Supercharging Engineering Simulations at Mercury Marine with NVIDIA GPUs

S4669: NVIDIA GPU Technology Conference

Arden Anderson

March 26, 2014

Overview

- High Performance Computing has been an enabling technology for Simulation Based Product Development at Mercury Marine
  - Cost effective
  - Enhances product insight
- A journey from desktop computing to HPC with hundreds of cores over the last 10 years
- GPU's are the most recent step to increased capability

Audience Questions (show of hands)

- How many people run numerical simulations like FEA, CFD, etc?
- How many compute cores do you use for your largest jobs?
  - Less than 8
  - 8-64
  - 64-256
  - More than 256
- How many people utilize GPU's for numerical simulation?

Outline

- About Mercury Marine
- Engineering simulation capabilities
- Progression of High Performance Computing (HPC)
- GPU Focus
- Summary

Mercury Marine Founded in Cedarburg, WI in 1939

- Mercury Marine began as the Kiekhaefer Corp. in 1939
  - Founded by E. Carl Kiekhaefer
- Mercury acquired by Brunswick Corporation in 1961
  - Leader in active recreation: marine engines, boating, bowling, billiards, and fitness

- Employs 5,400 People Worldwide
- Fond du Lac, WI campus includes
  - Corporate Offices
  - Technology Center, R&D Offices
  - Outboard and Sterndrive Manufacturing (Casting, Machining, Assembly to Distribution)

Mercury's 1st Patent

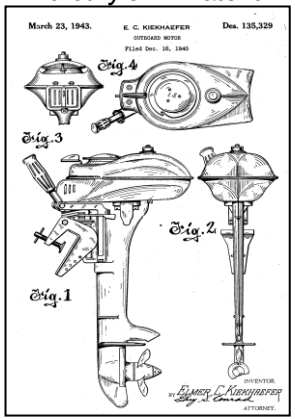

The Most Comprehensive Product Offering In Recreational Marine

Outboard Engines (2.5 hp to 350 hp)

All new or updated in last 5 years

Sterndrive Engines (135 hp to 1650 hp)

AXIUS

ZEUS

All updated to new emissions standard in last year

Props / Rigging / P&A

MERCURY Precision Parts

QUICKSILVER Marine Parts & Accessories

Land 'N' Sea / Attwood

attwood genuine ingenuity

MotorGuide NEVER STOP. motorguide.com

SEACHOICE PRODUCTS

Diversified, Quality Products, Connected to Parent Corporation

Outline

- About Mercury Marine
- Engineering simulation capabilities
- Progression of High Performance Computing (HPC)
- GPU Focus
- Summary

