ANSYS

Realize Your Product Promise™

Advances in Accelerator-based CFD Simulation

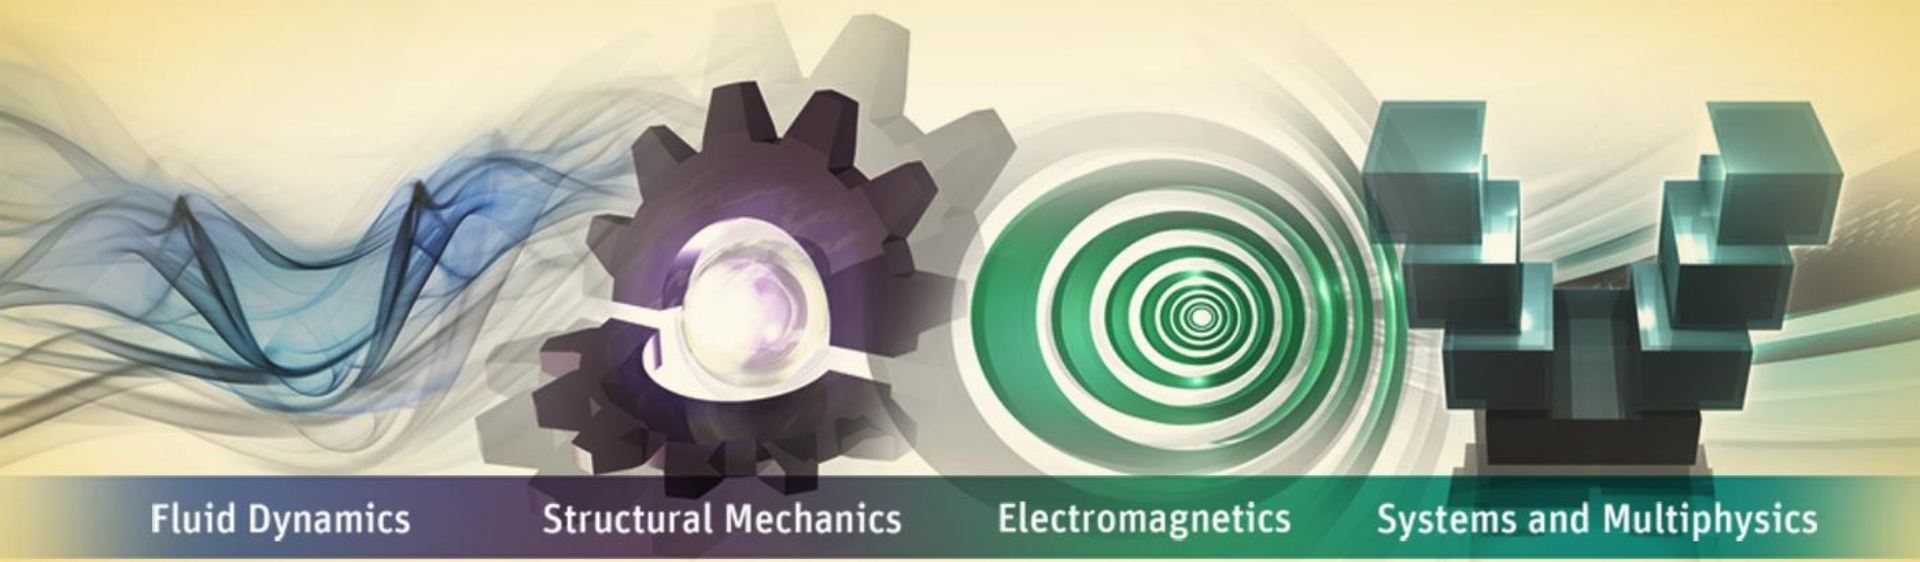

Wim Slagter, PhD

ANSYS, Inc.

TERATEC Forum, Workshop 6

June 24th, 2015

Overview

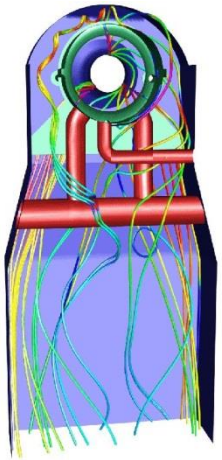

Introduction

Why Accelerator-based CFD

Status of Current Accelerator-based Solver Support

Guidelines

Licensing

Next Steps and Future Directions

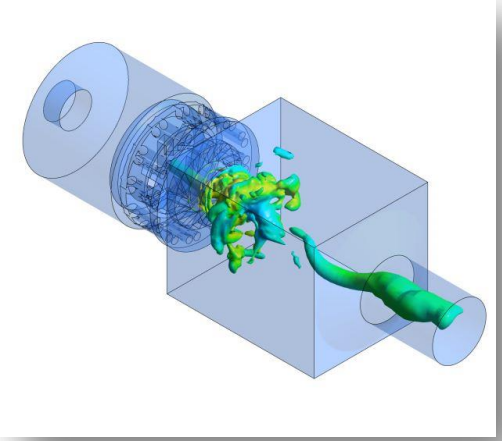

ANSYS, Inc.

FOCUSED

This is all we do.
Leading product technologies in all physics areas
Largest development team focused on simulation

TRUSTED

96 of the top 100

FORTUNE 500 Industrials
ISO 9001 and NQA-1 certified

CAPABLE

PROVEN

Recognized as one of the world's **MOST INNOVATIVE AND FASTEST-GROWING COMPANIES***

INDEPENDENT

Long-term financial stability
CAD agnostic

LARGEST

3x The size of our nearest competitor

**BusinessWeek, FORTUNE*

Industry Reach and Solution Offerings

Leading Disciplines

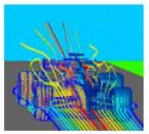
Fluids

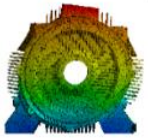
Structures

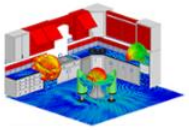
Electromagnetics

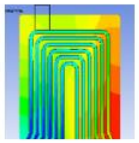
Thermal

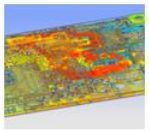
Power Integrity

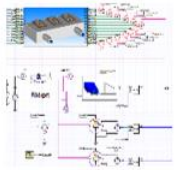
Systems

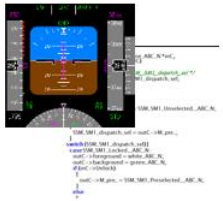
Embedded Software

Global Reach

PEOPLE **1,300+** **680+** **800+** **2,000+ CHANNEL PARTNERS**

CUSTOMERS **45,000+ GLOBALLY**

Industry Presence

Automotive

Aerospace & Defense

Academic

Construction

Consumer Goods

Electronics

Energy

Healthcare

Industrial Equip. & Rotating Machinery

Materials & Chemical Processing

Semiconductors

R&D at ANSYS

1,100+ software development professionals in 16 countries (>450 have PhDs)

Active research partnerships with >50 research institutions

180+ active partnerships with hardware and software providers

"Continual ANSYS research leads to advanced, more robust solutions for even the most complex problems. End users have confidence in analysis results, meaning they can rely less on costly physical testing."

Dr. Florian Menter
Research and Development Fellow
ANSYS

Why HPC?

