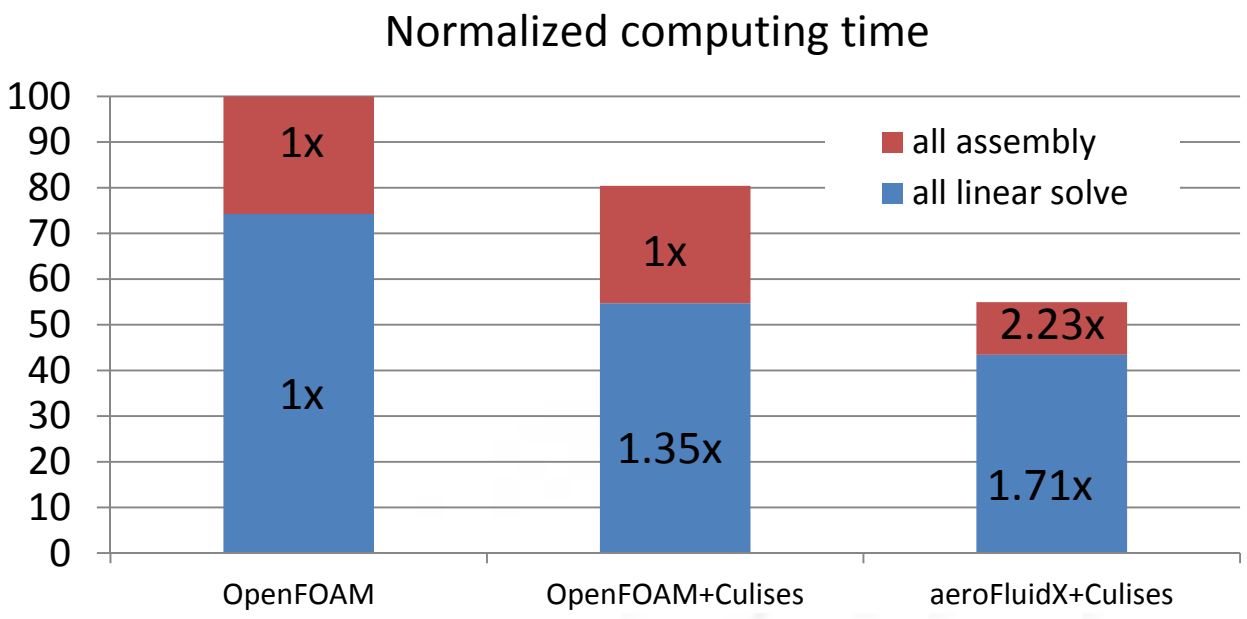

all assembly = assembly of all linear systems (pressure and velocity)
all linear solve = solution of all linear systems (pressure and velocity)

aeroFluidX
Release Planning

FluiDyna® INNOVATIVE ENGINEERING AND BEYOND

2014 | 2015 | 2016

V0.98

- Steady-state laminar flow
- Single-GPU
- Speedup* > 2x

V1.0

- Turbulent flow (RANS)
  - k-omega (SST)
  - Spalart-Allmaras
- Multi-GPU
  - Single-node
  - Multi-node

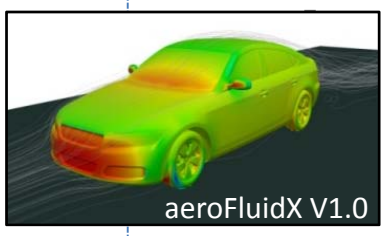

V1.2

- Unsteady flow
- Advanced turbulence modelling (LES/DES)
- Speedup* 2-3x

V2.0

- Basic support for moving geometries (MRF)
- Porous media
- Advanced model for rotating devices (sliding mesh approach)

* Speedup against standard OpenFoam

Summary

- **Culises** - hybrid approach for accelerated CFD applications (OpenFOAM)
  - **General applicability** for industrial cases including various existing flow models
  - Significant speedup (≥ 2x) of linear solver employing GPUs
  - **Moderate speedup** (≤ 1.6x) of total simulation
  - Culises V1.1 released: Commercial and academic licensing available
    Free testing & benchmarking opportunities at FluiDyna GPU-servers
- **aeroFluidX** - fully ported flow solver on GPU to harvest full GPU computing power
  - General applicability requires **rewrite of large portion of existing code**
  - Steady-state, incompressible **unstructured multigrid** flow solver established & validated
  - Significant speedup (≥ 2x) of matrix assembly; **without full code tuning/optimization!**
  - **Enhanced speedup** for total simulation

AGENDA

- GPUs in Enterprise Computing
- Business Challenges in Product Development
- NVIDIA GPUs for CAE Applications
- Computational Fluid Dynamics (CFD) Applications
- Computational Structural Mechanics (CSM) Applications
- Computational Electro Magnetics (CEM) Applications
- NVIDIA Solutions
- Value of GPU-accelerated Simulations

GPU-accelerated CSM Applications

APPLICATION	DESCRIPTION	SUPPORTED FEATURES	MULTI-GPU SUPPORT
Abaqus/Standard	Simulation and analysis tool for structural mechanics	Direct sparse solver	Yes
ANSYS Mechanical	Simulation and analysis tool for structural mechanics	Direct and iterative solvers	Yes
Impetus Afea	Predicts large deformations of structures and components exposed to extreme loading conditions	Linear equation solver	Yes
LS-DYNA Implicit	Multiphysics simulation package used	Linear equation solver	Yes
MSC Nastran	Simulation and analysis tool for structural mechanics	Direct sparse solver	Yes
MSC Marc	Simulation and analysis tool for structural mechanics	Direct sparse solver	Yes
OptiStruct	Simulation and analysis tool for structural mechanics	Direct and iterative solvers	Yes
DS Exsight	Uses Abaqus Standard for GPU computing	Direct sparse solver	Single only
DS DesignSight	Uses Abaqus Standard for GPU computing	Direct sparse solver	Single only
PAM-CRASH Implicit	Multiphysics simulation package used	Linear equation solver	Single only
NX Nastran	Simulation and analysis tool for structural mechanics	Linear equation solver	Single only