Standard FEA

Fatigue & Hardware Correlation

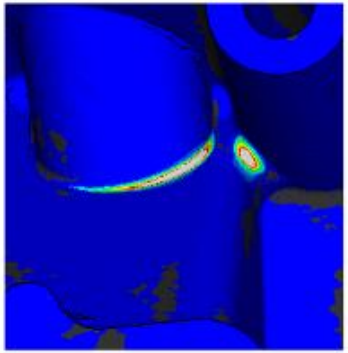

Non-Linear Gaskets

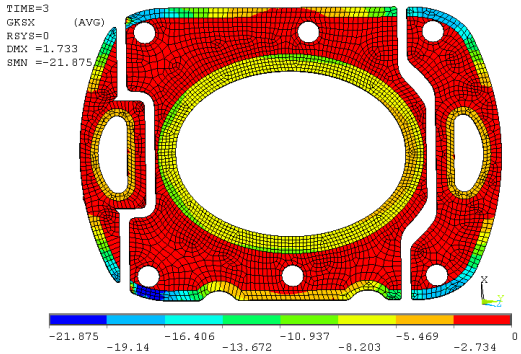

System Assemblies with Contact

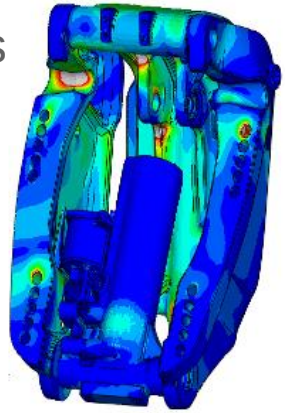

Sub-Modeling

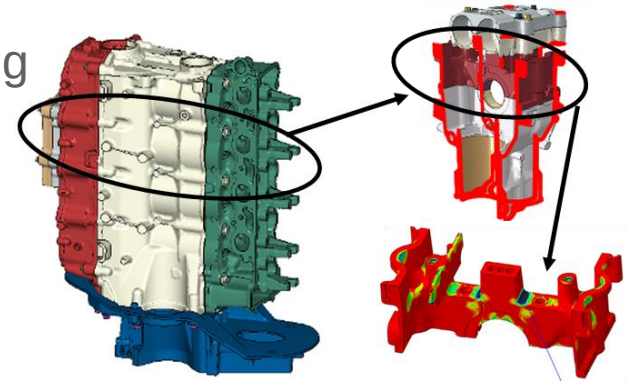

MERCURY

Explicit FEA

System level submerged object impact

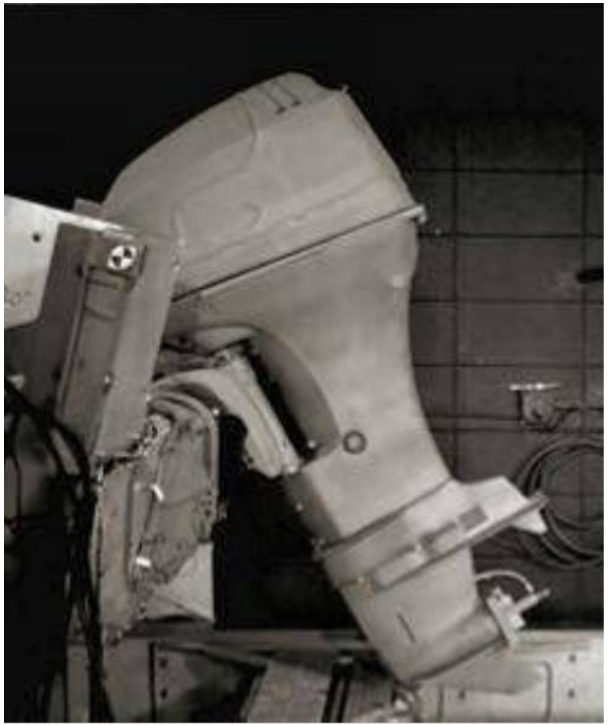

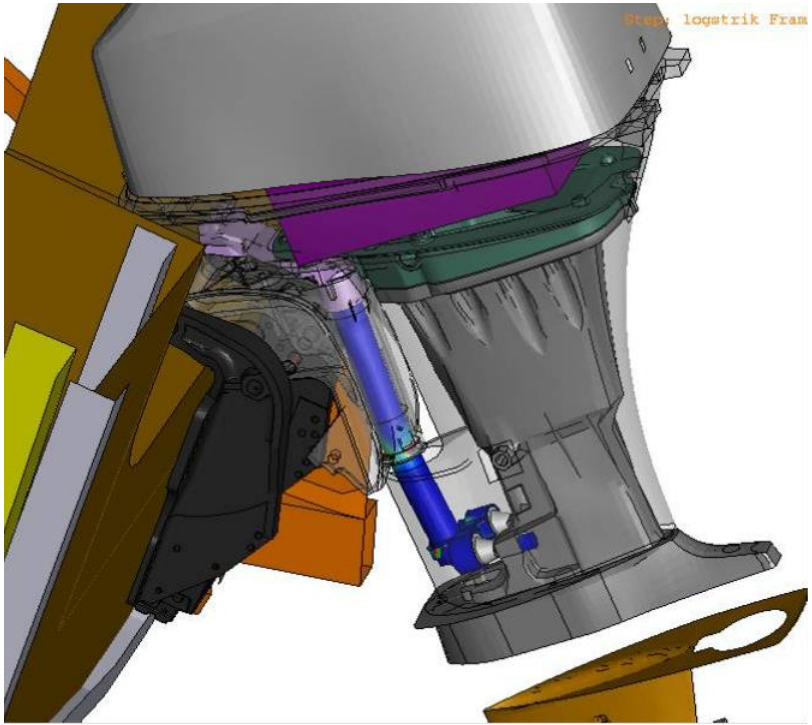

CFD

Transient Internal Flow

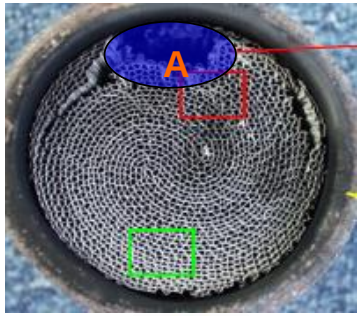

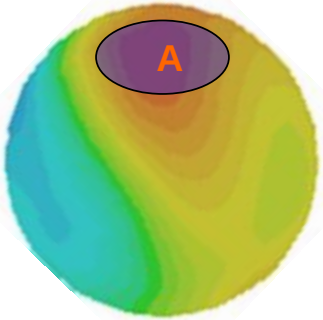

Flow distribution correlated to hardware

External Flow

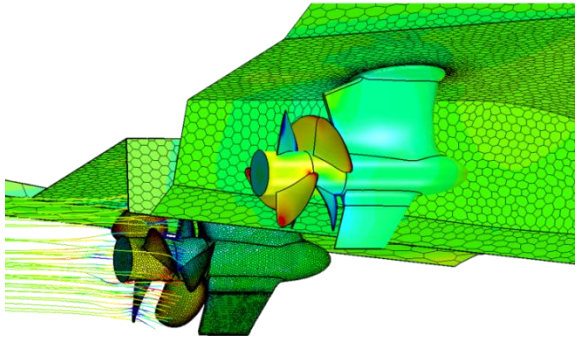

