ANSYS

Realize Your Product Promise™

ANSYS Improvements to Engineering Productivity with HPC and GPU-Accelerated Simulation

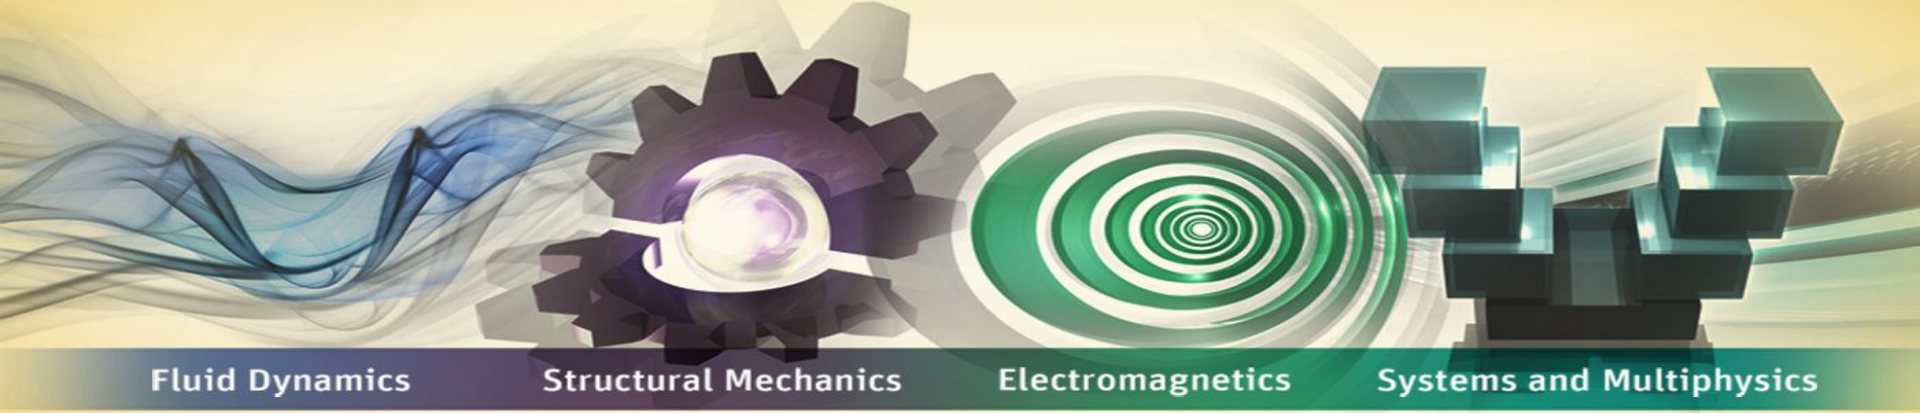

Ray Browell

nVidia Technology Theater

SC12 – November 13, 2012

HPC Revolution

- **Recent advancements have revolutionized the computational speed available on the desktop**
  - Multi-core processors
    - Every core is really an independent processor
  - Large amounts of RAM and SSD
  - GPUs