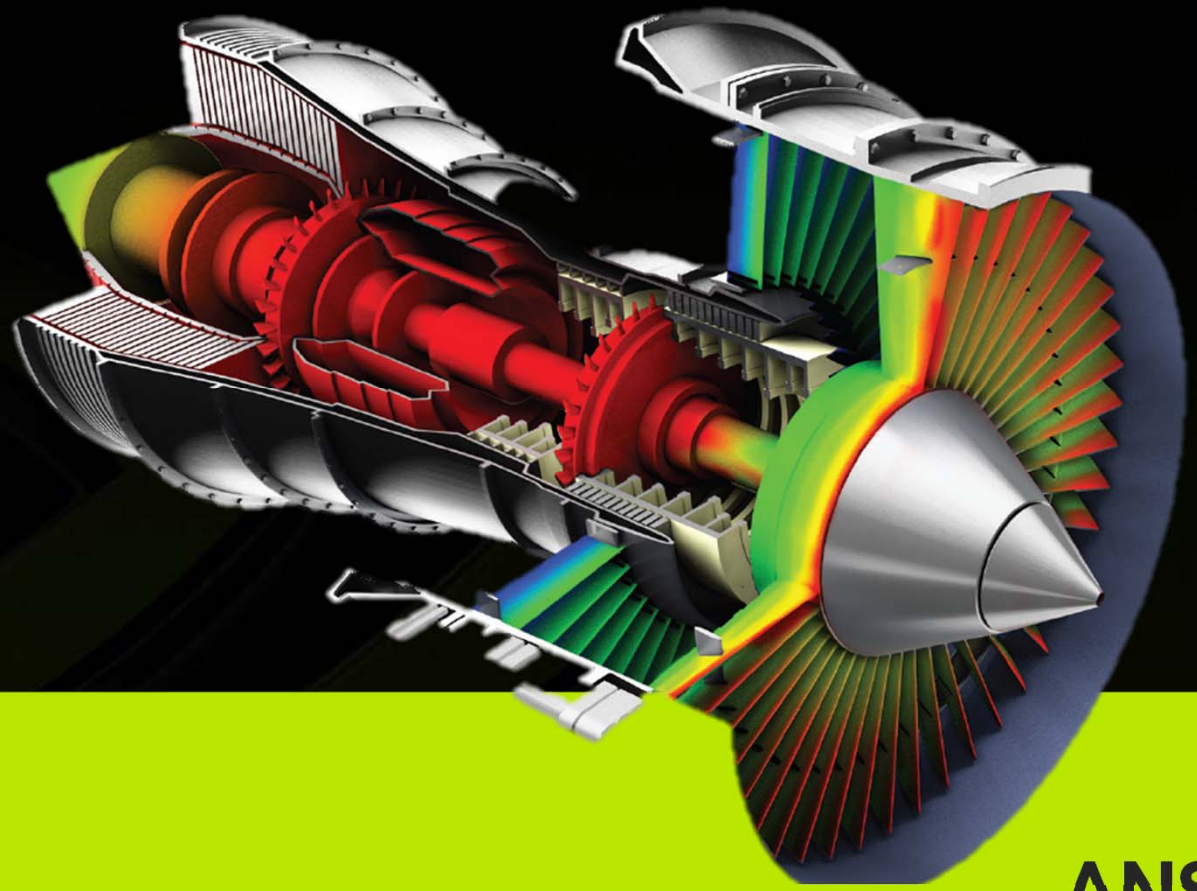

ANSYS® Mechanical15.0

GPU Acceleration in Mechanical 15.0

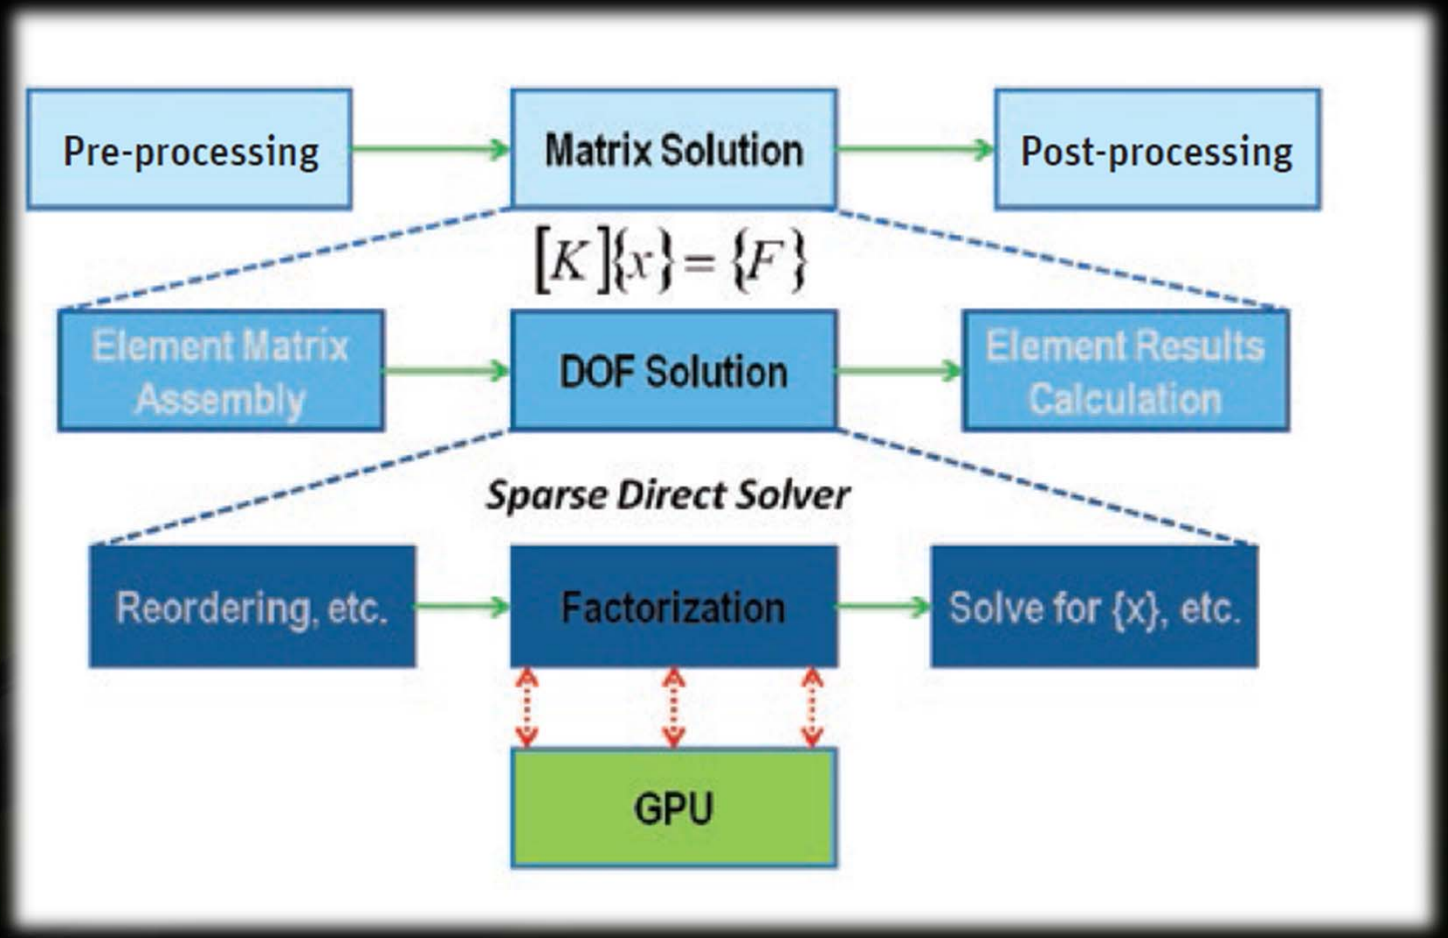

Source: Accelerating Mechanical Solutions with GPUs by Sheldon Imaka, ANSYS Advantage Volume VII, Issue 3, 2013

ANSYS Mechanical15.0 on Tesla K40

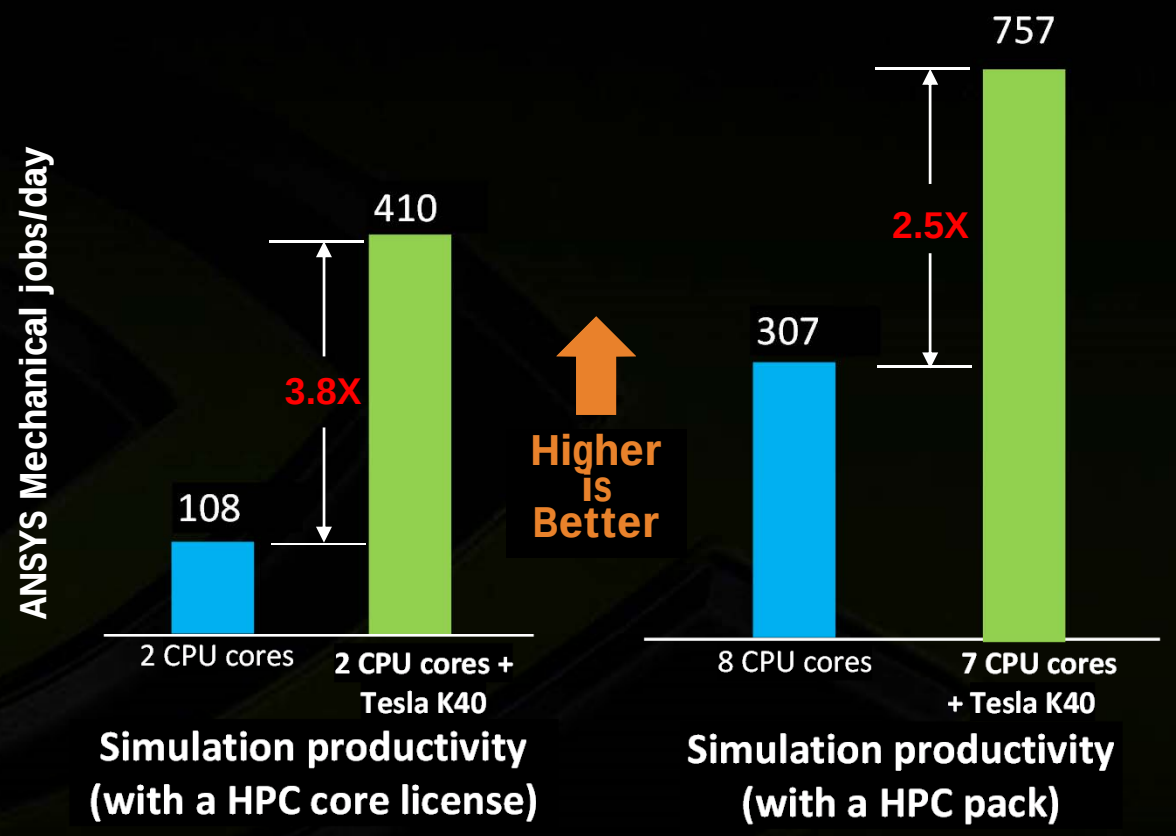

Distributed ANSYS Mechanical 15.0 with Sandy Bridge (Xeon E5-2687W 3.1 GHz) 8-core CPU and a Tesla K40 GPU with boost clocks; V145sp-5 model, Turbine geometry, 2.5 M DOF, and direct sparse solver.

V14sp-5 Model

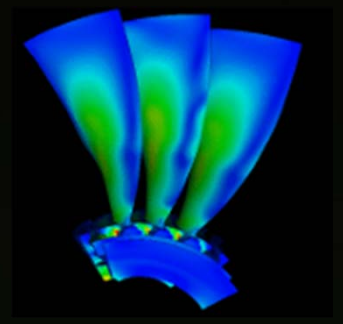

- Turbine geometry
- 2,100,000 DOF
- SOLID187 FEs
- Static, nonlinear
- Distributed ANSYS 15.0
- Direct sparse solver

GPU value proposition for Mechanical15.0

With one HPC license + Tesla K20

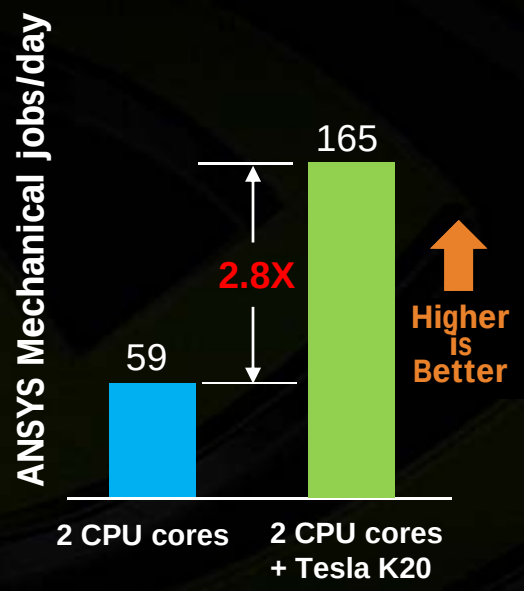

Simulation productivity

V14sp-6 benchmark, 4.9 M DOF, static non-linear analysis; direct sparse solver, distributed ANSYS Mechanical 15.0 with Intel Xeon E5-2697 v2 2.7 GHz CPU; Tesla K20 GPU.