Gear Case and Hull Flow

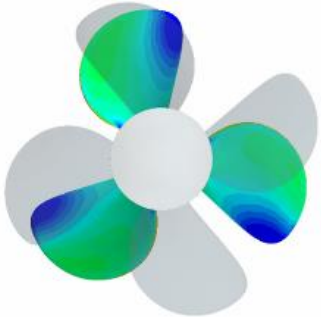

Moving Mesh Propeller

Two Phase Flow

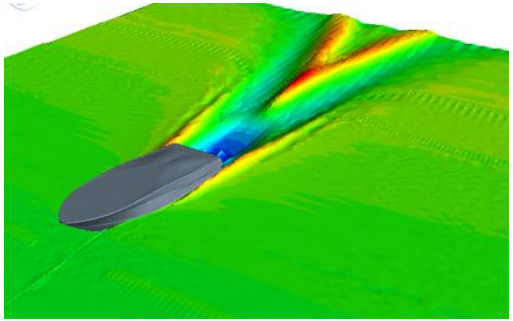

Vessel drag, heave, and pitch

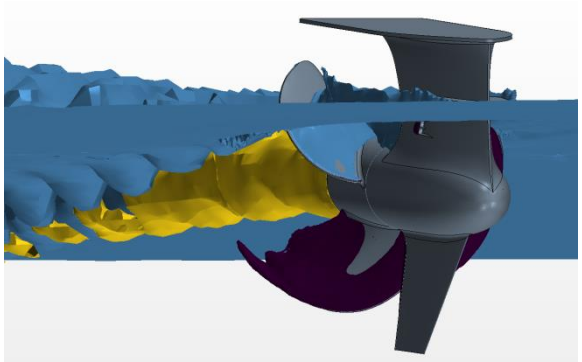

Surfacing Propeller

MERCURY

Heat Transfer

Enclosure Air Flow &
Component Temperatures

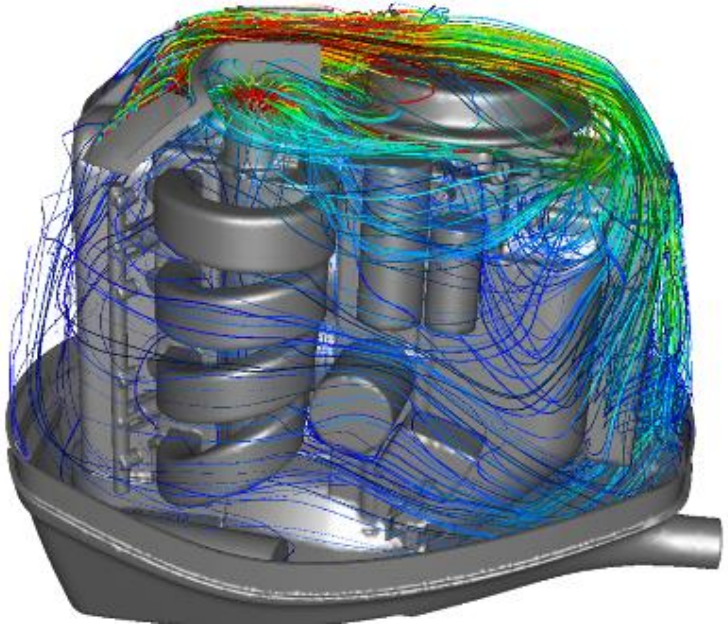

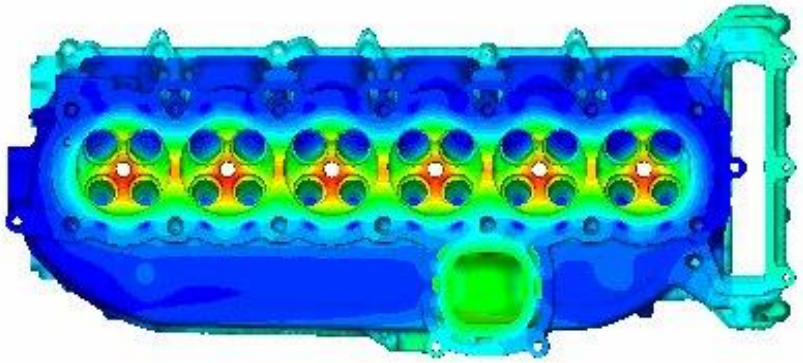

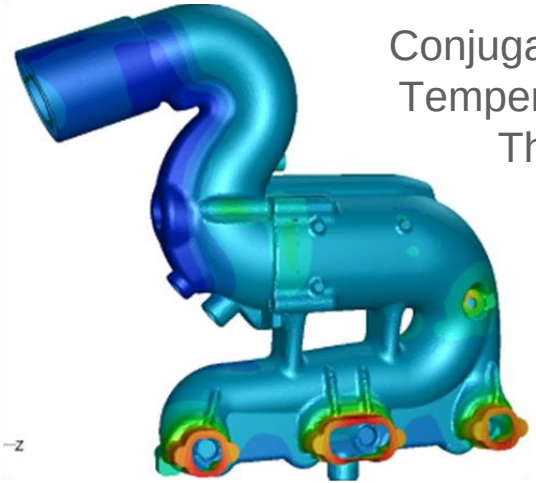

Conjugate Heat Transfer for
Temperature Distribution &
Thermal Fatigue

Outline

- About Mercury Marine
- Engineering simulation capabilities
- Progression of High Performance Computing (HPC)
- GPU Focus
- Summary