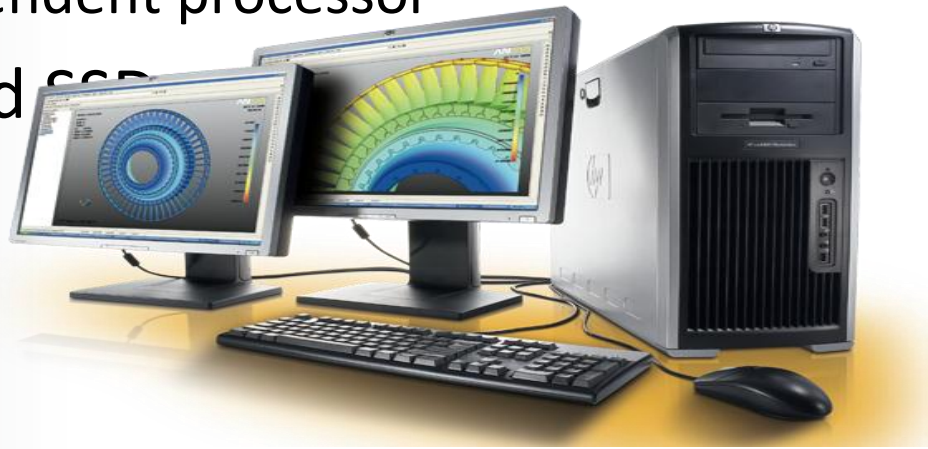

Mechanical GPU Accelerator Capability

- "Accelerate" Sparse direct equation solver (SMP & DMP)
  - GPU is used to factor many dense "frontal" matrices
  - Decision is made automatically on when to send data to GPU
    - "Frontal matrix" too small, too much overhead, stays on CPU
    - "Frontal matrix" too large, exceeds GPU memory, only partially accelerated
- "Accelerate" PCG/JCG iterative solvers (SMP & DMP)
  - GPU is only used for sparse-matrix vector multiply (SpMV kernel)
  - Decision is made automatically on when to send data to GPU
    - Model too small, too much overhead, stays on CPU
    - Model too large, exceeds GPU memory, only partially accelerated

Mechanical GPU Accelerator Capability

- **Supported hardware**
  - Currently support NVIDIA Tesla 20-series, Quadro 6000, and Quadro K5000 cards
  - Next generation NVIDIA Tesla cards (Kepler) should work with ANSYS R14.5
  - Installing a GPU requires the following:
    - Larger power supply (single card needs ~250W)
    - Open 2x form factor PCIe x16 2.0 (or 3.0) slot
- **Supported platforms**
  - Windows and Linux 64-bit platforms only
    - Does not include Linux Itanium (IA-64) platform

Mechanical GPU Accelerator Capability

- **Targeted hardware**

	NVIDIA Tesla C2075	NVIDIA Tesla M2090	NVIDIA Quadro 6000	NVIDIA Quadro K5000†	NVIDIA Tesla K10	NVIDIA Tesla K20†
Power (W)	225	250	225	122	250	250
Memory	6 GB	6 GB	6 GB	4 GB	8 GB	6 to 24 GB
Memory Bandwidth (GB/s)	144	177.4	144	173	320	288
Peak Speed SP/DP (GFlops)	1030/515	1331/665	1030/515	2290/95	4577/190	5184/1728

† These NVIDIA “Kepler” based products are not released yet, so specifications may be incorrect

Mechanical GPU Accelerator Capability

- Supports majority of ANSYS users
  - Covers both sparse direct and PCG iterative solvers
  - Only a few minor limitations
- Ease of use
  - Requires at least one supported GPU card to be installed
  - Requires at least one HPC pack license
  - No rebuild, no additional installation steps
- Performance
  - ~10-25% reduction in time to solution when using 8 CPU cores
  - *Should never slow down your simulation!*