Faster

Impact product design
Enable large models
Allow parametric studies

Courtesy of Red Bull Racing

Larger

Assemblies
CAD-to-mesh
Capture fidelity

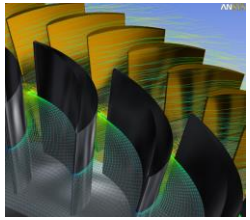

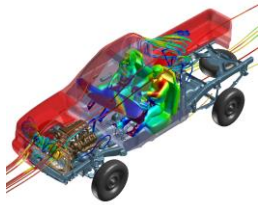

Extend

Turbulence
Combustion
Particle Tracking

More

Multiple design ideas
Optimize the design
Ensure product integrity

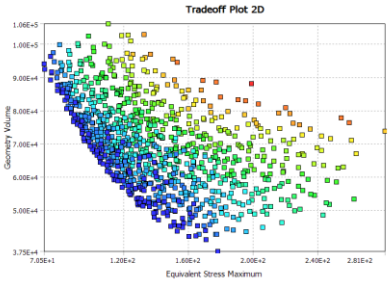

Scale-up of HPC

A Software Development Imperative

Clock Speed – Levelling off

Core Counts

- **Growing (Intel & AMD)**
- **Exploding (NVIDIA GPUs)**

Future performance depends on highly scalable parallel software

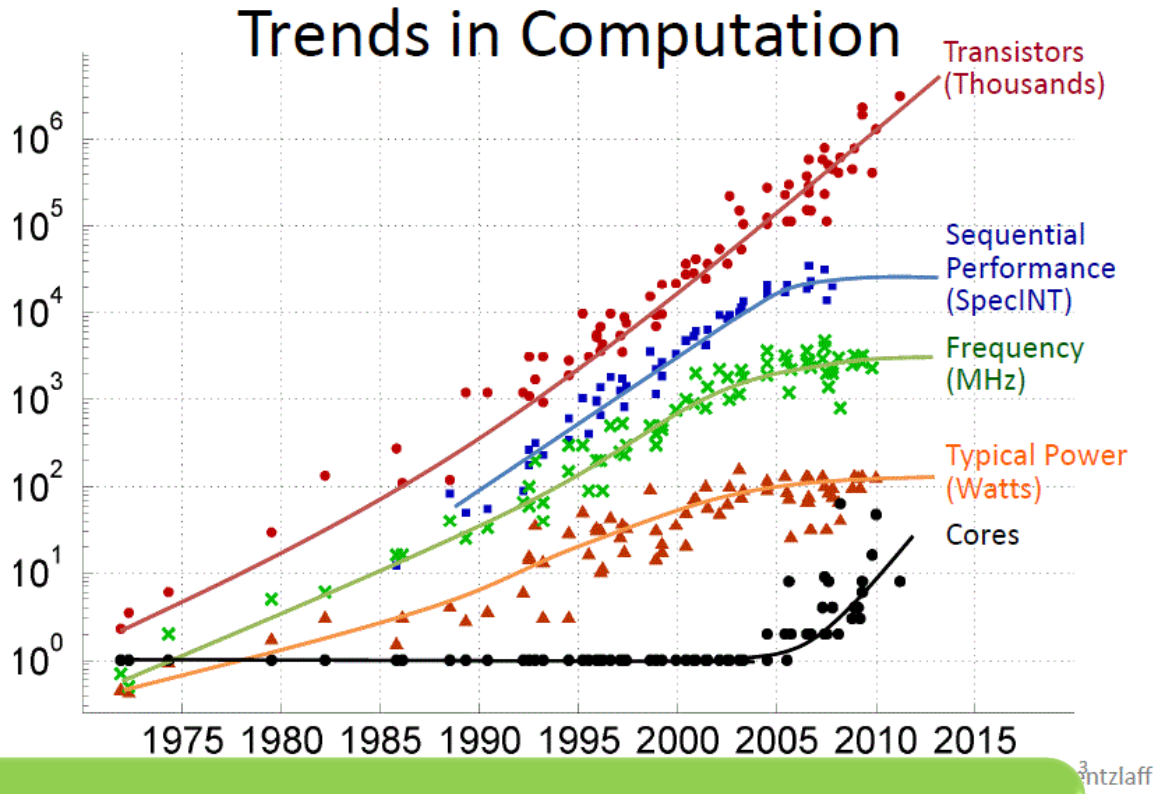

Today's multi-core / many-core hardware evolution makes HPC a software development imperative.

Evolvement of HPC Developments at ANSYS

2015

2010 - 2015
- ►Scaling to 36,864 cores (fluids)
- ►Hybrid parallelization (fluids)
- ►Network-aware partitioning (fluids)
- ►DDM for finite antenna arrays (HFSS 14)
- ►GPU acceleration with DMP(structures), AMG solver (fluids), and HFSS-Transient

2009
- ►Ideal scaling to 2048 cores (fluids)
- ►Teraflop performance at 512 core (structures)
- ►Parallel I/O (fluids)
- ►Domain Decomposition introduced (HFSS 12)

2007 - 2008
- ►Optimized performance on multicore processors
- ►1st One Billion cell fluids simulation

2005 -2007
- ►Distributed sparse solver
- ►Distributed PCG solver
- ►Variational Technology
- ►SYS released
- ►Distributed Solve (DSO) HFSS 10

2004
- ►1st company to solve 100M structural DOF

2000

1994
- ►Iterative PCG Solver Introduced for large structural analysis

1993
- ►1st general-purpose parallel CFD with interactive client-server user environment

2010 - 2015
- ►Scaling to 36,864 cores (fluids)
- ►Hybrid parallelization (fluids)
- ►Network-aware partitioning (fluids)
- ►DDM for finite antenna arrays (HFSS 14)
- ►GPU acceleration with DMP(structures), AMG solver (fluids), and HFSS-Transient

ANSYS is committed to maintaining performance leadership.

