Discover your design quicker as before with HPC

Alexander Dopf, CADFEM

PRACE Autumn School 2013 - Industry Oriented HPC Simulations, September 21-27, University of Ljubljana, Faculty of Mechanical Engineering, Ljubljana, Slovenia

ANSYS

Realize Your Product Promise™

Introduction

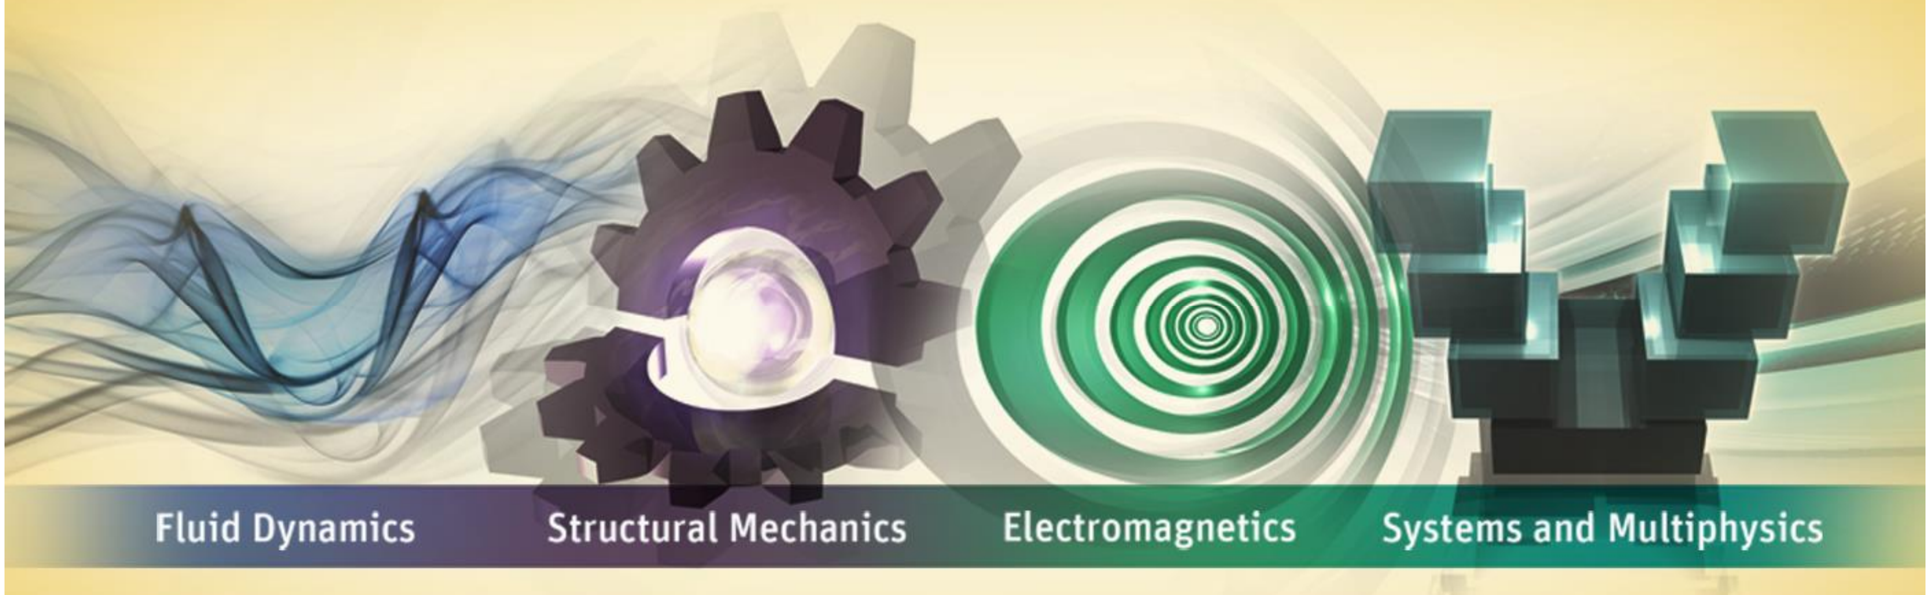

Need for Increased Performance & Robustness

Products in the headlines

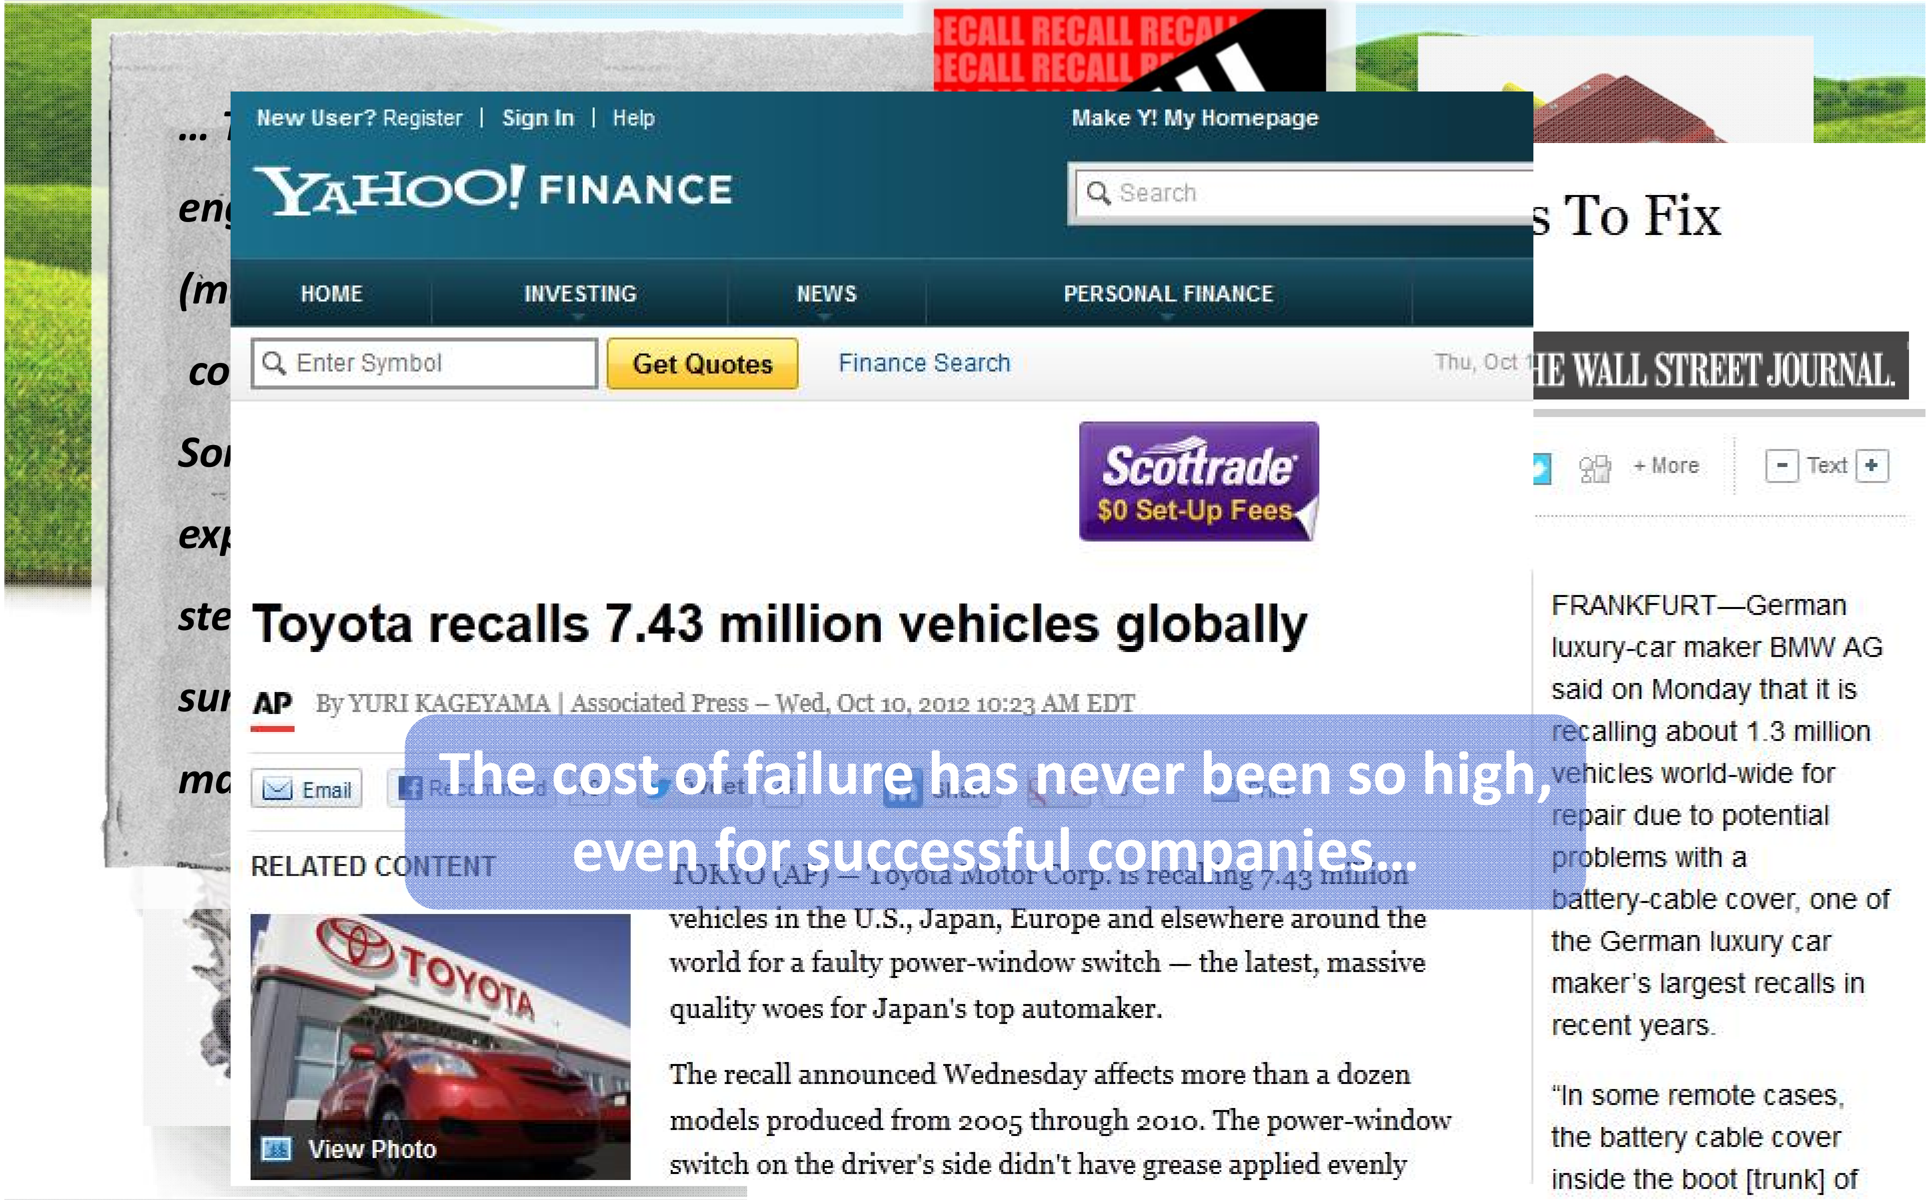

The cost of failure has never been so high, even for successful companies…

Product Integrity via Robust Design

Building in Product Reliability