Motivation

- Request from management to reduce simulation turn around time
  - Some simulations were taking 2+ weeks as runs became more detailed and complex
- Quicker feedback to avoid late tooling changes
- Need to minimize manpower down time
- Large software costs
  - **Need to maximize software investment**

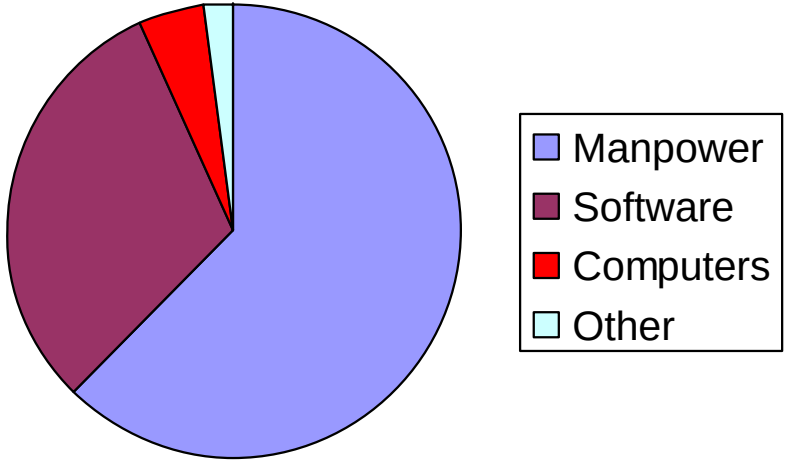

Evolution of Computing at Mercury Marine

2004

Desktop Environment

- Pre and post processing on Windows PC, 2 GB RAM
- Computing on HP Unix Workstation
- ~$200k for 10 boxes
- Memory limitations on pre-post and limited model size
- Minimal parallel-processing
- A to B comparative