V14sp-6 Model

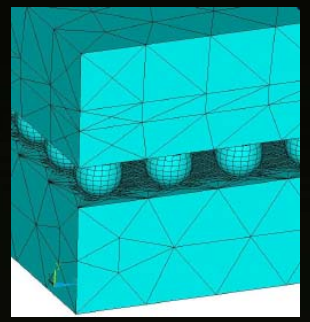

7X

Additional productivity from a GPU and one HPC license

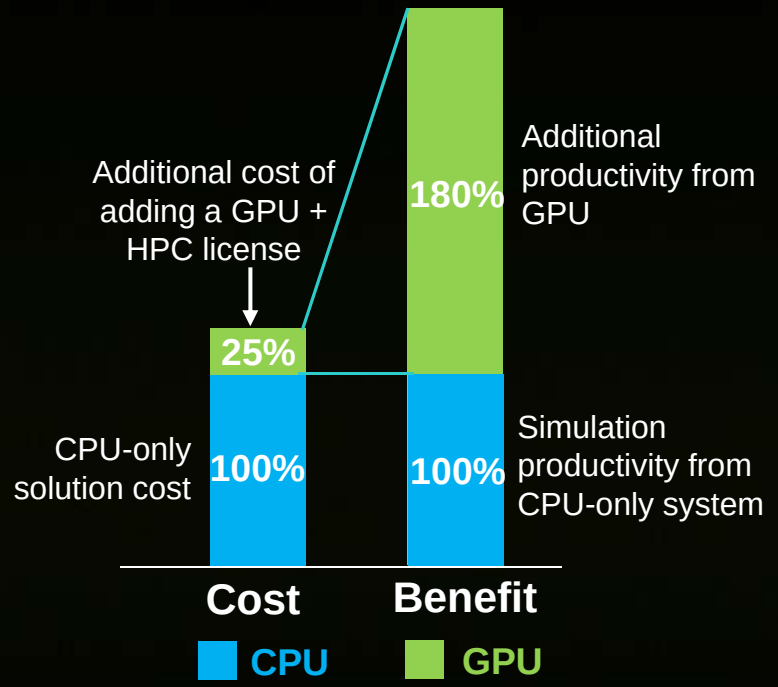

CPU-only solution cost is approximated and includes both hardware and software license costs. Benefit is based on the number of completed Mechanical jobs/day.

GPU value proposition for Mechanical15.0

With an HPC pack + Tesla K20

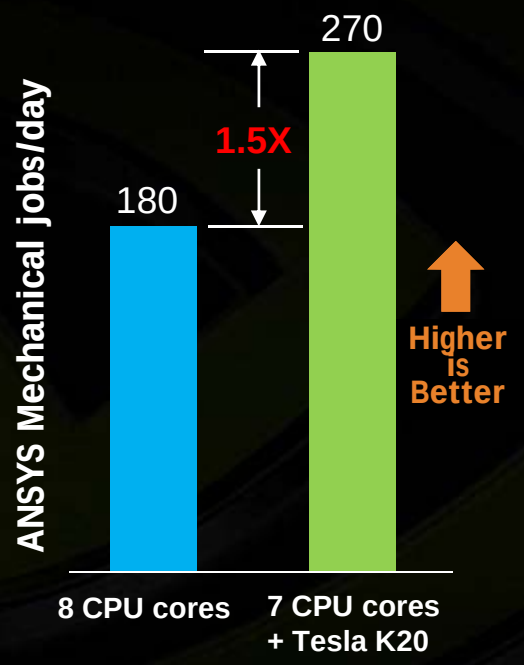

Simulation productivity

V14sp-6 benchmark, 4.9 M DOF, static non-linear analysis; direct sparse solver, distributed ANSYS Mechanical 15.0 with Intel Xeon E5-2697 v2 2.7 GHz CPU; Tesla K20 GPU.

V14sp-6 Model

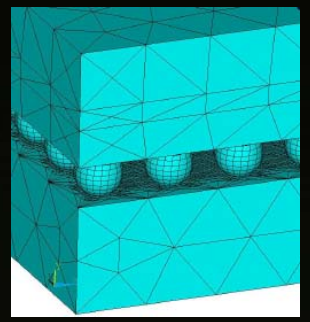

4X

Additional productivity from a GPU and HPC pack

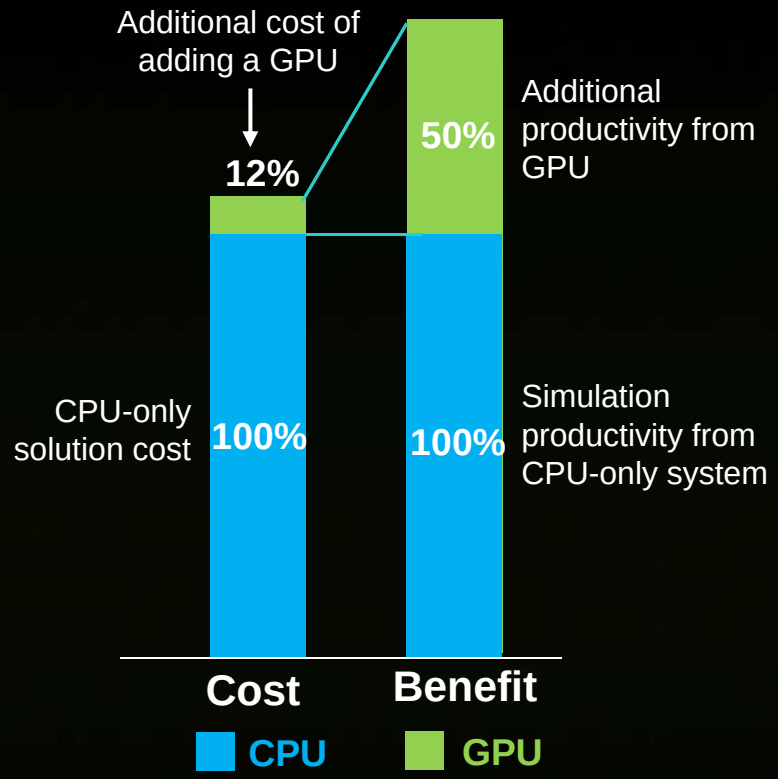

CPU-only solution cost is approximated and includes both hardware and software license costs. Benefit is based on the number of completed Mechanical jobs/day.

ANSYS Mechanical 15.0 Success Story: PSI Inc.

- ANSYS V14.5.7
- Windows 7 Pro SP1
- 5.8 Million DOF
- 8 Cores (Xeon E5-2687)
- 64 GB Ram
- NVIDIA C2075 GPU
- Solution Time :
  - CPU-only: ~ 4 hours
  - With GPU: ~ 2 hours

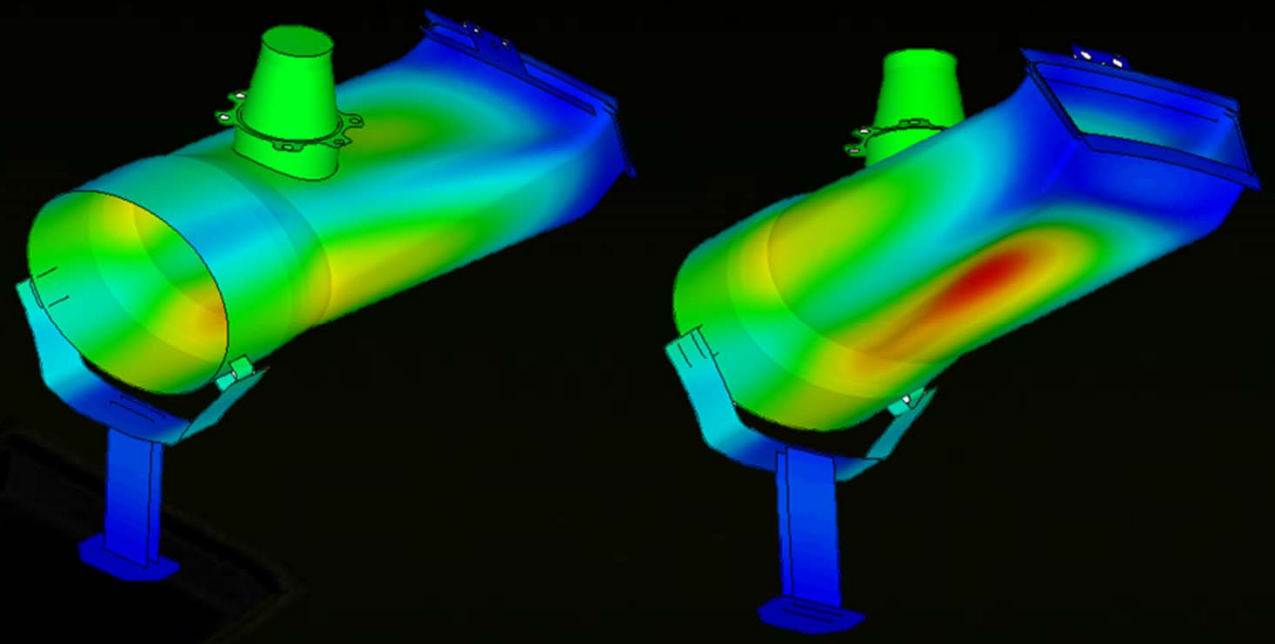

Images courtesy of Parametric Solutions, Inc.