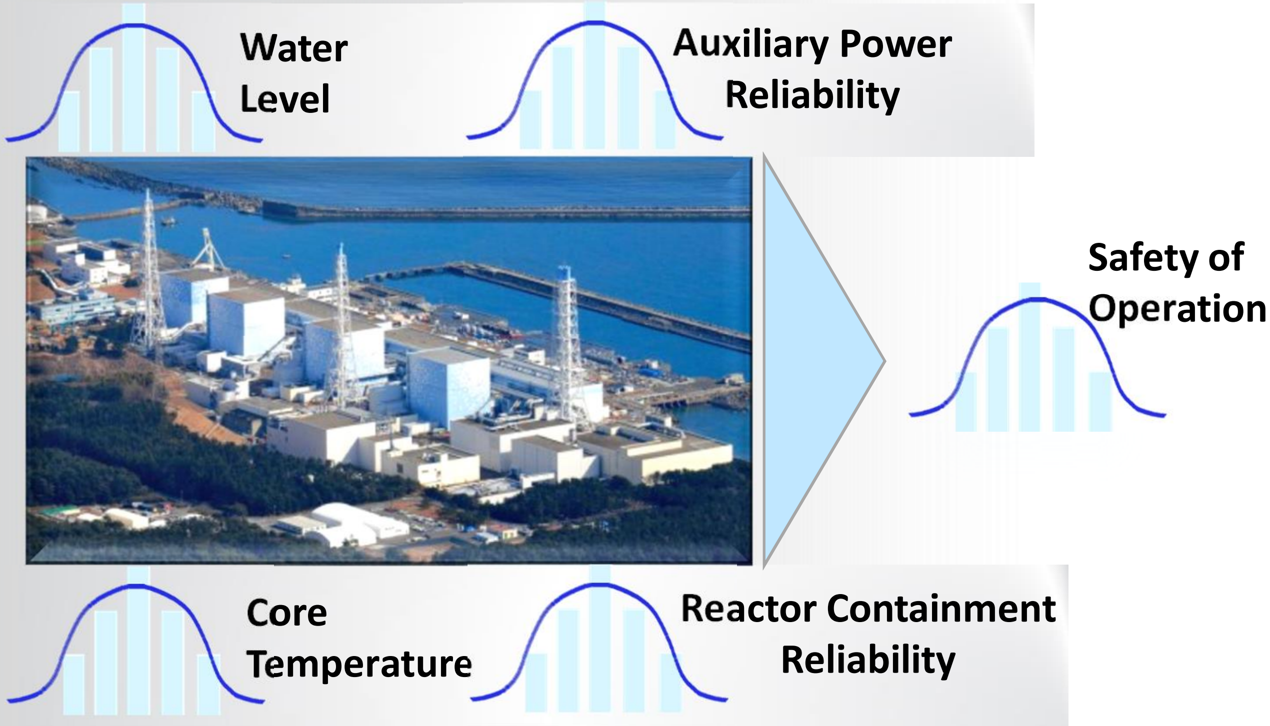

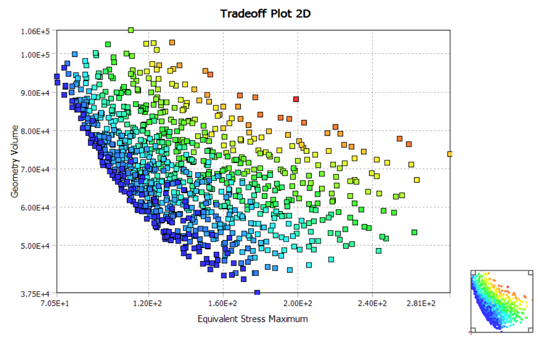

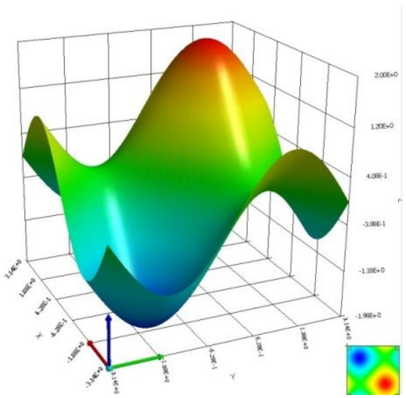

Variations in operating conditions, manufacturing processes and material properties create uncertainty in the overall success of a product design.