ANSYS Mechanical 14.5 Preview

Results for Distributed ANSYS 14.5 Preview and Xeon 8-Core CPUs

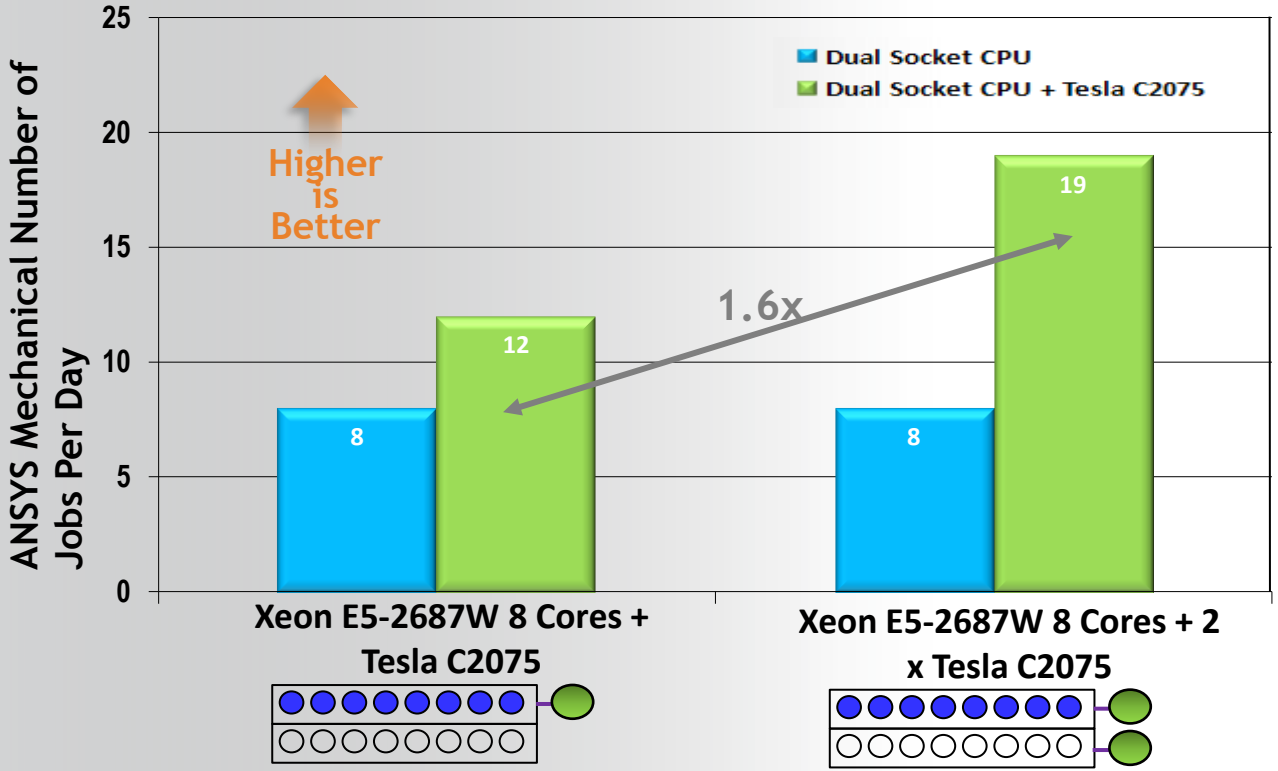

V14sp-5 Model

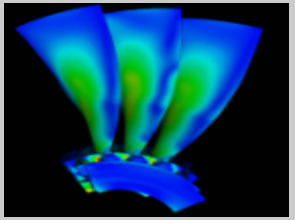

- Turbine geometry
- 2,100 K DOF
- SOLID187 FEs
- Static, nonlinear
- One iteration
- ANSYS Mechanical14.5
- Direct sparse solver

Results from HP Z820; 2 x Xeons (16 Cores, use of only 8) 128GB memory, Win7; 2 x Tesla C2075

Structural GPU Accelerator Capability

- GPUs can offer significantly faster time to solution

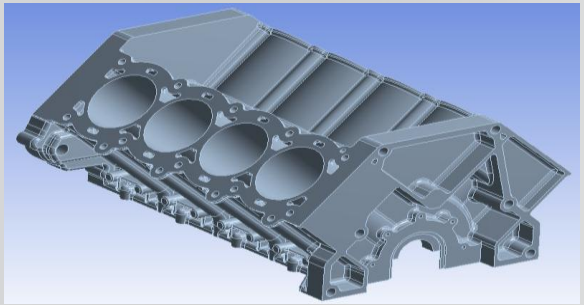

- 6.5 million DOF
- Linear static analysis
- Sparse solver (DMP)
- 2 Intel Xeon E5-2670 (2.6 GHz, 16 cores total), 128 GB RAM, SSD, 4 Tesla C2075, Win7