Scaling Improvements Release by Release

Continuous development effort to improve HPC scaling in Fluent

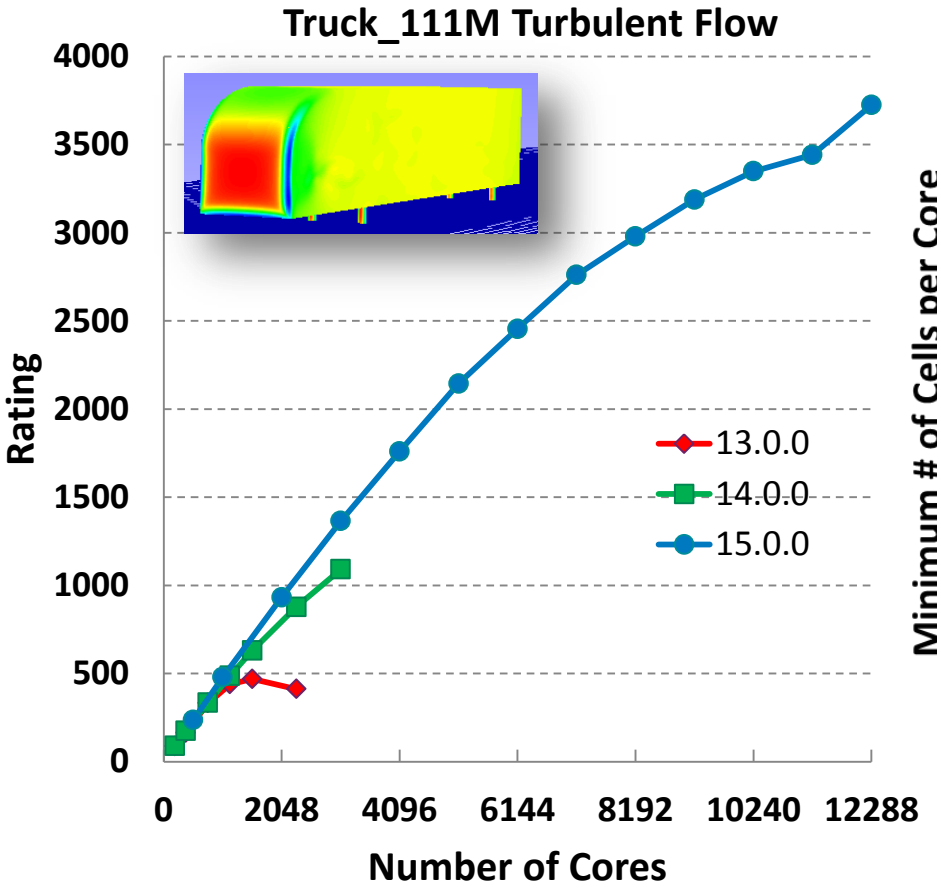

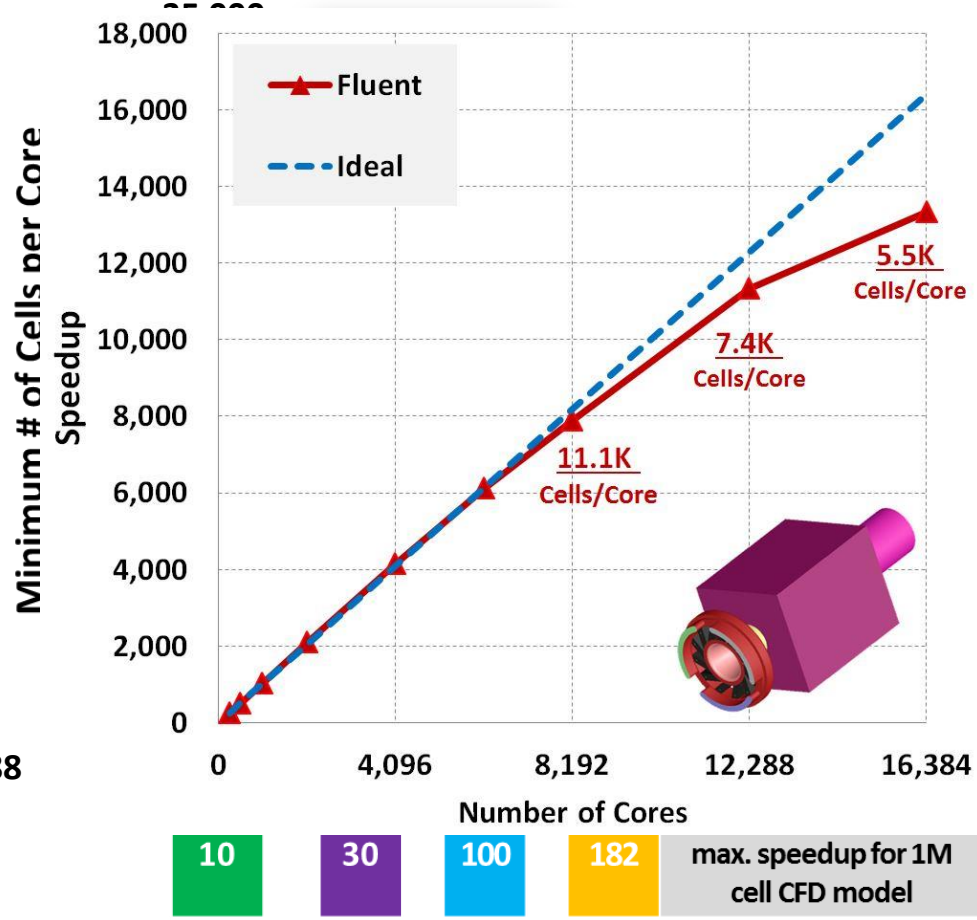

Rating is jobs per day.
A higher rating means faster performance.

Extreme Scaling – April 2014

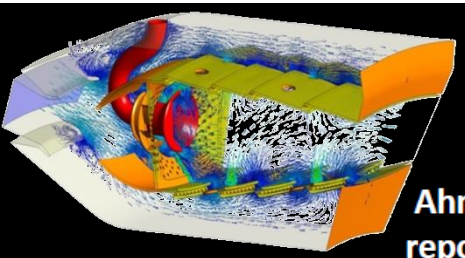

Scaling Breakthrough

Ahmed A. Taha, National Center for Supercomputing Applications (NCSA) reported scalability > 80% up to 20,480 cores for a 830 M case (April 2014)

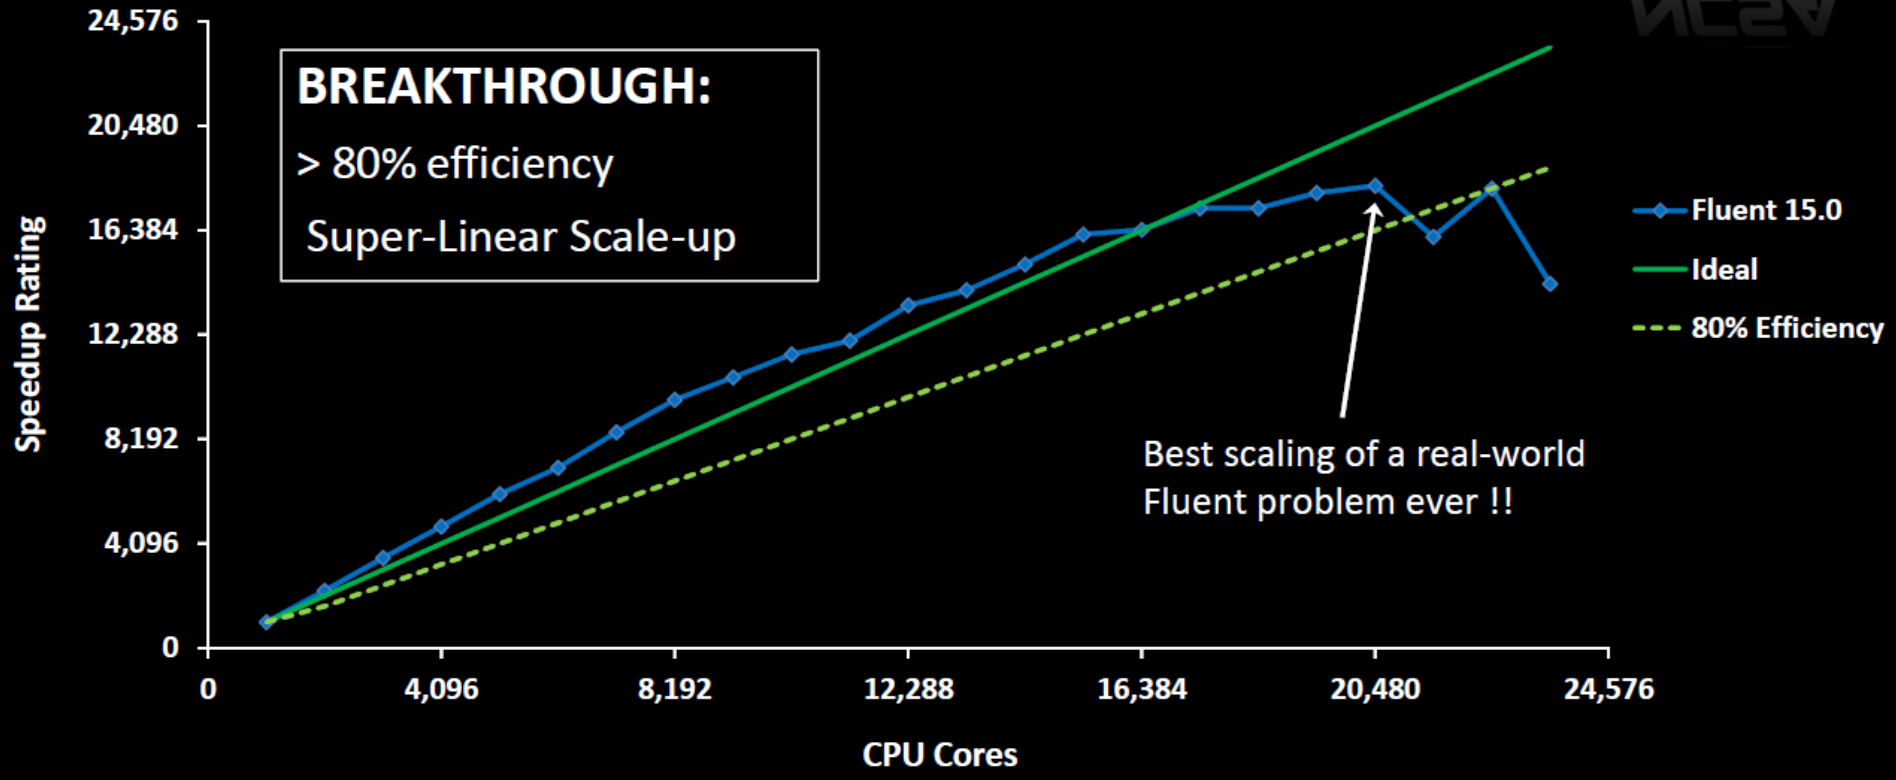

Source: "Industrial HPC Applications Scalability and Challenges", Seid Korić, ISC 2014