GPU Produces a 50% Reduction in Solution Time

ANSYS Mechanical 15.0 Resources

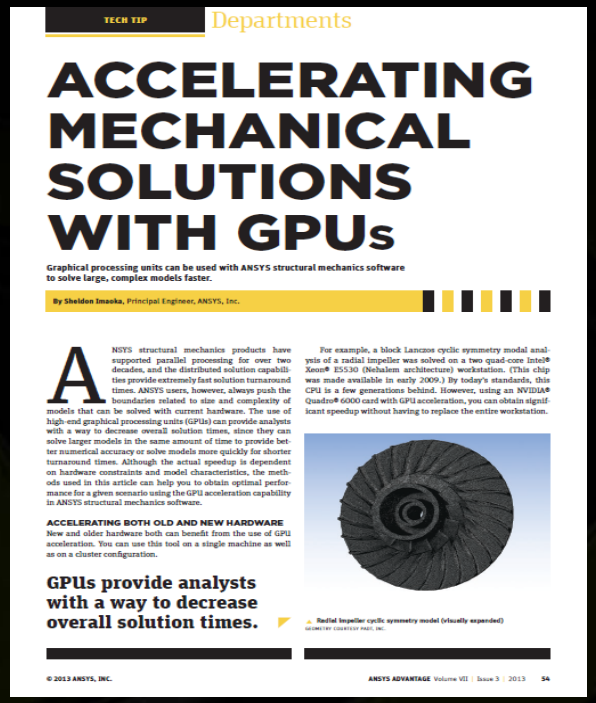

TECH TIP Departments

ACCELERATING MECHANICAL SOLUTIONS WITH GPUs

Graphical processing units can be used with ANSYS structural mechanics software to solve large, complex models faster.

By Sheldon Imaoka, Principal Engineer, ANSYS, Inc.

ANSYS structural mechanics products have supported parallel processing for over two decades, and the distributed solution capabilities provide extremely fast solution turnaround times. ANSYS users, however, always push the boundaries related to size and complexity of models that can be solved with current hardware. The use of high-end graphical processing units (GPUs) can provide analysts with a way to decrease overall solution times, since they can solve larger models in the same amount of time to provide better numerical accuracy or solve models more quickly for shorter turnaround times. Although the actual speedup is dependent on hardware constraints and model characteristics, the methods used in this article can help you to obtain optimal performance for a given scenario using the GPU acceleration capability in ANSYS structural mechanics software.

ACCELERATING BOTH OLD AND NEW HARDWARE

New and older hardware both can benefit from the use of GPU acceleration. You can use this tool on a single machine as well as on a cluster configuration.

GPUs provide analysts with a way to decrease overall solution times.

For example, a block Lanczos cyclic symmetry modal analysis of a radial impeller was solved on a two quad-core Intel® Xeon® E5530 (Nehalem architecture) workstation. (This chip was made available in early 2009.) By today's standards, this CPU is a few generations behind. However, using an NVIDIA® Quadro® 6000 card with GPU acceleration, you can obtain significant speedup without having to replace the entire workstation.

Radial Impeller cyclic symmetry model (visually expanded)

GEOMETRY COURTESY PADT, INC.

© 2013 ANSYS, INC.

ANSYS ADVANTAGE Volume VII | Issue 3 | 2013 54

ANSYS ADVANTAGE Volume VII | Issue 3 | 2013

- ANSYS solution web page at NVIDIA – http://www.nvidia.com/ansys
- ANSYS Advantage magazine article titled "Accelerating Mechanical Solutions with GPUs" for download at http://www.ansys.com
- Previously recorded ANSYS IT webcast series webinar titled "How to Speed Up ANSYS 15.0 with NVIDIA GPUs", which is available at http://www.ansys.com