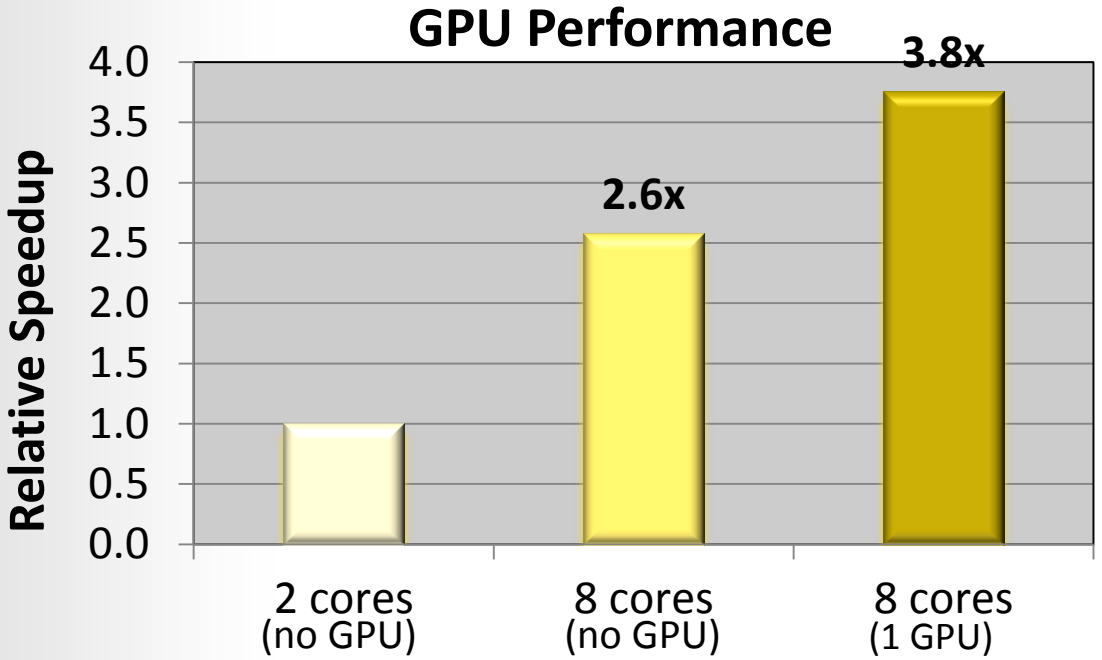

Structural GPU Accelerator Capability

- GPUs can offer significantly faster time to solution

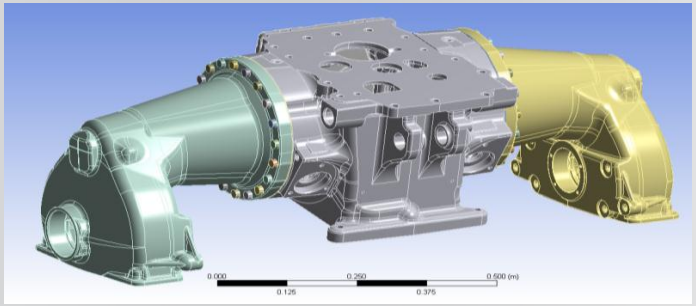

- 11.8 million DOF
- Linear static analysis
- PCG solver (DMP)
- 2 Intel Xeon E5-2670 (2.6 GHz, 16 cores total), 128 GB RAM, SSD, 4 Tesla C2075, Win7