2005

Compute Server

- Four node server with 10 compute cores
- $125k server
- Increased model sizes with larger memory
- Parallel processing for FEA & CFD
- Approximately the same number of processors as previous system with large increases in speed and capability

2007

Carpet Cluster

- Updated desktop machines to Linux workstations
- Increased computing by clustering the new Linux machines
- Continued use of the compute server
- **New methods are developed and validated, but more speed is still needed to meet the pace of product development**

Move from desktop computing to "carpet cluster"…

MERCURY

Evolution of Computing at Mercury Marine

2009

HPC Server

- Four 8-core nodes
- Job scheduler increases efficiency
- Hardware costs offset by rebalancing software costs, parallel processing has much more efficient license structure
- Small incremental cost
- Most simulations are meeting the pace of product development; increased appetite for analysis

2010

Augment HPC Server

- Add six 8-core nodes (ten nodes, 80 core total)
- Infiniband switch enhances multi-node runs
- Additional method development, particularly for CFD
- Up to 8x improvement in runtime, bringing even more simulations up to the pace of product development

2012

Upgrade: HPC + GPU

- Refresh the HPC system with ten new 12-core nodes
- **Four nodes include 2 NVIDIA Tesla GPU's for additional performance**
- Approximately cost neutral for 50% more cores (80 to 120)
- **Analysis work is tightly integrated into product development with many success stories**

HPC deployed in 2009 and progresses through 2014…

Outline

- About Mercury Marine
- Engineering simulation capabilities
- Progression of High Performance Computing (HPC)
- GPU Focus
- Summary

GPU Acceleration Learning Curve

- Abaqus first released GPU acceleration in their commercial code in 2011
  - Additional capability in each new release (multi-gpu per host, multi-host, more routines, etc)
- Special thanks to NVIDIA and Abaqus for helping Mercury Marine evaluate the new capability
- Initial evaluation used a Lenovo D20 Maximus workstation
  - Dual hex-core X5670 CPUs, 48 GB RAM
  - Tesla C2075 GPU (compute only)
  - Quadro 6000 (graphics and compute tasks)
- Positive results from the workstation evaluation
- For Mercury Marine, GPU's are more beneficial as a shared resource in the cluster
- GPU's were deployed for 2012 upgrade to our computing cluster
  - Details on the following slides…

Mercury Marine Benchmark for Abaqus/Standard

- Similar to Abaqus S4b benchmark, but larger
- Block + Head + Bedplate
- Preload + Thermal + Reciprocating Forces
- 8,800,000 DOF
- 55Gb Memory

** Picture of Abaqus benchmark S4b*

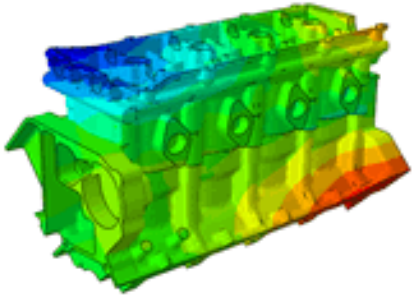

- Before hardware acquisition Mercury Marine ran several trials
- Also evaluated Abaqus benchmark data published on their website

GPU Hardware Costs

For the hardware that Mercury Marine considered in 2012:

- A compute host with dual hex-core CPU (12 core) cost approximately the same amount as two Tesla M2070 GPU's

Either/or:

- (8) compute hosts with 12 core each
- (4) compute hosts with 12 core PLUS 2 Tesla M2070 GPU's each

- The decision was not clear if considering only hardware cost and speed
  - A representative simulation showed that two hosts of CPU-only was approximately the same speed as a single host of CPU + GPU