Extreme Scaling – December 2014

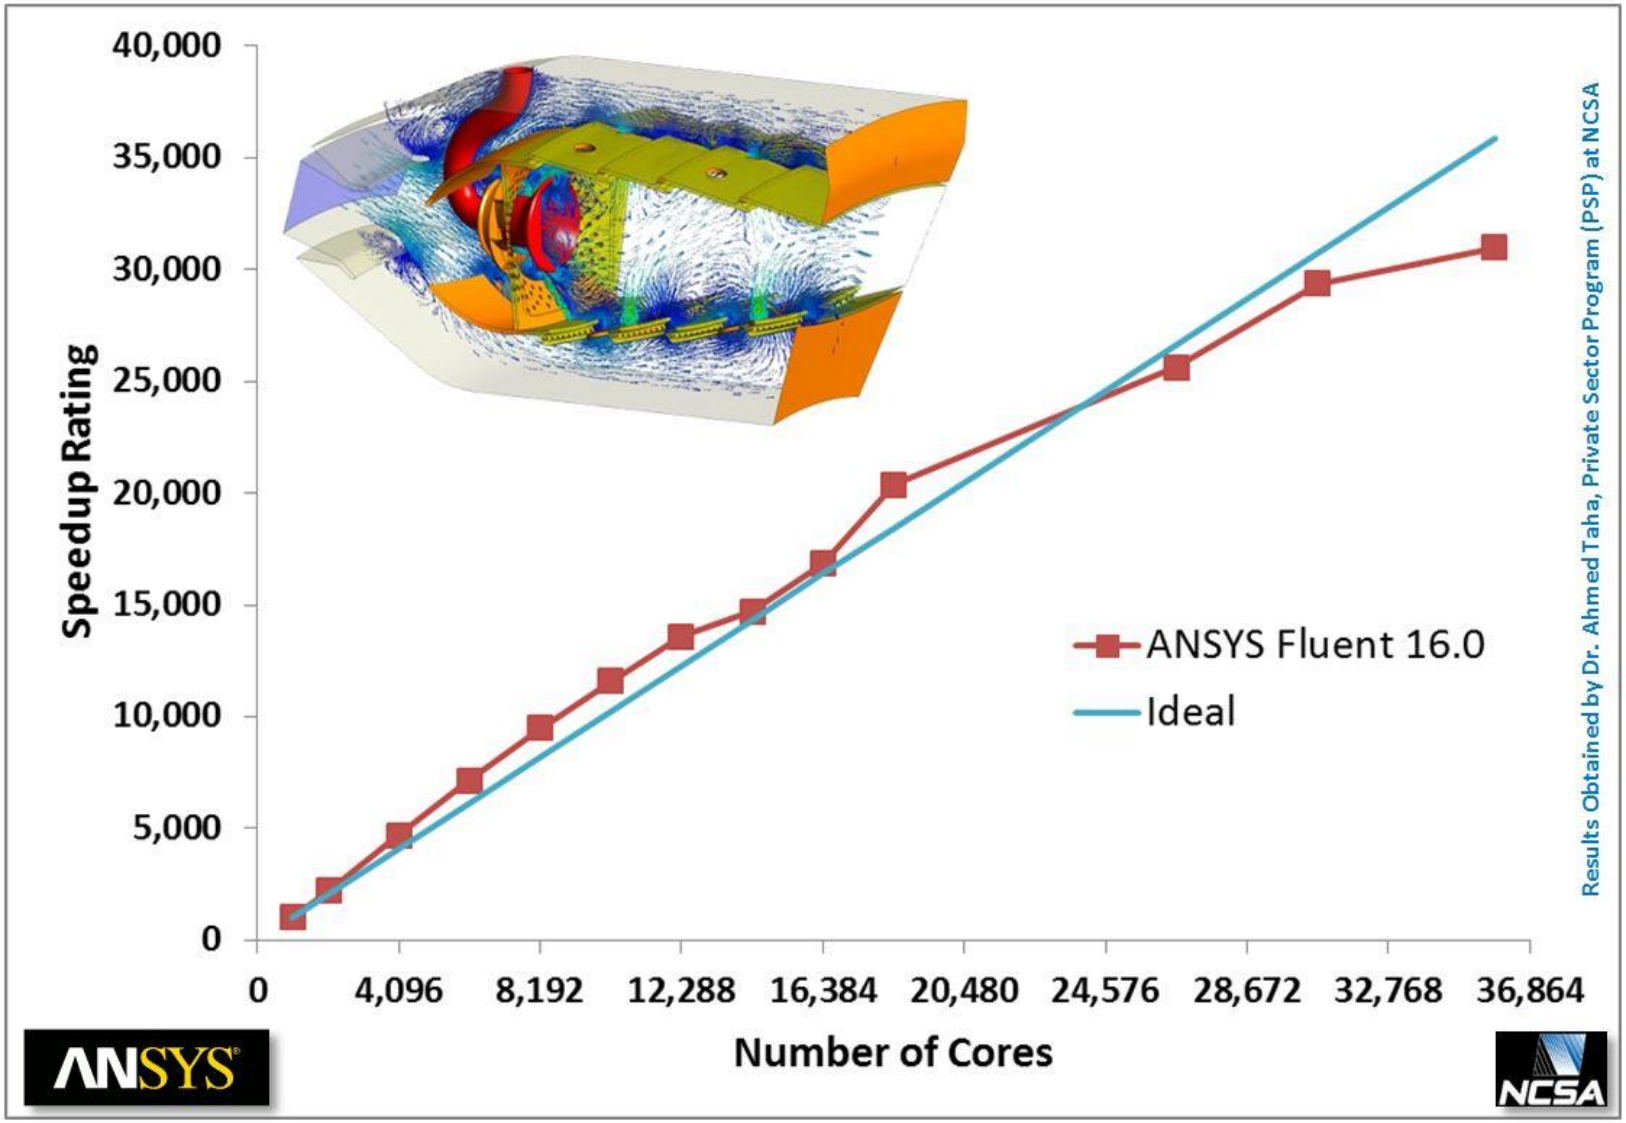

Click here for more info

Industry's HPC Challenges Go Even Further…

LES of a Powered Aircraft Configuration Across the Full Flight Envelope

- **Assess the ability to use CFD over the entire flight envelope, including dynamic maneuvers**
- **Assess the ability of CFD to accurately predict separated turbulent flows**
  - Monitor increasing LES region for hybrid RANS-LES simulations
  - Evaluate success of WMLES
  - Determine future feasibility of WRLES
- **Assess the ability to model or simulate transition effects**
- **Project future reductions in wind tunnel testing**

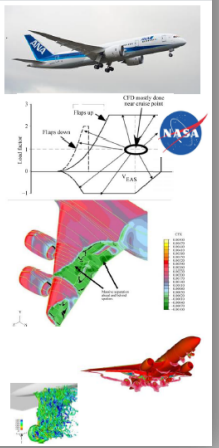

Source: "NASA Vision 2030 CFD Code – Final Technical Review", Contract # NNL08AA16B, November 14, 2013, NASA Langley Research Center