SIMULIA Abaqus with NVIDIA GPUs

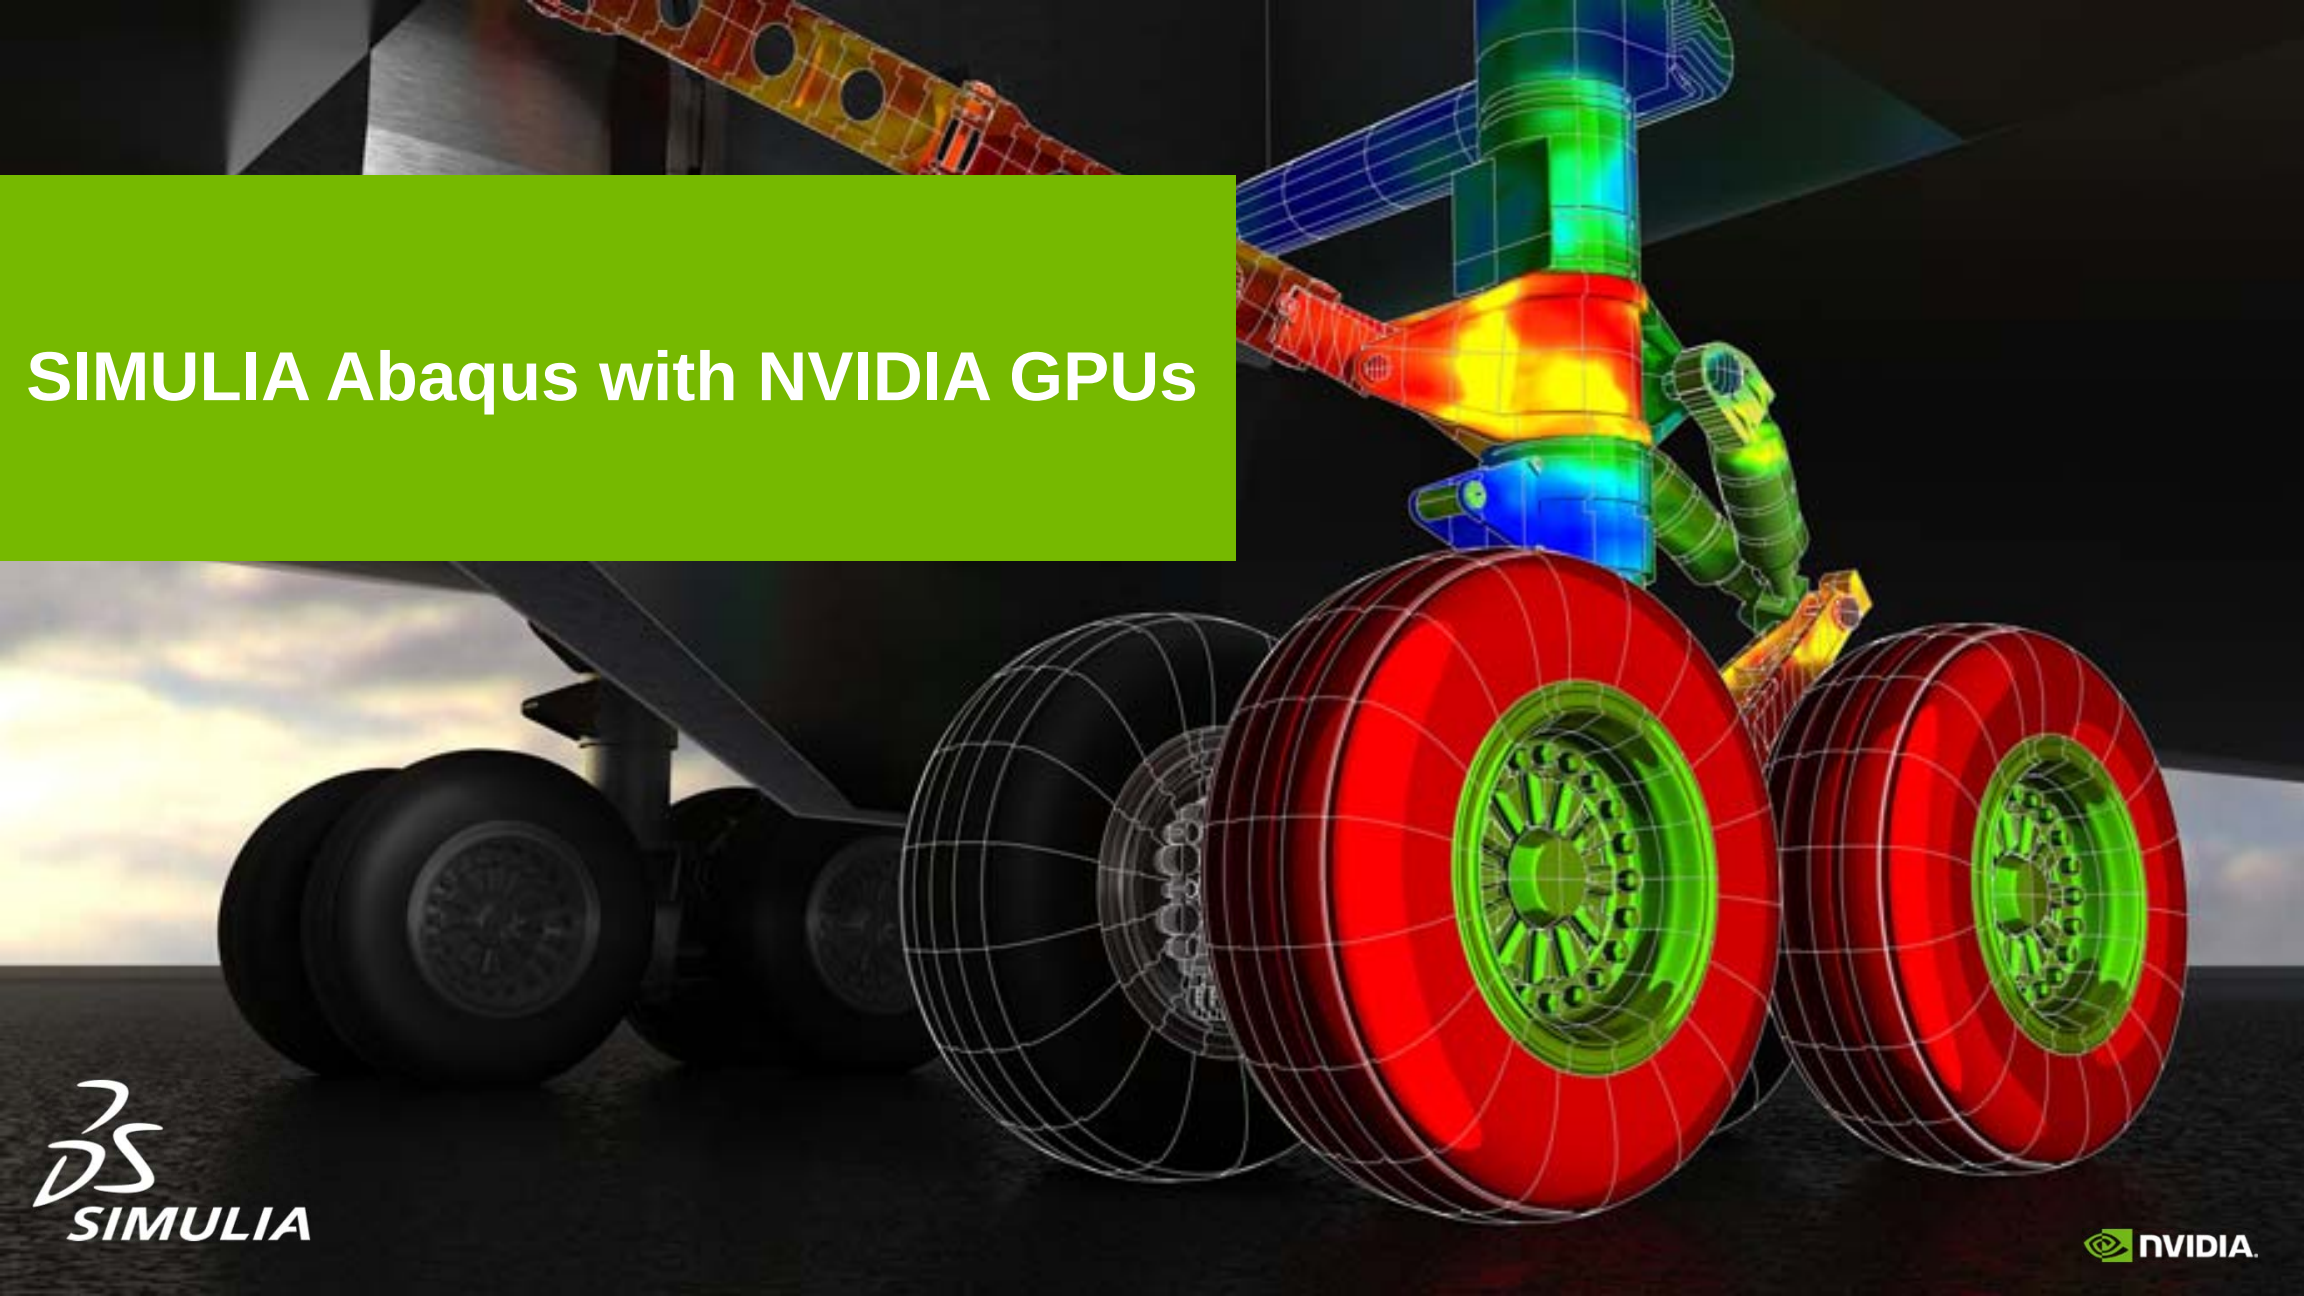

Abaqus/Standard GPU Computing

- **Abaqus 6.11, June 2011**
  - Direct sparse solver is accelerated on single GPU
- **Abaqus 6.12, June 2012**
  - Multi-GPU/node; multi-node DMP clusters
- **Abaqus 6.13, June 2013**
  - Un-symmetric sparse solver on GPU
  - Official Kepler (Tesla K20/K20X) support
- **Abaqus 6.14, July 2014***
  - Direct Sparse Solver
    - Relaxation of memory requirements for GPU
    - Improved performance / DMP split
  - AMS Eigensolver
    - GPUs used in the AMS reduced eigen solution phase. Note: Relevant only for models with ~10,000 or more modes

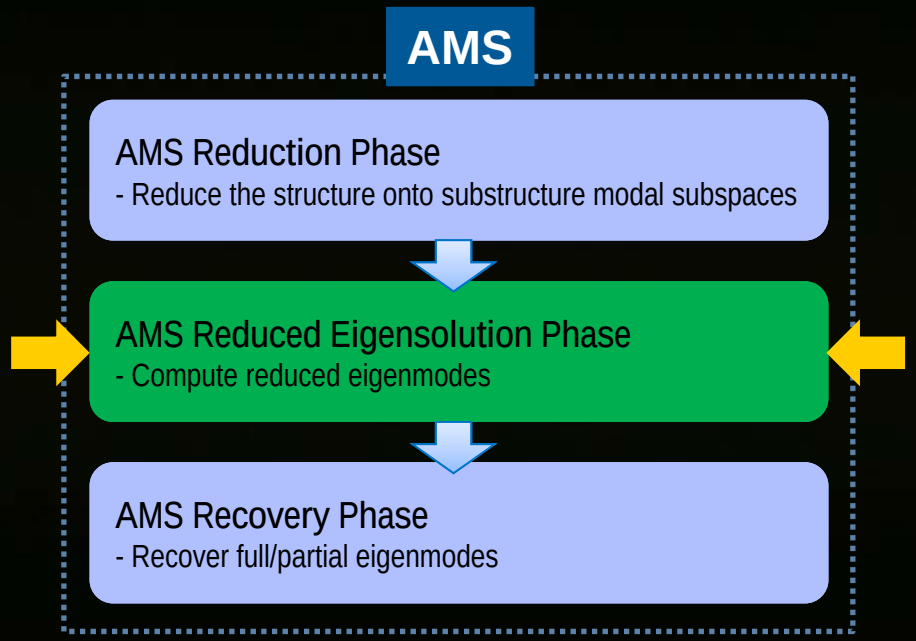

Abaqus Performance with GPU

Customer: Rolls Royce

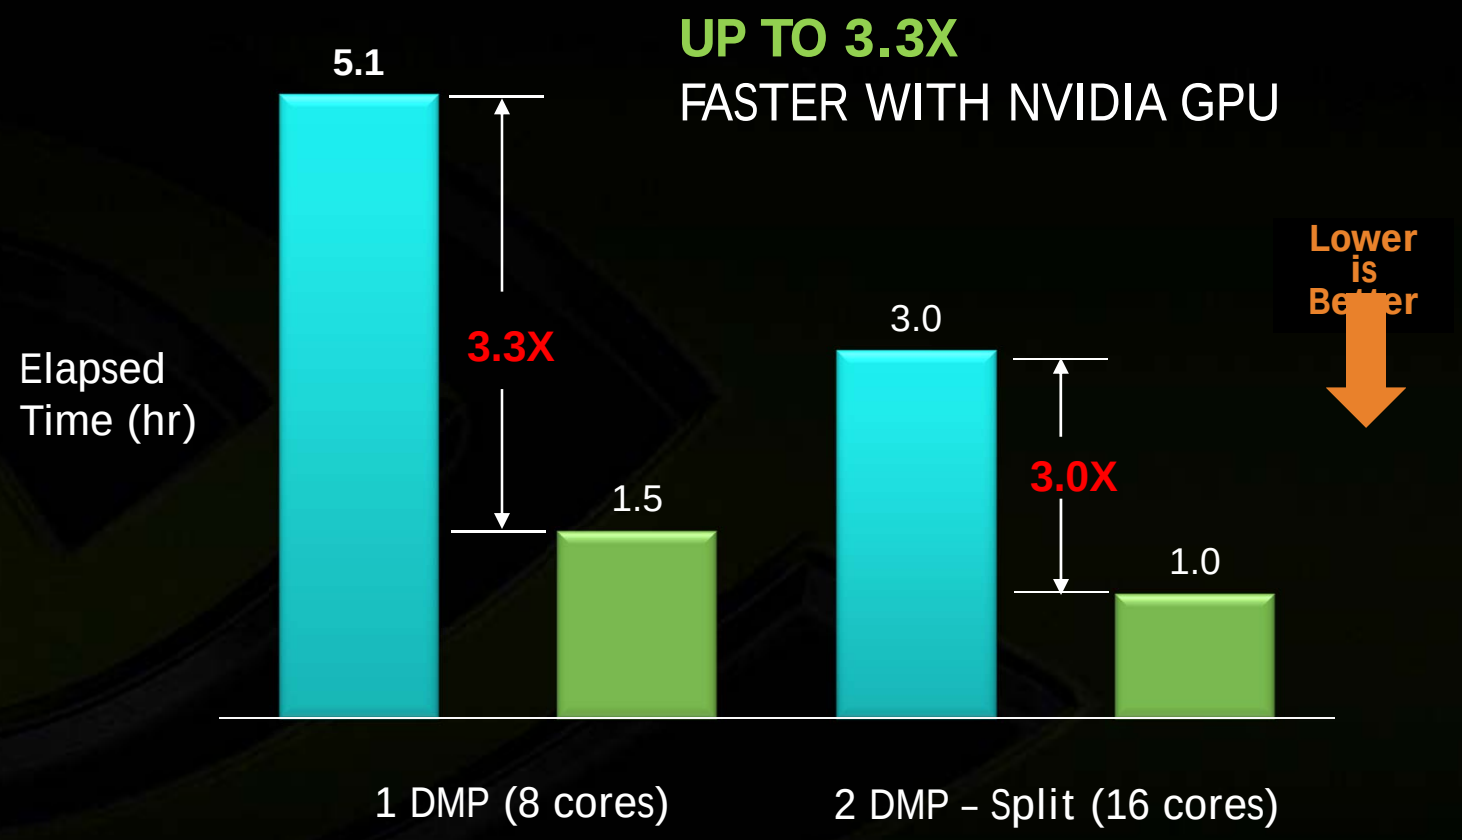

Large Model (~77 TFLOPs), 4.71M DOF, Nonlinear Static, Direct Sparse Solver
Abaqus 6.14-PR2 with Intel Xeon E5-2690v2, 3.0 GHz CPU, 128 GB memory; Tesla K20x GPU