Need for HPC

Larger

Assemblies
CAD-to-mesh
Capture fidelity

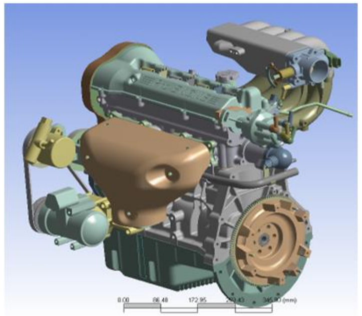

Faster

Impact product design
Enable large models
Allow parametric studies

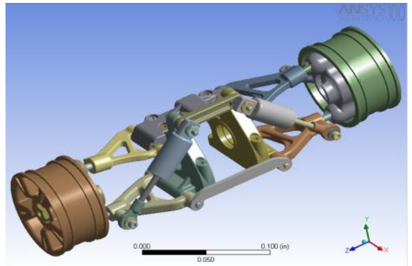

Extend

Modal
Nonlinear
Multiphysics
Dynamics

More

Multiple design ideas
Optimize the design
Ensure product integrity

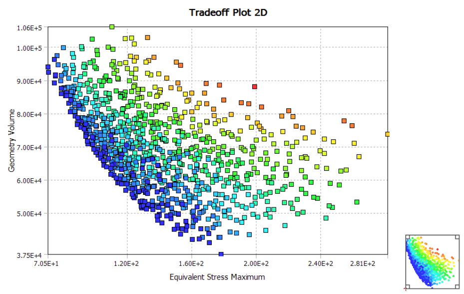

Summary

It's all about getting *better* insight into product behavior *quicker*!

HPC enables high-fidelity

Å Include details - for reliable results

Å "Getting it right the first time"