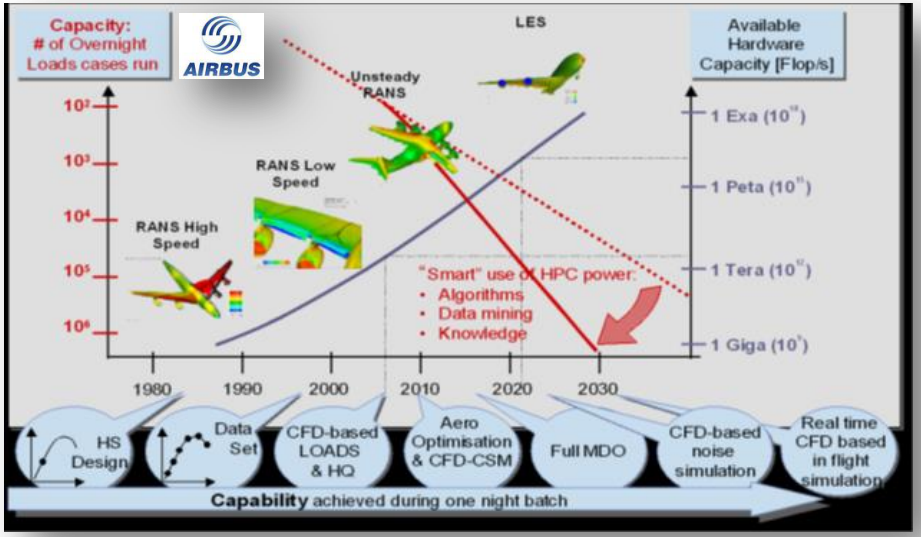

Source: "Exascale Challenges of European Academic & Industrial Applications", S. Requena, ISC'14, 22-26 June 2014, Leipzig

Active Nozzle – a key component for package sterilisation in new filling machine

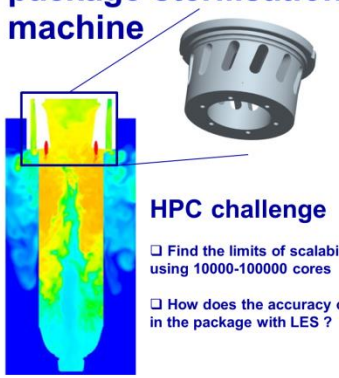

HPC challenge

- ❑ Find the limits of scalability with Large Eddy Simulation (LES) using 10000-100000 cores
- ❑ How does the accuracy change by fully resolve the boundary layer in the package with LES ?

LES simulation

MO 2012-04-16

Our HPC Ambitions

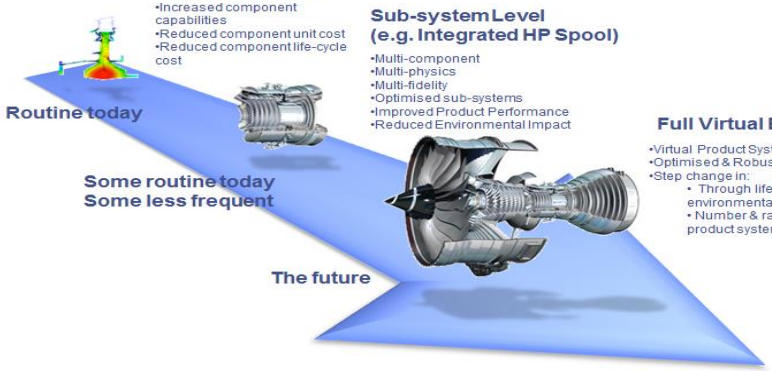

Rolls-Royce

Source: "Computational Science and Engineering Grand Challenges in Rolls-Royce", Leigh Lapworth, Networkshop42, 1-3 April 2014, University of Leeds

HPC Challenges & Emerging Technologies

Traditional CPU technology may be no longer capable of scaling performance sufficiently to address industry's HPC demand

HPC hardware challenges include:

- **Power consumption (limited)**
- **Energy efficiency ("Green Computing")**
- **Cooling (the lower the better)**

Hence the evolution of:

- **Quantum or bio-computing…**
- **Hardware accelerators:**
  - **Graphics Processing Units (GPUs) from NVIDIA and AMD**
  - **Intel® Xeon Phi™ coprocessors (previously called Intel MIC)**

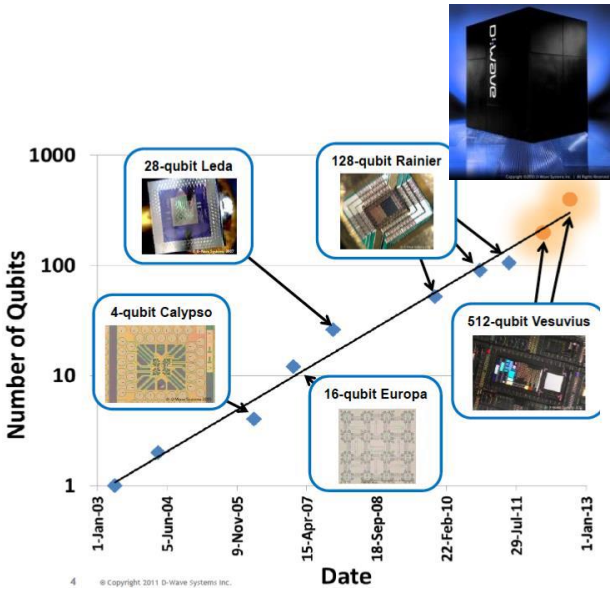

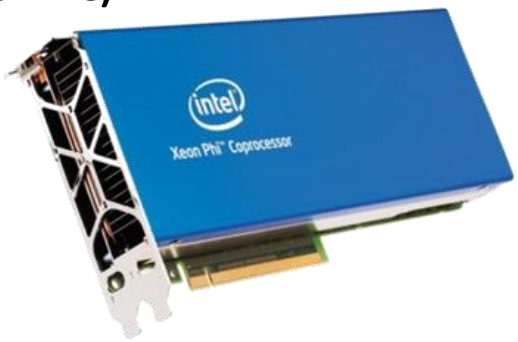

Motivations for Accelerator-Based CFD

- Accelerators are getting more powerful, e.g.
  - Number of GPU cores are increasing
  - GPU memory is getting bigger to the point where it can fit a large CFD problem
  - Intel Knight's Landing come into the market (addressing memory and I/O performance challenges)
- Problems do exist in CFD which can use large computing power
  - Coupled solver takes 60-70% time in solving the linear equation system
  - Stiff chemistry problems in species can take 90-95% time in ODE solver
  - Radiation models depending on their complexity can consume majority of the processing time
- HPC industry is moving toward heterogeneous computing systems, where CPUs and accelerators work together to perform general-purpose computing tasks
  - Supercomputing centers have been driving adoption of new accelerators for Top500-class machines
  - Delivering the highest performance energy efficiency

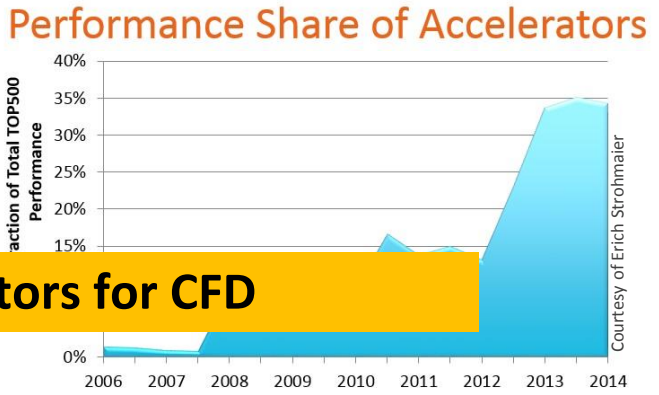

All good reasons to explore Accelerators for CFD

Evolution of GPU-Accelerated Solver Support

Beta Release in ANSYS Fluent 14.5

Full Product Support in Release ANSYS Fluent 15.0

GPU-based Model:	**Radiation Heat Transfer using OptiX, Product in R14.5**
GPU-based Solver:	**Coupled Algebraic Multigrid (AMG) PBNS linear solver**
Operating Systems:	**Both Linux and Win64 for workstations and servers**
Parallel Methods:	**Shared memory in R14.5; distributed memory in R15.0**
Supported GPUs:	**Tesla K40, Tesla K80 and Quadro 6000**
Multi-GPU Support:	**Single GPU for R14.5; full multi-GPU, multi-node R15.0**
Model Suitability:	**Size of 3M cells or less in R14.5; unlimited in R15.0**