Abaqus Performance with GPU

Customer: Rolls Royce

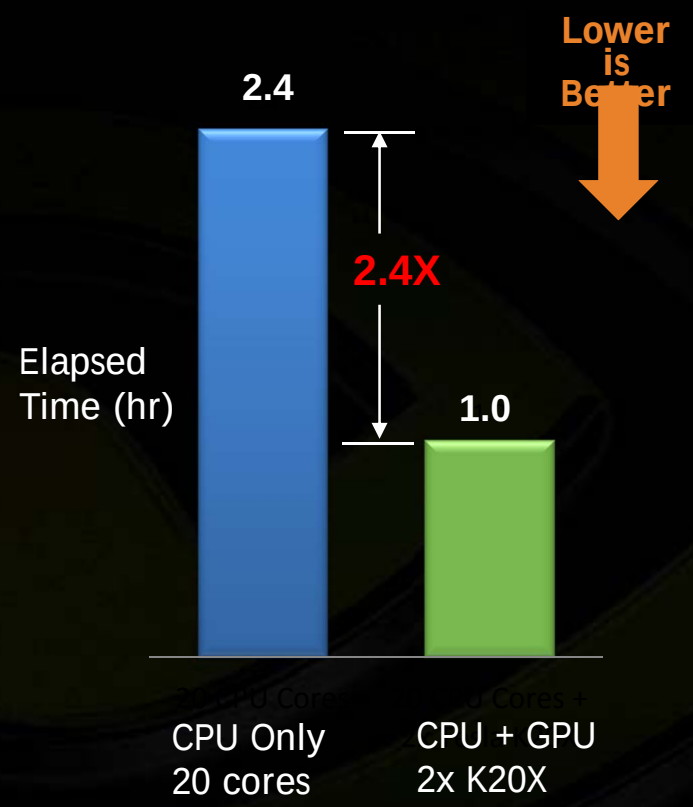

9X

Additional productivity/$ spent on GPUs

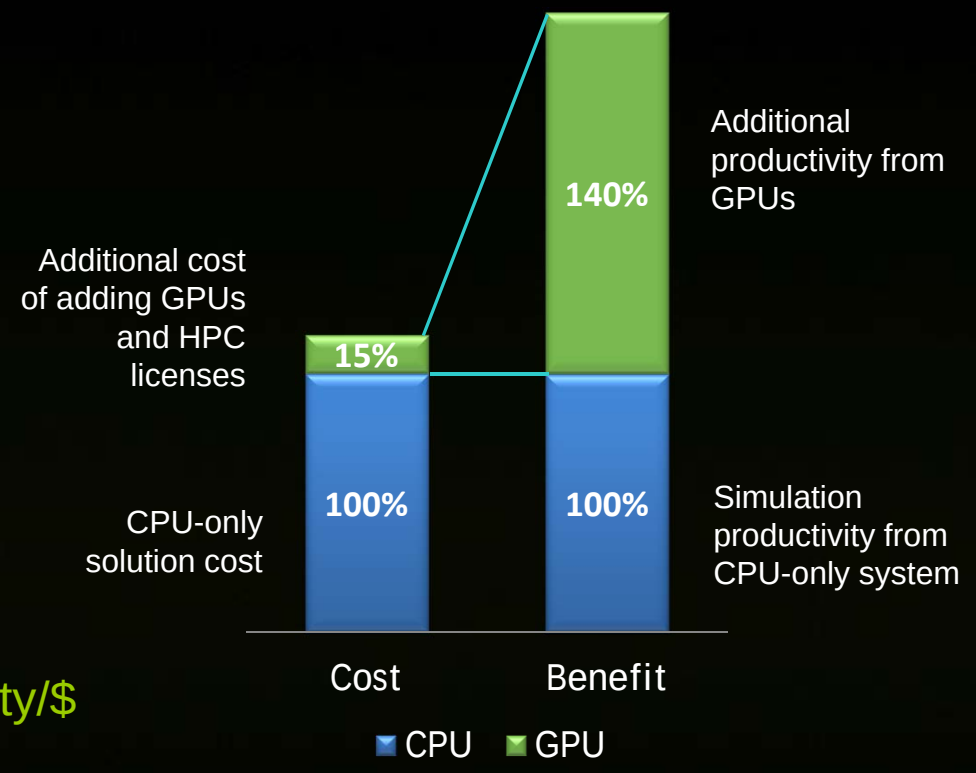

Large Model (~77 TFLOPs), 4.71M DOF, Nonlinear Static, Direct Sparse Solver, 2 DMP - Split
Abaqus 6.14-PR2 with Intel Xeon E5-2690v2, 3.0 GHz CPU, 128 GB memory; Tesla K20X

CPU-only solution cost is approximated and includes both hardware and paid-up software license costs.
Benefit/productivity is based on the number of completed Abaqus jobs/day.