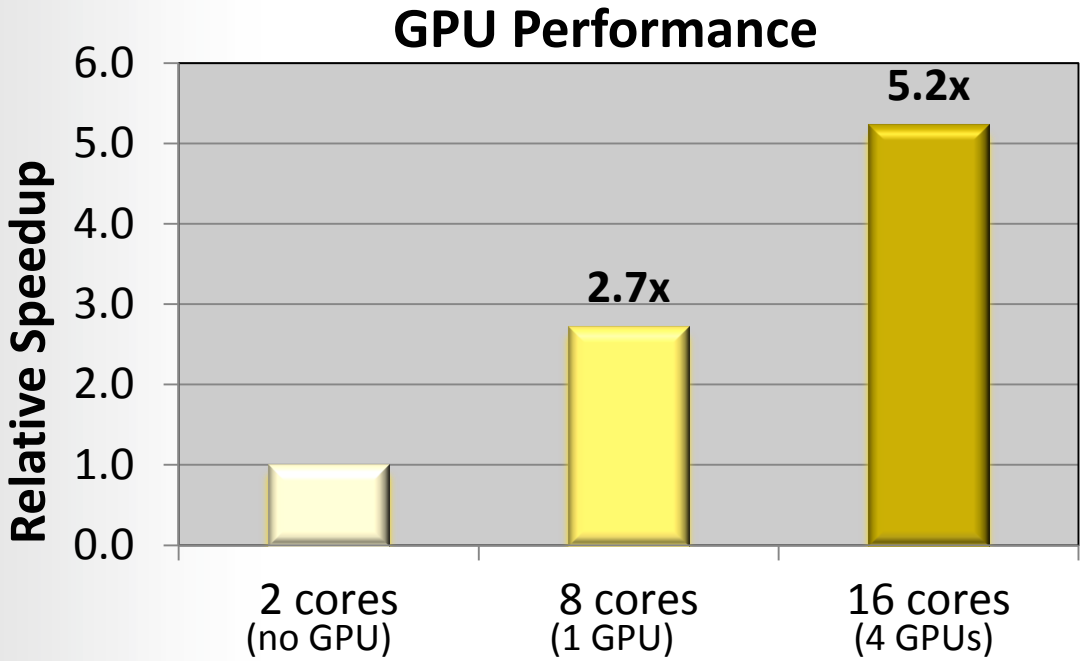

ANSYS and NVIDIA Collaboration

Release	ANSYS Mechanical 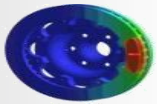	ANSYS Fluent 
13.0 Dec 2010	Shared Memory Solvers; Single Node/ Single GPU	
14.0 Dec 2011	+ Distributed ANSYS; Multi-node / 1 GPU/node	Radiation Heat Transfer (beta)
14.5 Nov 2012	+ Multi-GPU / node; + Hybrid PCG;	+ GPU AMG Solver (beta), Single GPU

Fluent Radiation Modeling on GPUs

VIEWFAC

- Utility to compute view factors
- Hybrid MPI-OpenMP-OpenCL parallel implementation
- Works on CPUs, GPUs or both