GPU Value Proposition

- ANSYS Fluent 16.0

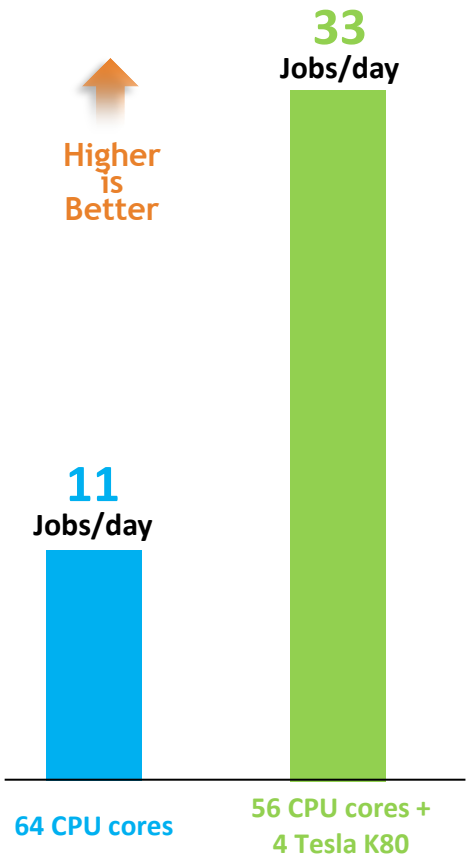

Truck Model

- External aerodynamics
- 14 million cells
- Steady, k-ε turbulence
- Coupled PBNS, DP
- 2 nodes each with dual Intel Xeon E5-2698 V3 (16 CPU cores) and dual Tesla K80 GPUs

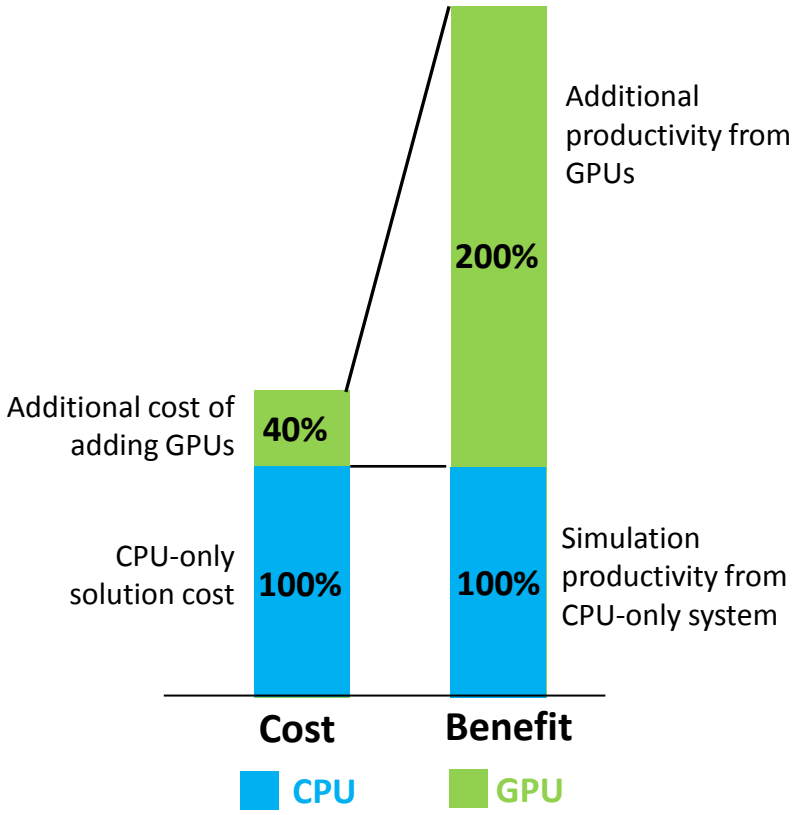

Simulation productivity (with an HPC Workgroup 64 license)