GPU Accelerated Performance

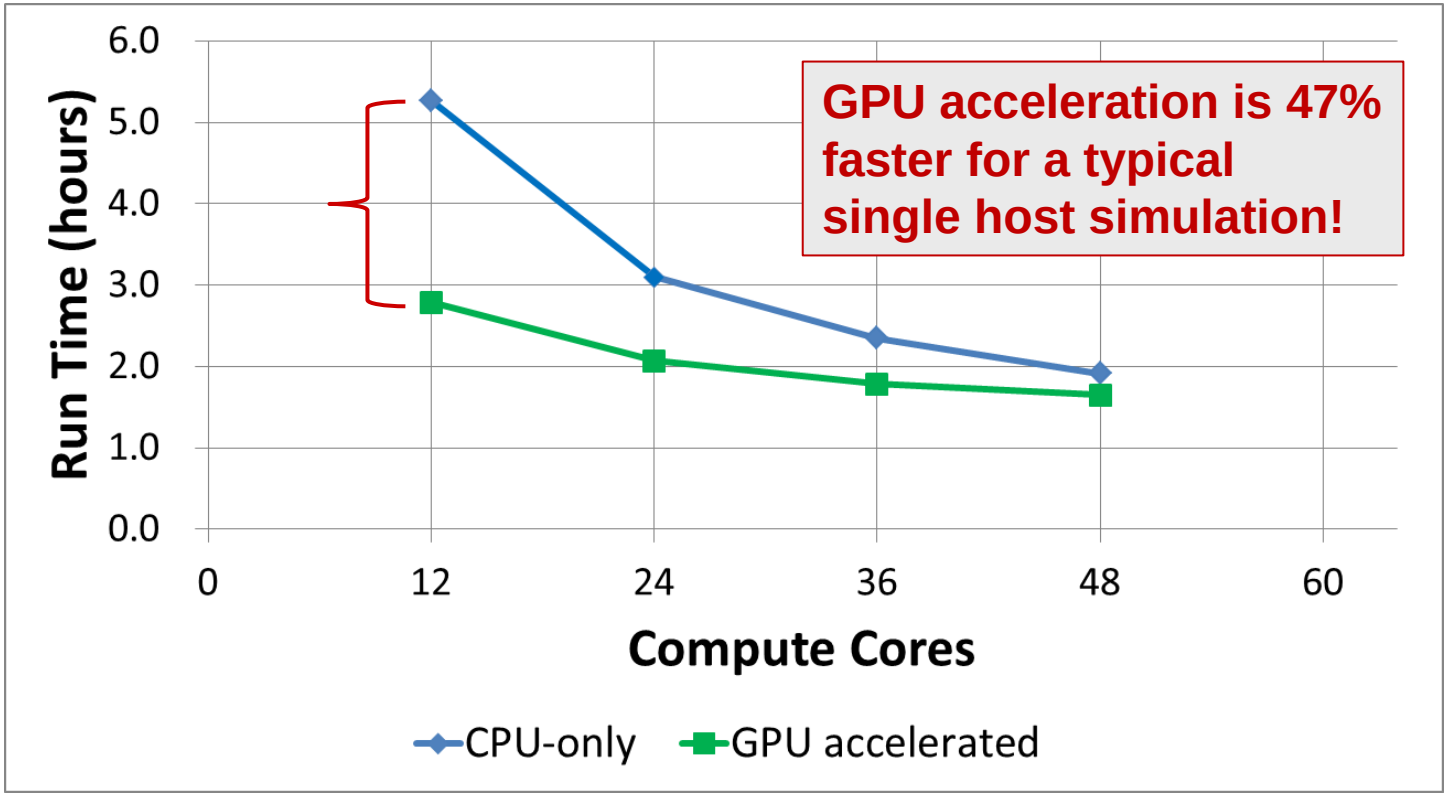

* 'GPU accelerated' data includes 2 GPUs per 12 core host

MERCURY

Analysis Software Cost

- Abaqus FEA token system is based on number of cores used
  - The curve flattens as number of cores increase
- Abaqus token structure considers 1 gpu the same as 1 compute core
- Need to consider hardware AND software costs