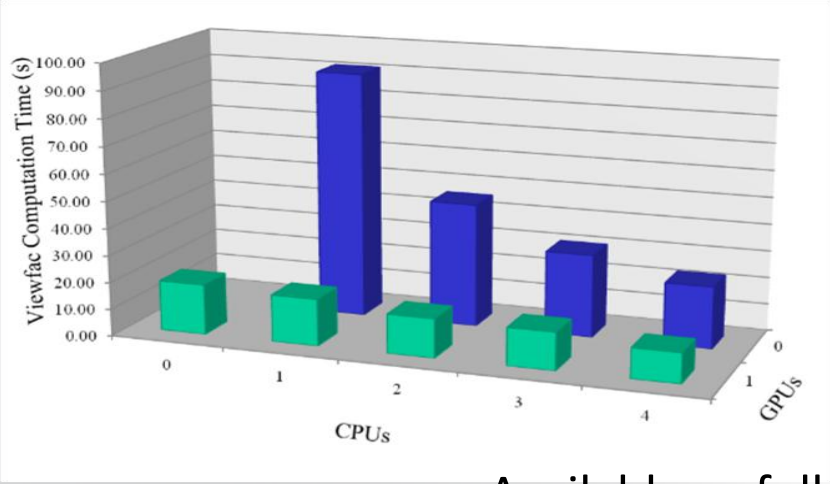

RAY TRACING

- Utility to compute view factors
- Uses Optix on NVIDIA C2070

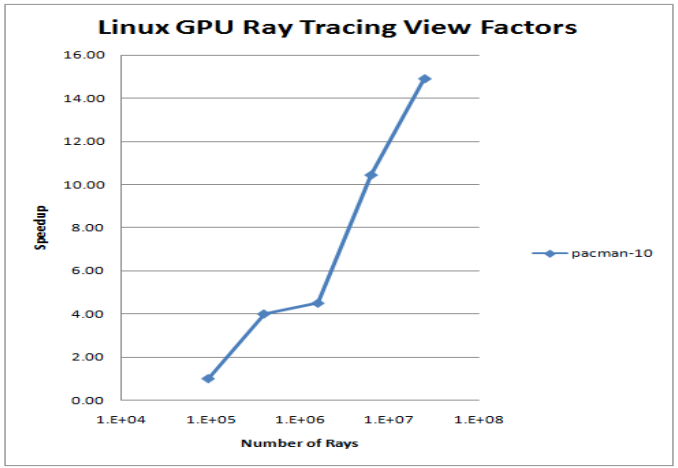

Available as full features in 14.5

Fluent AMG Solver on GPUs

Work-in-Progress

NVAMG Project - Preview of ANSYS Fluent 14.5 Performance

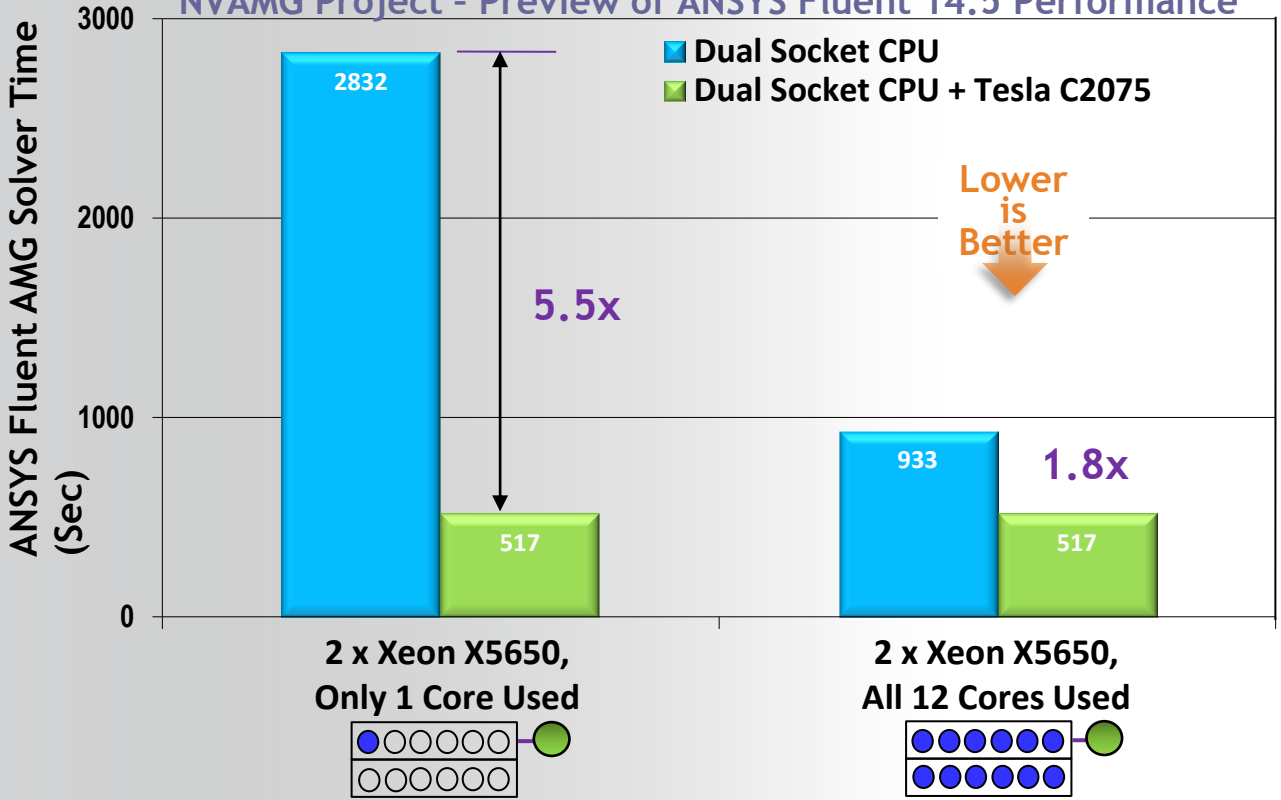

Helix Model

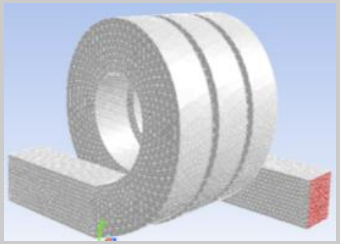

- Helix geometry
- 1.2M Hex cells
- Unsteady, laminar
- Coupled PBNS, DP
- AMG F-cycle on CPU
- AMG V-cycle on GPU

NOTE: All jobs solver time only, ~65% of total time