GPU Guidelines for ANSYS Fluent

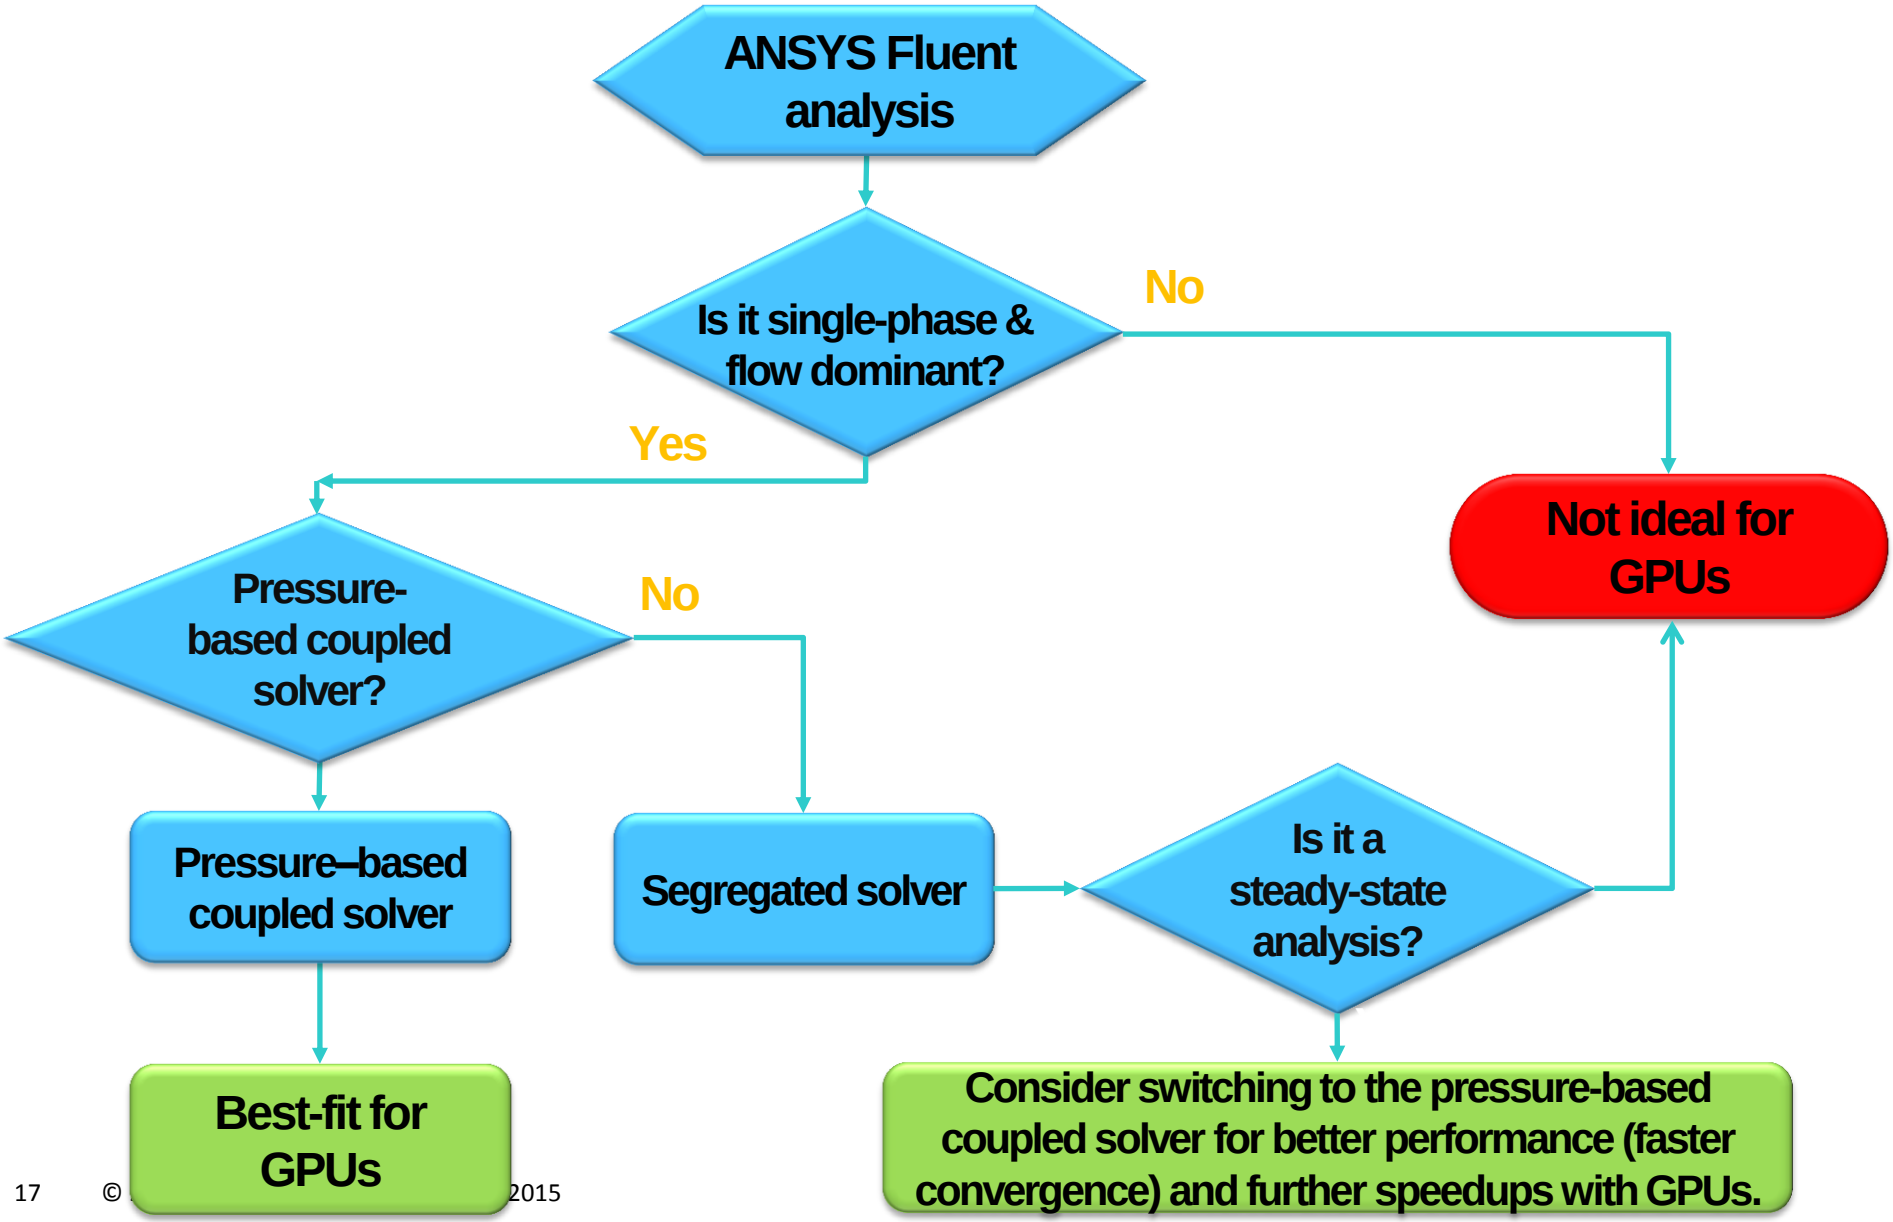

GPU Guidelines for ANSYS Fluent

GPUs accelerate the AMG solver portion of the CFD analysis, thus benefit problems with relatively high %AMG

- **Coupled solvers have high %AMG in the range of 60-70%**
- **Fine meshes and low-dissipation problems have high %AMG**

In some cases, pressure-based coupled solvers offer faster convergence compared to segregated solvers (problem-dependent)

The whole problem must fit on GPUs for the calculations to proceed

- **In pressure-based coupled solver, each million cells need approx. 4 GB of GPU memory**
- **High-memory cards such as Tesla K80 or Quadro K6000 are recommended**

Moving scalar equations such as turbulence may not benefit much because of low workloads (using 'scalar yes' option in 'amg-options')