Å Innovate with confidence

HPC enables design exploration & optimization

Å Consider multiple design ideas

Å Optimize the design

Å Ensure performance across range of conditions

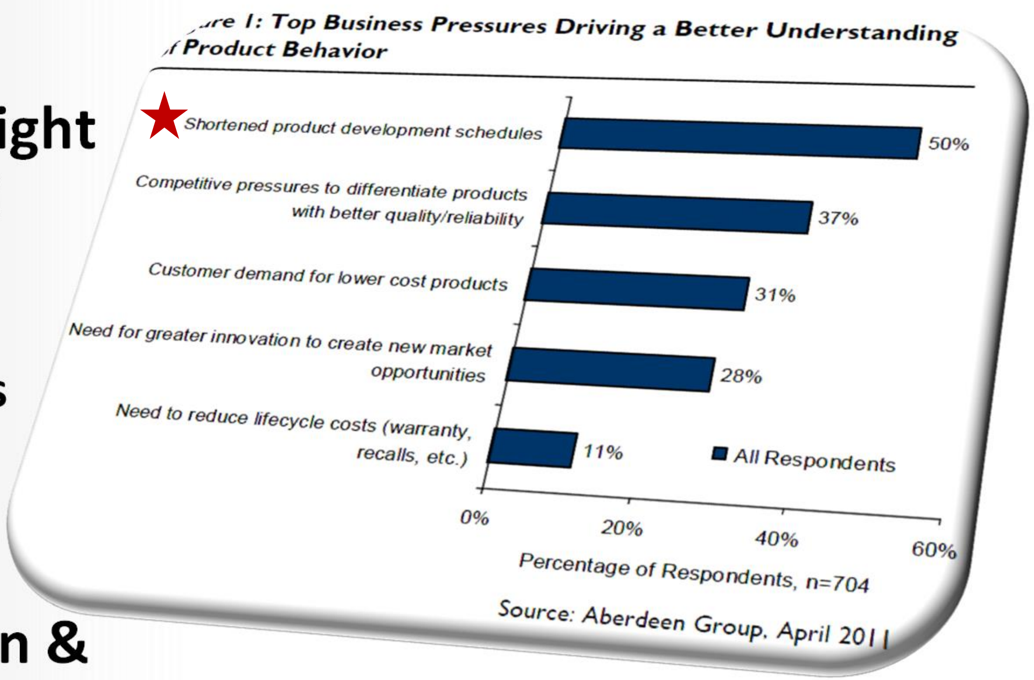

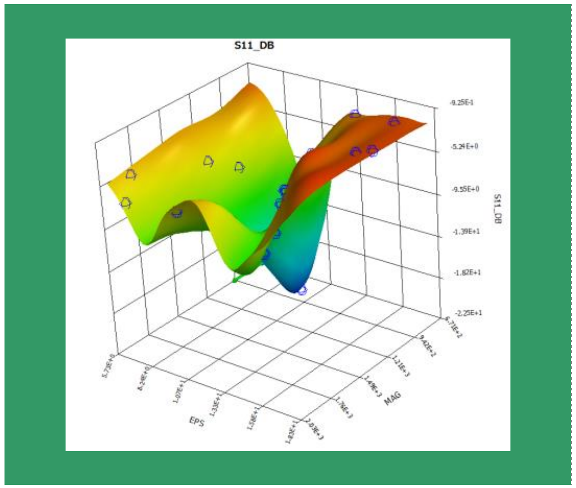

Realize Your Product Promise™

Larger Simulations