Design Optimization

- **How will you use all of this computing power?**

Higher fidelity

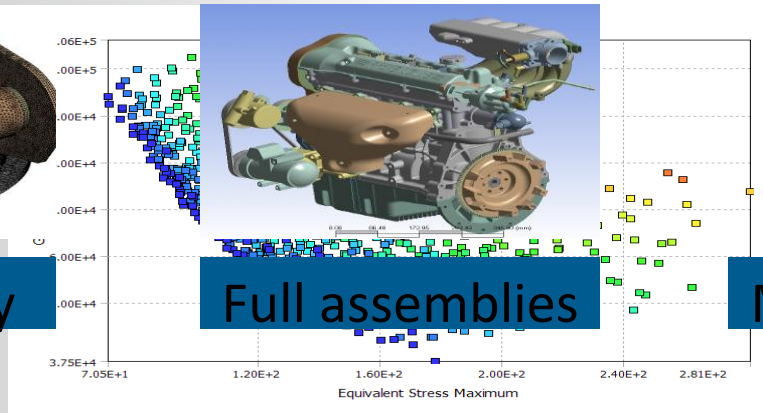

Full assemblies

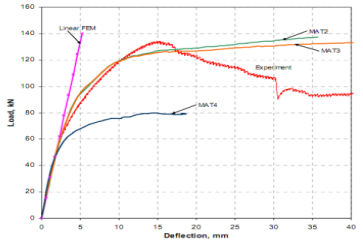

More nonlinear

Design Optimization Studies

HPC Revolution

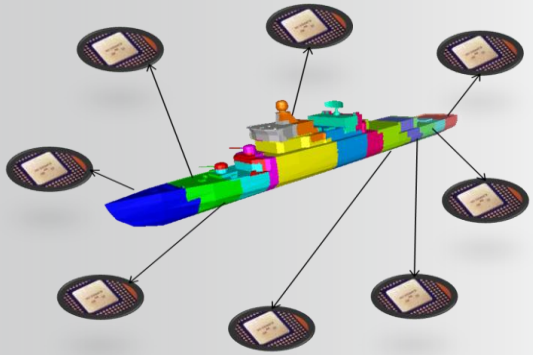

SMP vs. DMP

HDD vs. SSDs

GPUs

The right combination of algorithms and hardware leads to maximum efficiency

Clusters

Interconnects

Improving Engineering Productivity with HPC and GPU-Accelerated Simulation

Thank You!

Questions?

Raymond Browell

724.514.3070

Ray.Browell@ansys.com