Better performance on lower CPU core counts

- **A ratio of 3 or 4 CPU cores to 1 GPU is recommended**

ANSYS Fluent 15.0

- Power Consumption Study

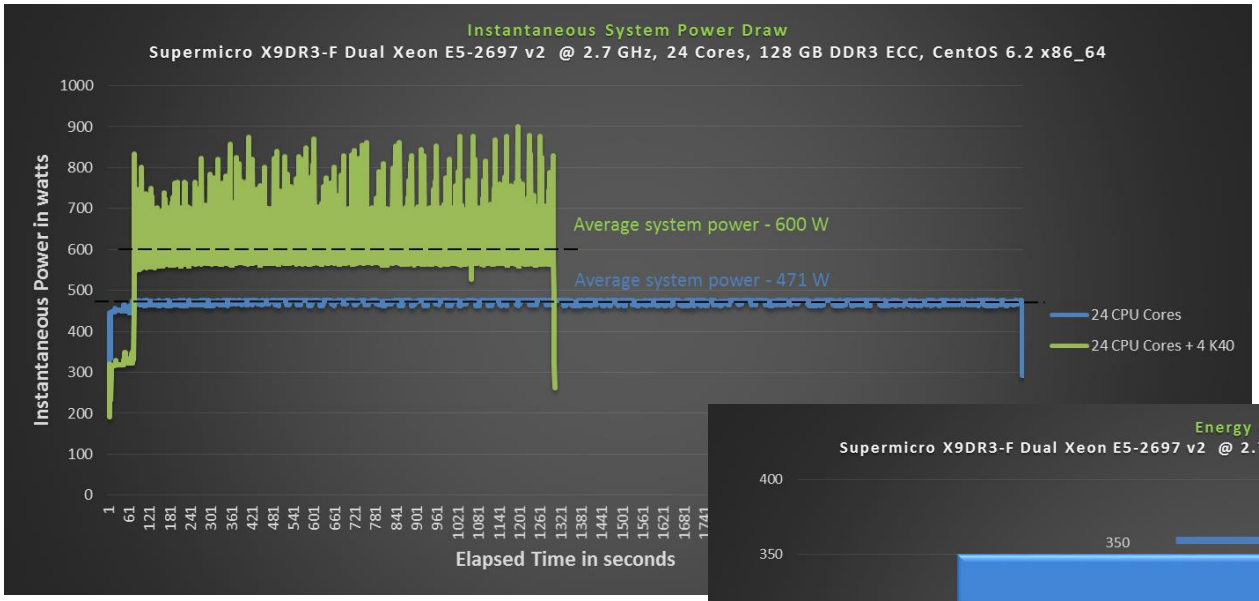

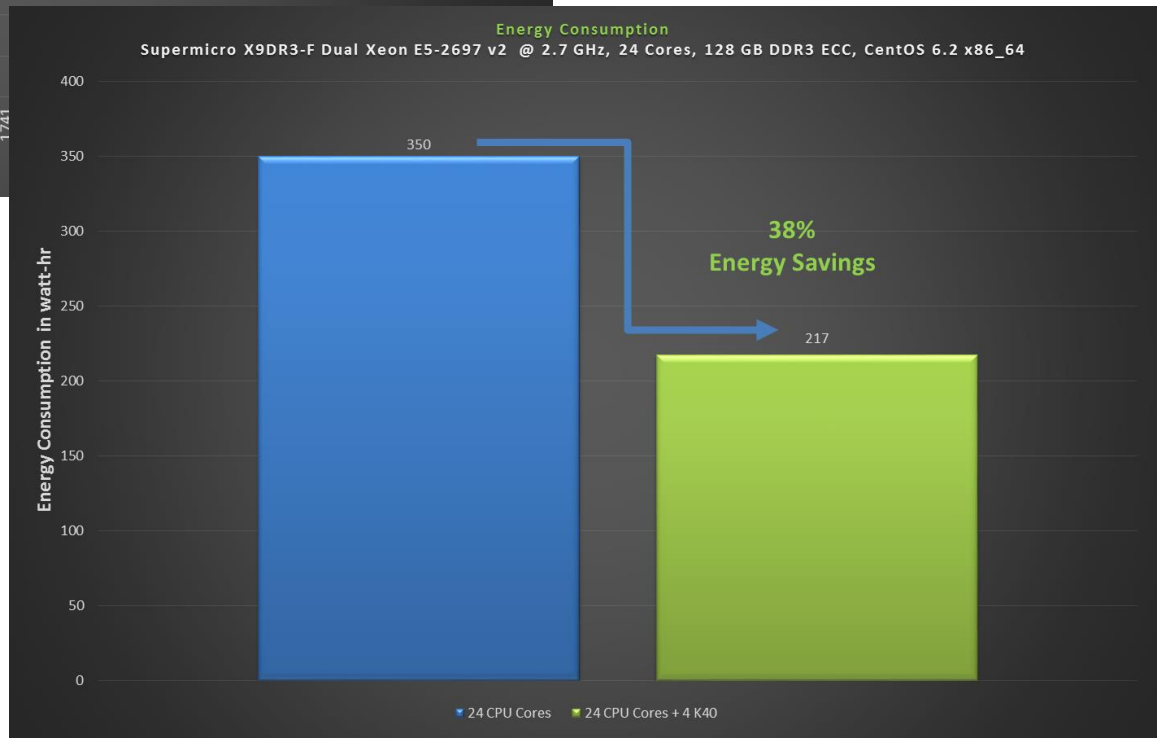

- **Adding GPUs to a CPU-only node resulted in 2.1x speed up while reducing energy consumption by 38%**

Benefit of GPU-Accelerated Workstation

- Shorter Time to Solution

Application Example

Objective

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

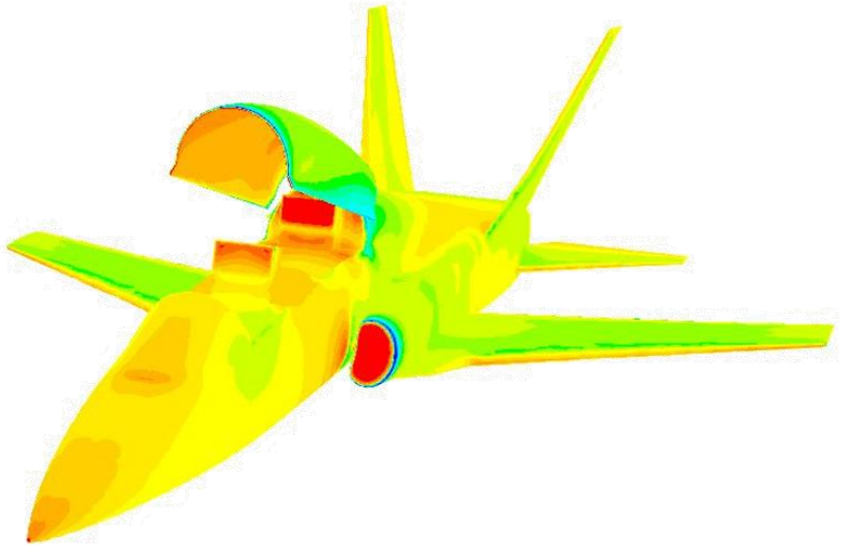

ANSYS Solution

- 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.

Intel Xeon Phi-Accelerated Solver Developments

Full Release in ANSYS Mechanical 16.0

Beta Release in ANSYS Fluent 16.0

MIC-based Solver:	**Sparse Direct Solver; all but DPM, raytracing most suited**
Operating Systems:	**Linux and Windows platform support; ditto R16.0**
Parallel Methods:	**SMP & DMP; distributed memory in R16.0**
Supported cards:	**Xeon Phi models 7120, 5110, 3120; ditto R16.0**
Multi-MIC Support:	**Multi-cards; and multi-node in R16.0**
Model Suitability:	**Model size is bound by physical memory on the card**

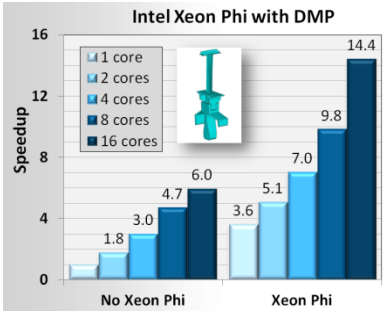

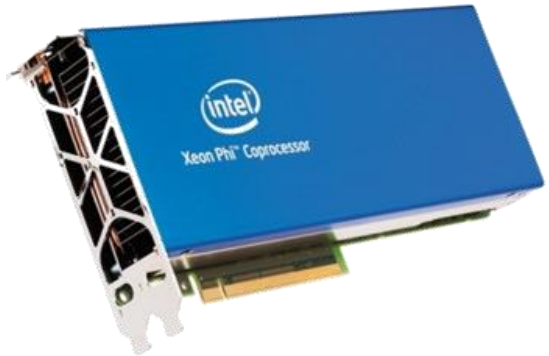

Intel Xeon Phi Coprocessor Product Lineup

7 Family
Highest Performance
Most Memory

Performance

16GB GDDR5

352GB/s

>1.2TF DP

300W TDP

7120P
MM# 927499

7120A
MM# 934878

7120D
(Dense Form Factor)
MM# 932330

7120X
(No Thermal Solution)
MM# 927498

5 Family
Optimized for High Density Environments
Performance/Watt

8GB GDDR5

>300GB/s

>1TF DP

225-245W TDP

5110P
MM# 924044

5120D (no thermal)
MM# 927503

3 Family
Parallel Computing Solution

Performance/$

6GB GDDR5

240GB/s

>1TF DP

300W TDP

3120P
MM# 927501

3120A
MM# 927500

HPC Licensing Enabling GPU Acceleration - One HPC Task Required to Unlock one GPU!

Licensing Examples:

1 x ANSYS HPC Pack

Total 8 HPC Tasks (4 GPUs Max)

2 x ANSYS HPC Pack

Total 32 HPC Tasks (16 GPUs Max)

.
.
.
.

Example of Valid Configurations:

6 CPU Cores + 2 GPUs

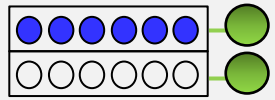

4 CPU Cores + 4 GPUs

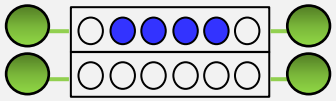

24 CPU Cores + 8 GPUs

(Total Use of 2 Compute Nodes)

(Applies to all license schemes: ANSYS HPC, ANSYS HPC Pack, ANSYS HPC Workgroup)

Next Steps and Future Directions

Next steps on "How to use GPUs"

- Article: "Accelerating ANSYS Fluent with NVIDIA GPUs"
- Recorded webinar: "How to Speed Up ANSYS 15.0 with GPUs"
- Technical brief: "Accelerating ANSYS Fluent 15.0 Using NVIDIA GPUs"
- Recorded webinar: "Understanding Hardware Selection for ANSYS 15.0"

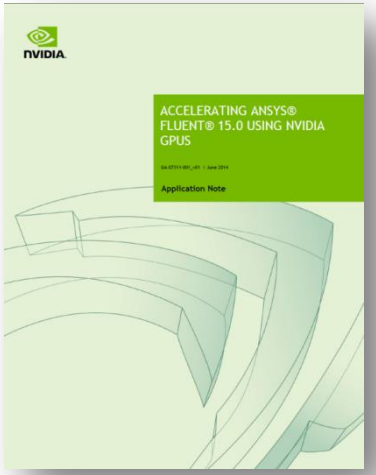

Future directions:

- Accelerate radiation modeling with discrete ordinate method by using AmgX
- Provide user control to pick and choose which equation to run on GPU
- Explore possibilities of further improvements via use of advanced AmgX features like direct GPU communication
- Explore possibilities of performance improvements for segregated solver

Thank You!

- **Connect with Me**
  - wim.slagter@ansys.com
- **Connect with ANSYS, Inc.**
  - LinkedIn ANSYSInc
  - Twitter @ANSYS_Inc
  - Facebook ANSYSInc
- **Follow our Blog**
  - ansys-blog.com