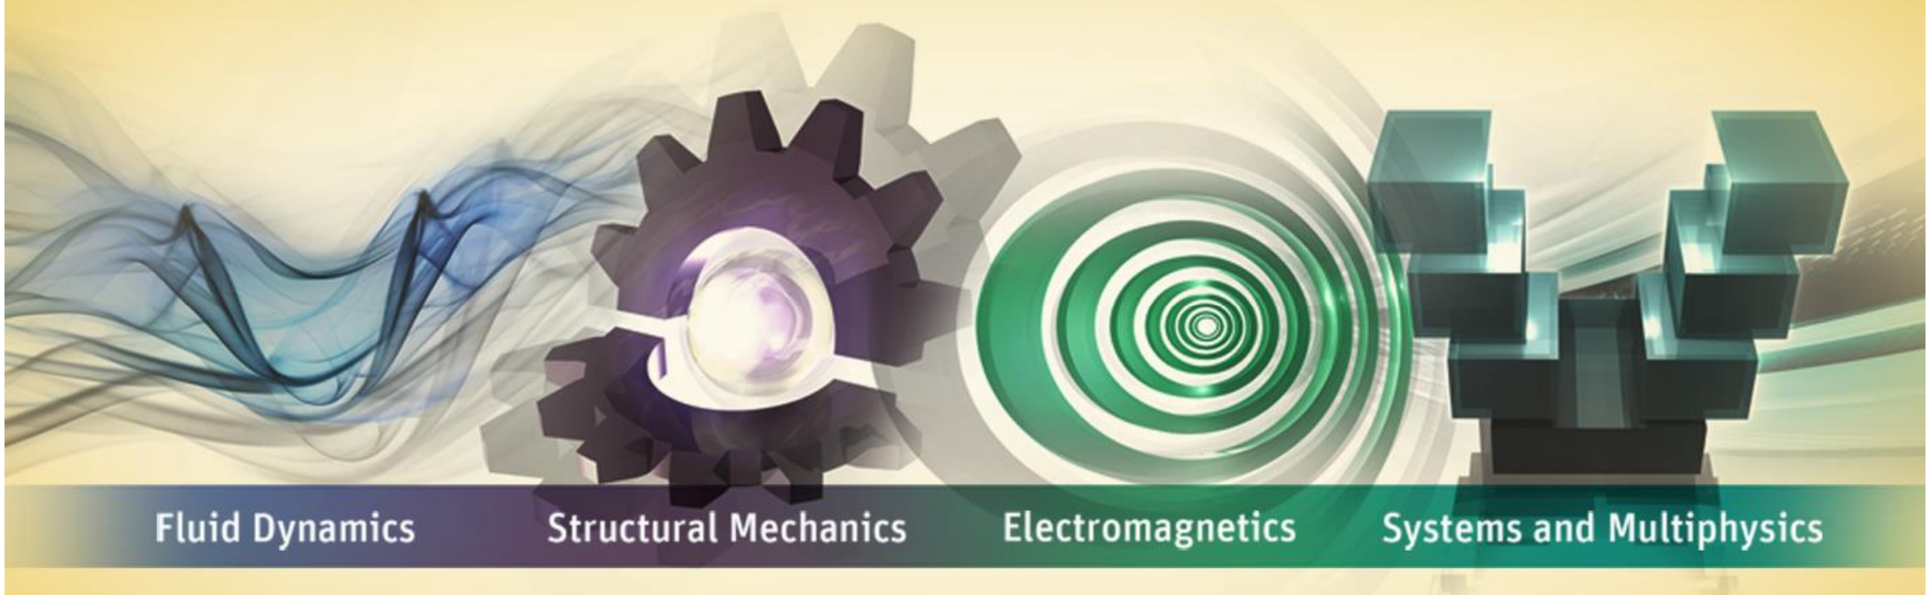

HPC – A Software Development Imperative

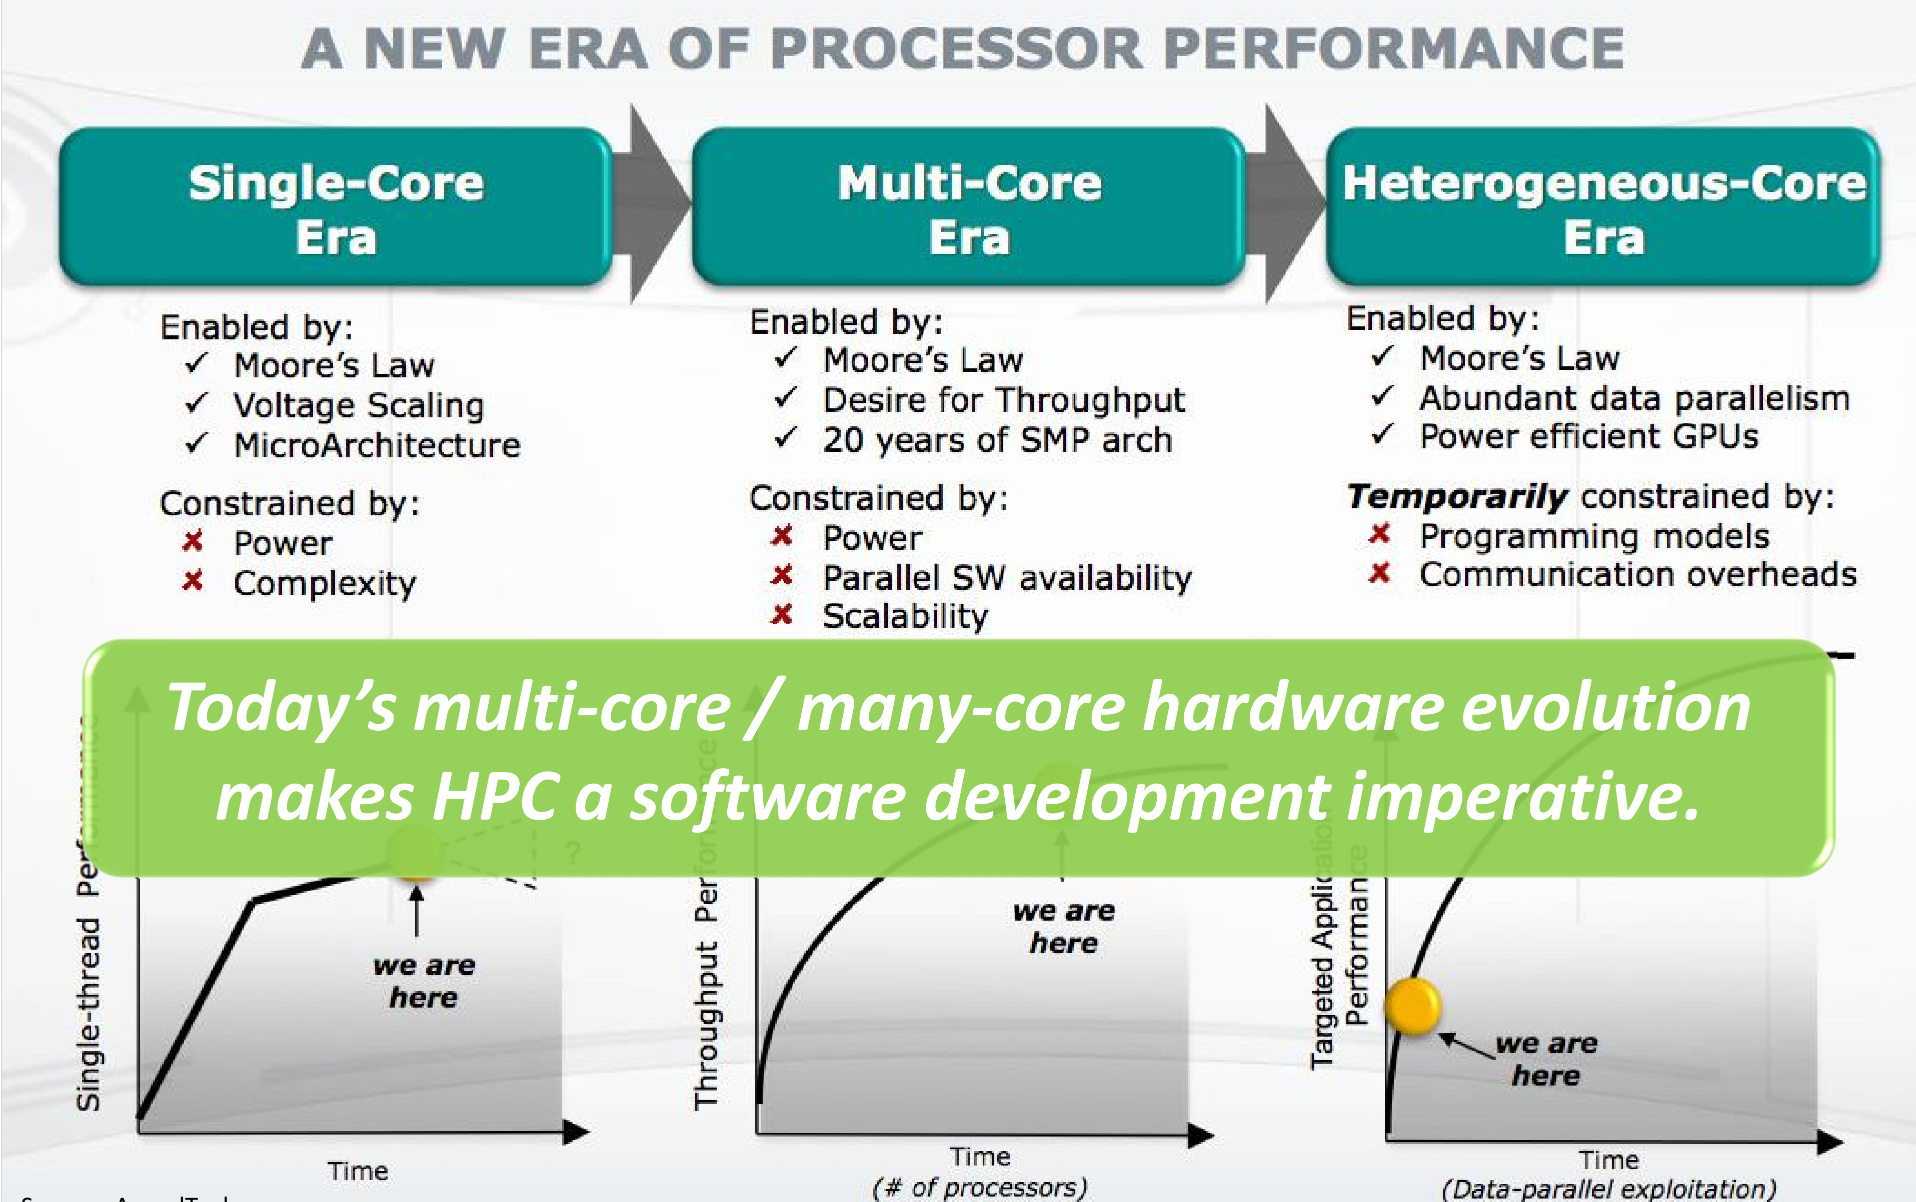

Today's multi-core / many-core hardware