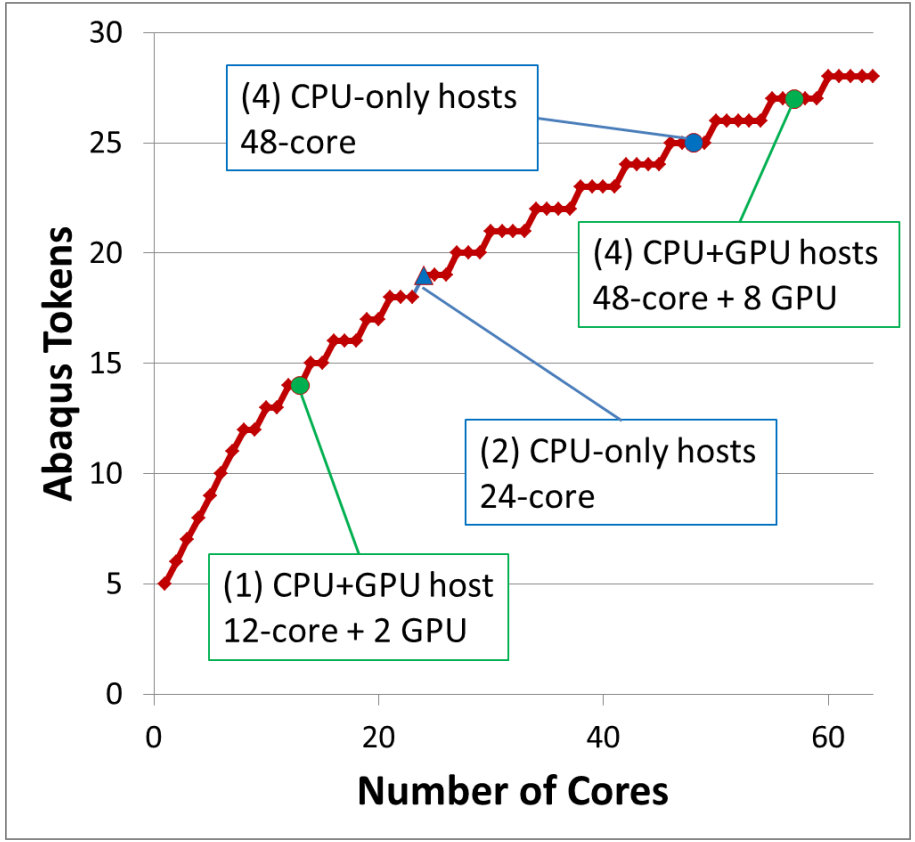

Software Advantage of GPU Accelerated FEA

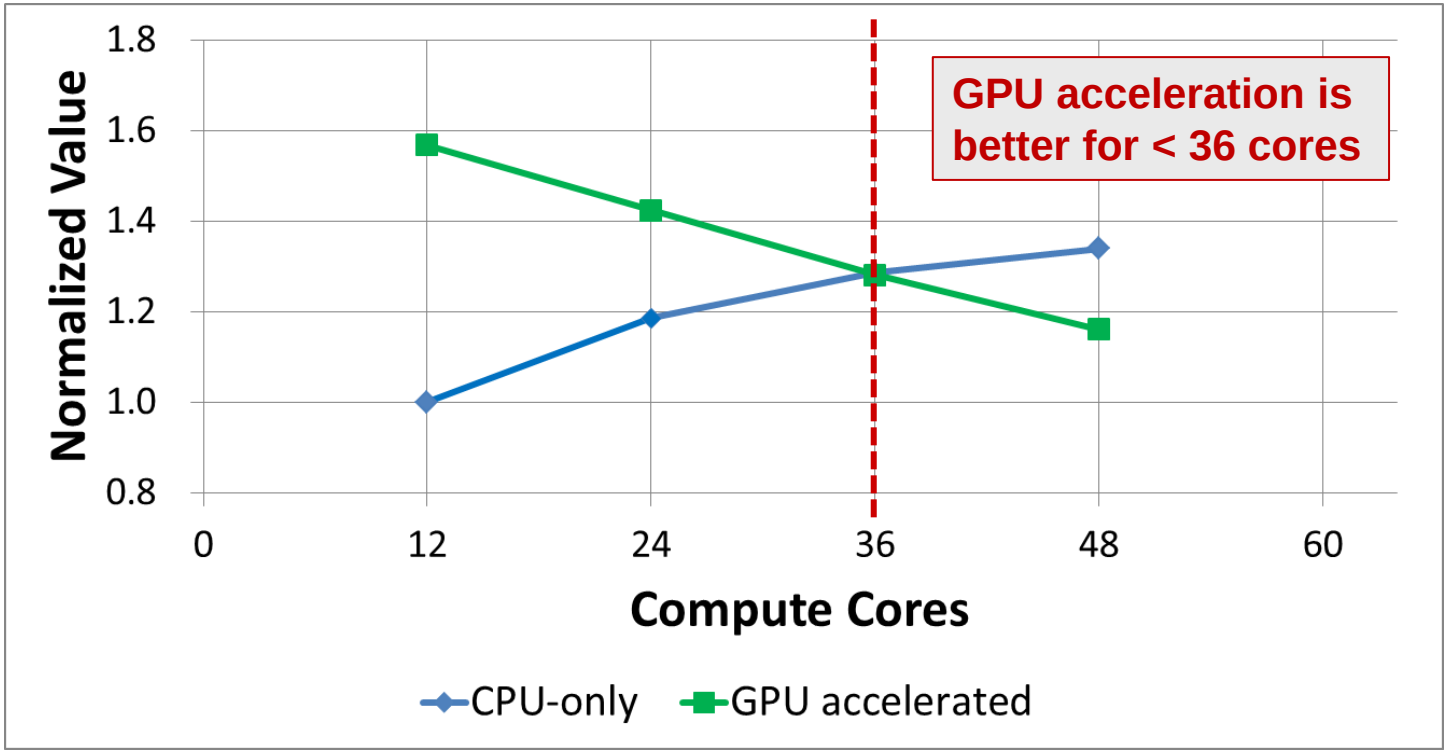

* 'GPU accelerated' data includes 2 GPUs per 12 core host

** From previous slide: Normalized Value = (Runtime x Normalized Cost). High is good.

MERCURY

Mercury Marine Benchmark for Abaqus/Standard

- Similar to Abaqus S4b benchmark, but larger
- Block + Head + Bedplate
- Preload + Thermal + Reciprocating Forces
- 8,800,000 DOF
- 55Gb Memory

** Picture of Abaqus benchmark S4b*

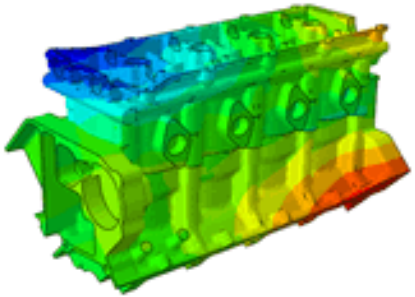

		CPU-only				GPU-accelerated (2 GPU per Host)			
Hosts	Cores	Abaqus tokens	Runtime (hours)	Normalized Cost *	Normalized Value **	Abaqus tokens	Runtime (hours)	Normalized Cost *	Normalized Value **
1	12	14	5.3	1.0	**1.00**	15	2.8	1.2	**1.57**
2	24	19	3.1	1.4	1.19	20	2.1	1.8	1.42
3	36	22	2.3	1.7	1.29	24	1.8	2.3	1.28
4	48	25	1.9	2.1	1.34	27	1.6	2.8	1.16

* Normalized Cost = hardware($) + software tokens($) / (Cost of 12 core cpu-only system). Low is good.

** Normalized Value = (Value of 12 core cpu-only system) / (Runtime * Normalized Cost). High is good.

MERCURY

GPU Use at Mercury Marine

- ✓ Deployed GPU accelerated system in 2012 upgrade
  - ✓ (4) 2x6-core hosts with dual Tesla M2070 (Fermi) GPU's
- ✓ Obtaining additional GPU hardware in 2014 upgrade
  - ✓ Multiple 2x8-core hosts with dual Tesla K20 (Kepler) GPU's
- ✓ Abaqus/Standard FEA uses GPU acceleration

- Awaiting GPU acceleration for Abaqus/Explicit FEA
- Evaluating GPU acceleration for CFD
  - Several routines use GPU acceleration already, but not Mercury Marine's core solvers at this time
- Virtualization?

Outline

- About Mercury Marine
- Engineering simulation capabilities
- Progression of High Performance Computing (HPC)
- GPU Focus
- Summary

Summary

- High Performance Computing has been an enabling technology at Mercury Marine
- Each incremental step enables Mercury Marine to better meet the demands of a fast paced product development cycle
- GPU's are the most recent step for our continually improving HPC capabilities
- Mercury Marine's GPU accelerated system delivers up to 47% faster run times than the comparable CPU-only system
- The GPU accelerated system has up to 57% higher value than the comparable CPU-only system
- Mercury Marine continues to evaluate new computing technology: **THE RACE NEVER STOPS!**

MERCURY

MERCURY