Symmetric Solver Speed-up with DMP Split

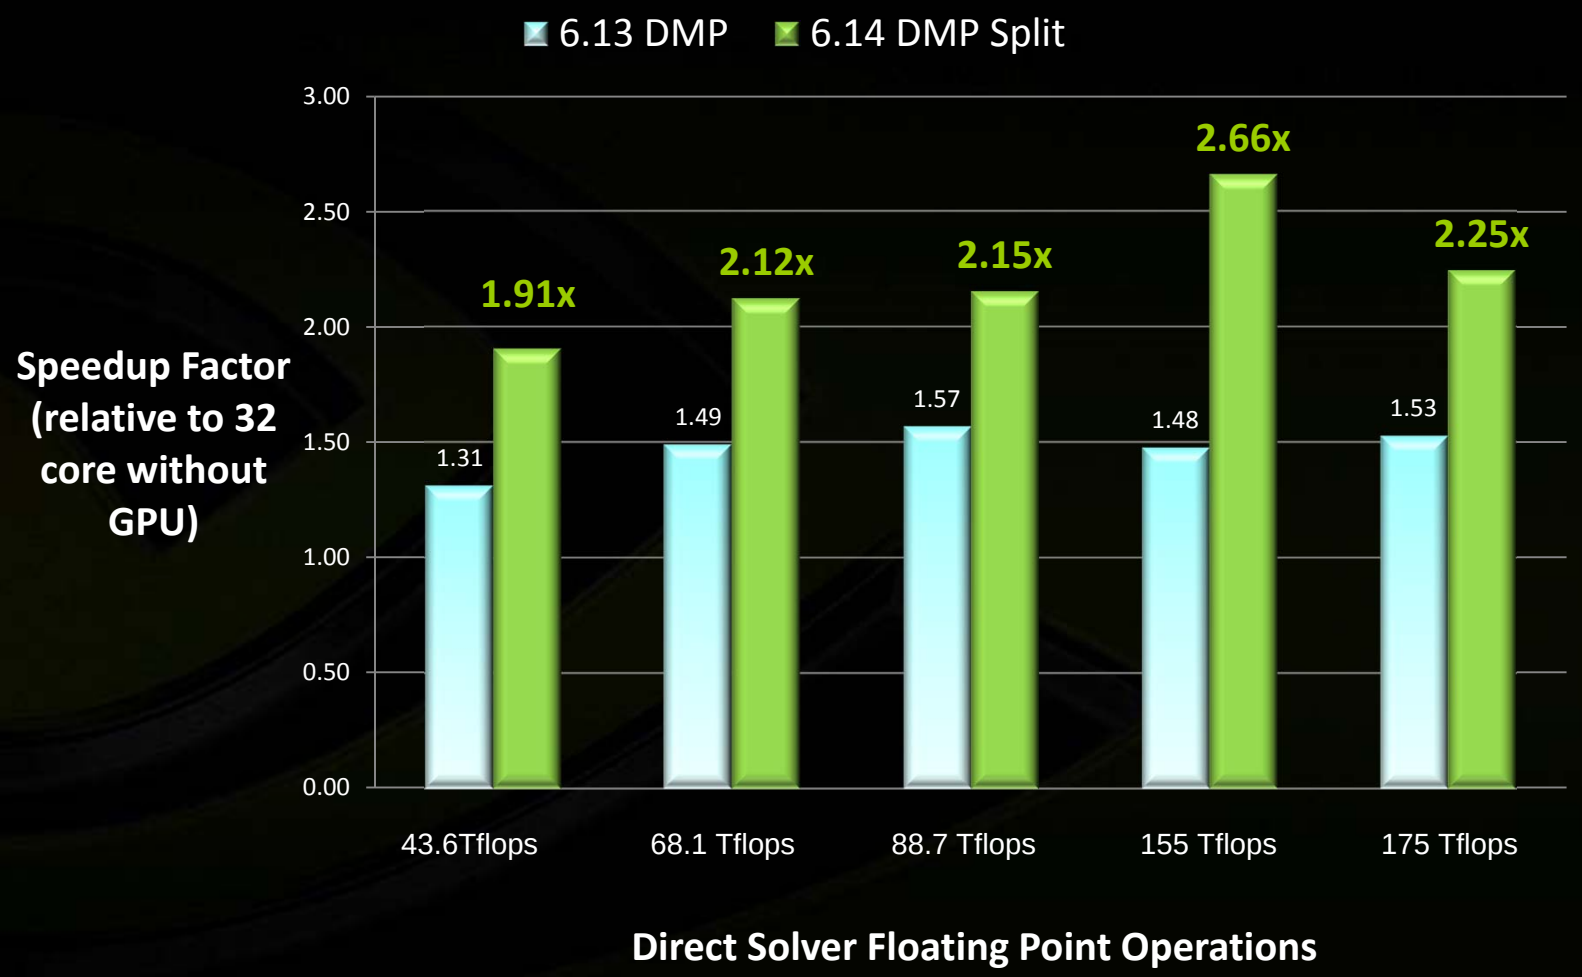

2 HP SL250, 2 Intel E5-2660 cpus (32 cores)

2 Nvidia K20m GPUs per compute node

128 Gb memory per compute node

Customer-confidential Auto model

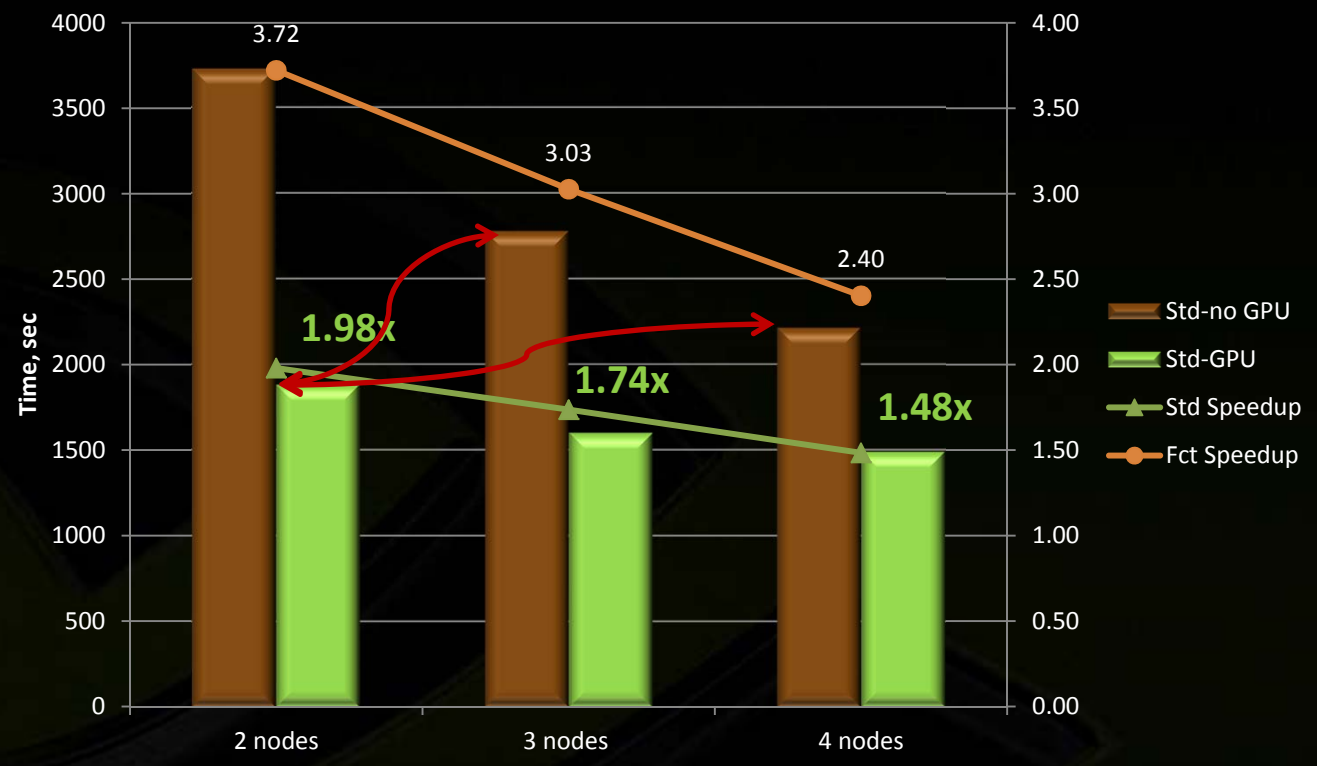

Up to 50% time savings!

In addition to time saving, there is license cost advantage of running on 2 nodes with GPUs compared to 3 and 4 nodes without GPUs.

2 MPI processes per compute node

Accelerated DMP execution mode (an optional feature in 6.14)

11M DOF on Sandybridge (16 cores+256 GB) and two K20 cards

MSc Nastran with NVIDIA GPUs

MSC Nastran GPU Computing

- **MSC Nastran direct equation solver is GPU accelerated**
  - Sparse direct factorization (MSCLDL, MSCLU)
  - Handles very large fronts
  - Impacts several solutions
    - High (SOL101 – Static Stress, SOL108 – Direct Freq Response)
    - Medium (SOL103 – Modal Analysis)
    - Low (SOL111 – Modal Freq Response , SOL400 – NL Static & Dynamic)

- **Support of NVIDIA multi-GPUs**
  - Tesla K20/K20X, Tesla K40, Quadro K6000 (compute & pre/post), Tesla 20-series

- **GPU Licensing**
  - Separate license feature supports unlimited GPU cores

MSC Nastran GPU Computing

Timeline

- **MSC Nastran 2012.1, 2H 2011**
  - Real & symmetric sparse direct solver is accelerated on the GPU
- **MSC Nastran 2012.x, 2012**
  - Complex & unsymmetric sparse direct solver is accelerated on the GPU
- **MSC Nastran 2013 & 2013.1, 2013**
  - Vastly reduced use of pinned host memory
  - Ability to handle arbitrarily large fronts, for very large models (> 15M DOF) on a single GPU with 6GB device memory

MSC Nastran 2013.1

SMP SOL101 and SOL103

30-70% time savings!

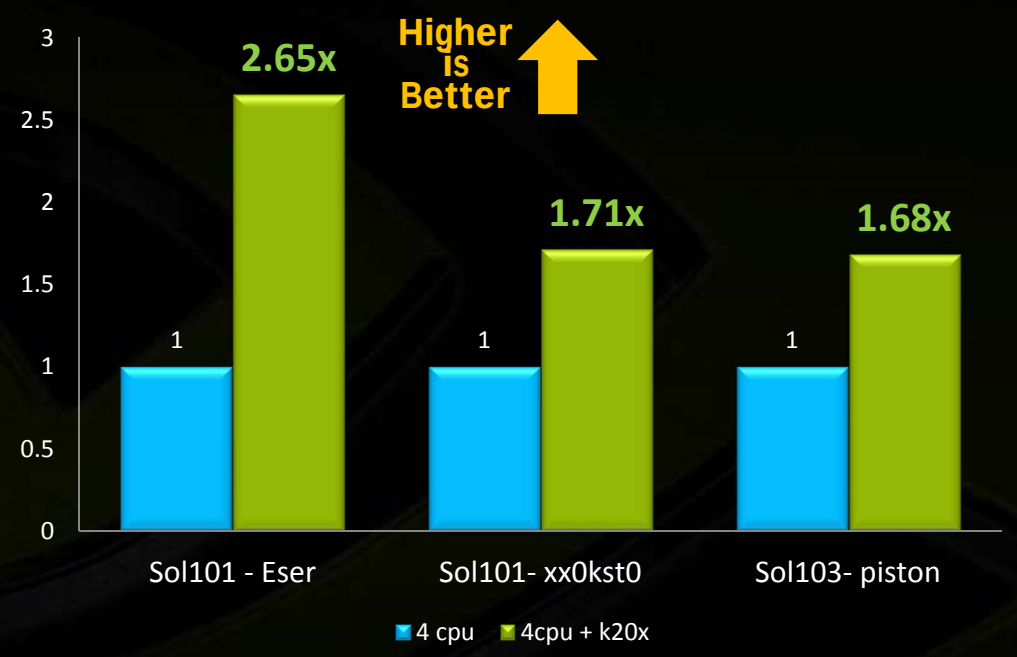

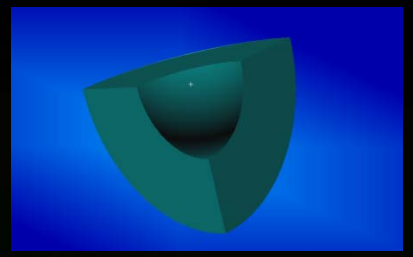

Hollow Sphere

Piston

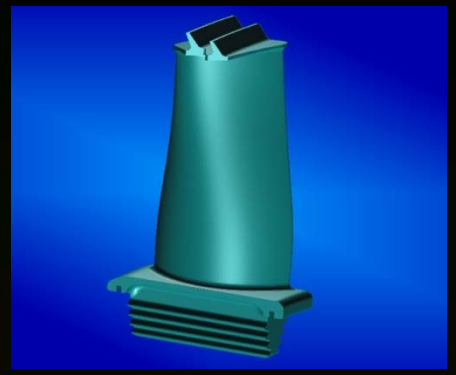

Turbine Blade

Server node: Ivy Bridge E5-2697v2 (2.7GHz), Tesla K20X GPU, 128 GB memory

AGENDA

- GPUs in Enterprise Computing
- Business Challenges in Product Development
- NVIDIA GPUs for CAE Applications
- Computational Fluid Dynamics (CFD) Applications
- Computational Structural Mechanics (CSM) Applications
- Computational Electro Magnetics (CEM) Applications
- NVIDIA Solutions
- Value of GPU-accelerated Simulations

GPU-accelerated CEM Applications

APPLICATION	DESCRIPTION	SUPPORTED FEATURES	MULTI-GPU SUPPORT
Agilent Technologies ADS	Simulation tool for design of RF, microwave and high speed digital circuits	Signal integrity simulation	Single only
Agilent Technologies EMPro	Modeling and simulation environment for analyzing 3D EM effects of high speed and RF/Microwave components	FDTD solver	Yes
ANSYS Nexxim	Circuit simulation engine for RF/analog/ mixed-signal IC design; IBIS-AMI analysis speedup with GPU computing.	AMI analysis	Single only
ANSYS HFSS	Simulation tool for modeling 3-D full-wave electromagnetic fields in high-frequency and high-speed electronic components.	Transient solver	Yes
CST Microwave Studio (MWS)	3D EM modeling and simulation	Transient solver	Yes
Delcross Savant	Simulation tool for installed antenna performance and antenna-to-antenna coupling	High-frequency solver	Yes
EMSS FEKO	3D EM modeling and simulation	MoM solver	Single only
Remcom XFdtd	3D EM modeling and simulation	FDTD solver	Yes
SPEAG SEMCAD-X	3D EM modeling and simulation	FDTD solver	Yes
Rocketick RocketSim	Verilog simulation	Verilog simulation	Yes

GPU-acceleration of ANSYS HFSS

- Average speed up of 2.41x and a maximum of 5.21x on a Tesla K20

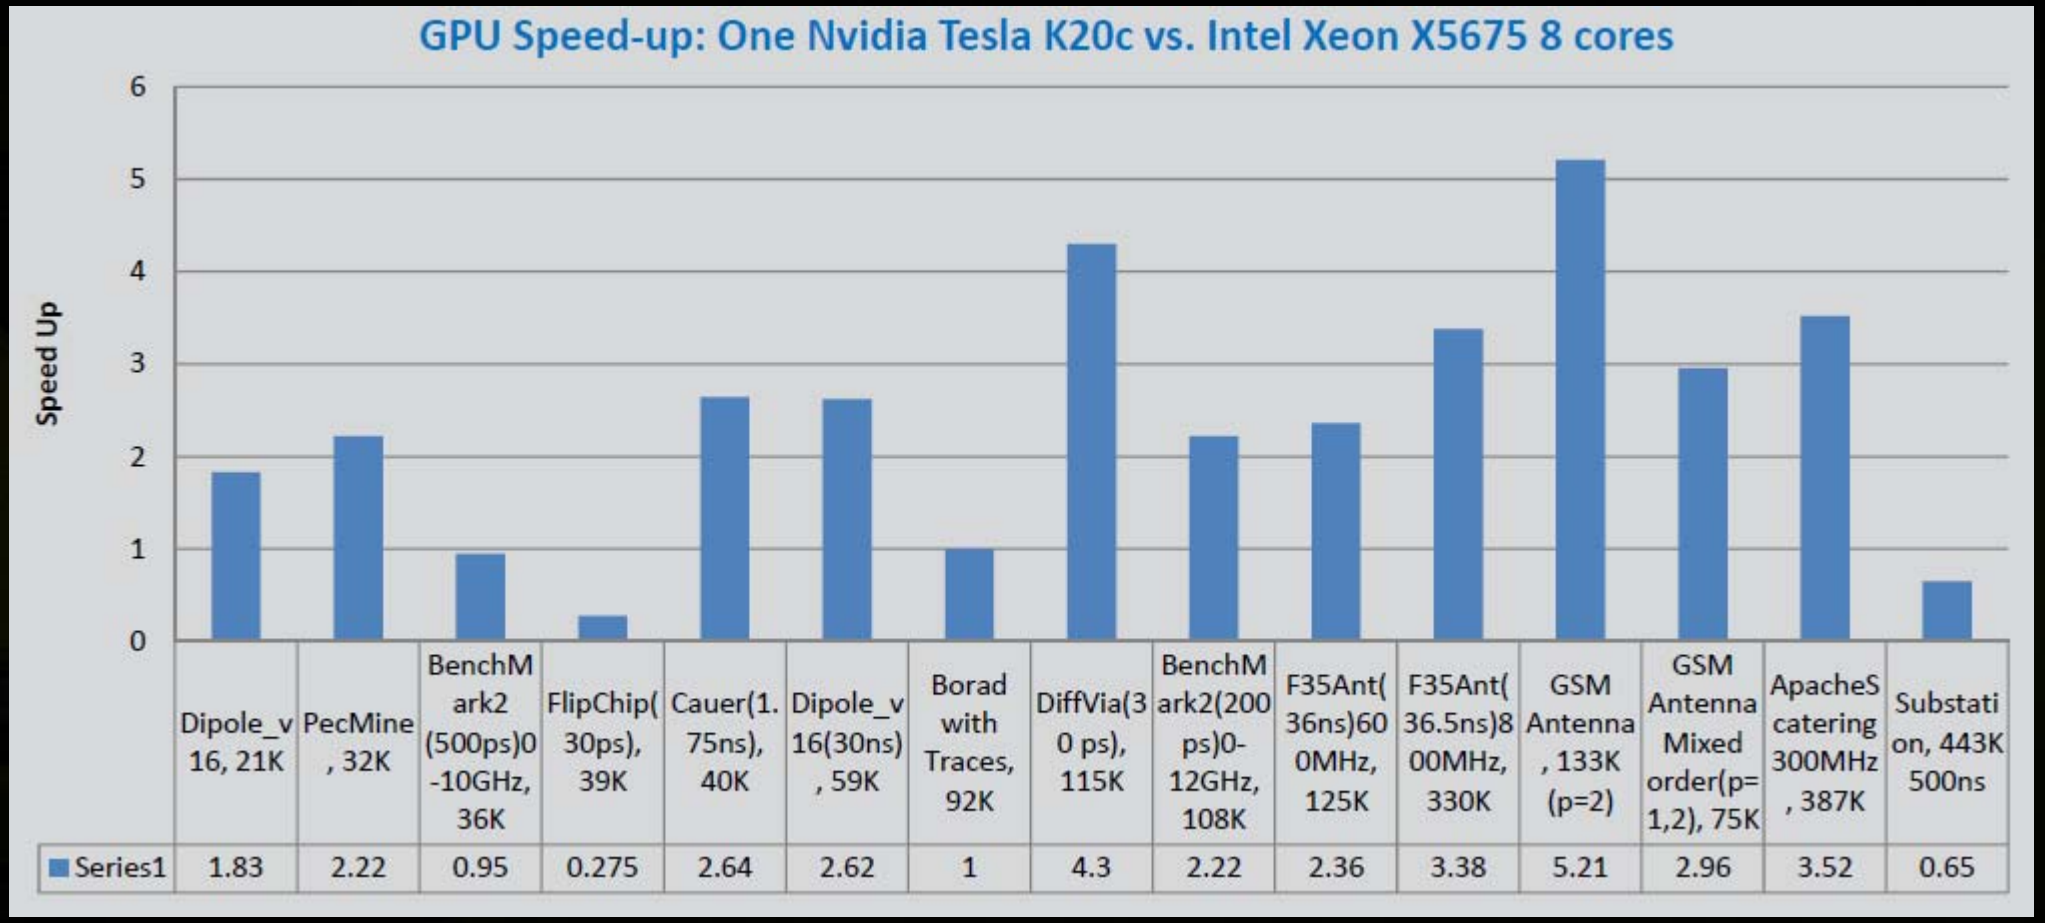

AGENDA

- GPUs in Enterprise Computing
- Business Challenges in Product Development
- NVIDIA GPUs for CAE Applications
- Computational Fluid Dynamics (CFD) Applications
- Computational Structural Mechanics (CSM) Applications
- Computational Electro Magnetics (CEM) Applications
- NVIDIA Solutions
- Value of GPU-accelerated Simulations

NVIDIA Kepler family GPUs for CAE simulations

K20 (5 GB)
K20X (6 GB)
K40 (12 GB)

K6000 (12 GB)

MAXIMUS Solution for Workstations

NVIDIA® MAXIMUS

Visual Computing	Parallel Computing
CAD Operations	FEA
Pre-processing	CFD
Post-processing	CEM

NVIDIA MAXIMUS™ TECHNOLOGY

- Intelligent GPU job Allocation
- Unified Driver for Quadro + Tesla
- ISV Application Certifications
- HP, Dell, Lenovo, others
- Now Kepler-based GPUs

Available Since November 2011

AGENDA

- GPUs in Enterprise Computing
- Business Challenges in Product Development
- NVIDIA GPUs for CAE Applications
- Computational Fluid Dynamics (CFD) Applications
- Computational Structural Mechanics (CSM) Applications
- Computational Electro Magnetics (CEM) Applications
- NVIDIA Solutions
- Value of GPU-accelerated Simulations

Benefits of GPU-accelerated simulations

More simulations in the same amount of time or same number of simulations in less amount of time

More design points can be analyzed for better-quality products without slipping project schedules

Simulation times can be cut into half, thus shorter product development times

Mesh sizes can be doubled or advanced models can be used without increasing simulation times

Thank you

Bhushan Desam
bdesam@nvidia.com