evolution makes HPC a software development imperative.

Source: AnandTech

ANSYS Mechanical Scaling Achievement @ 13.0

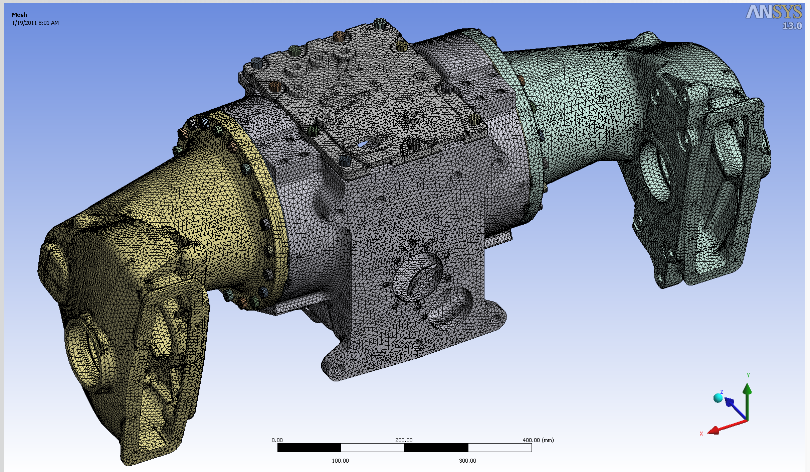

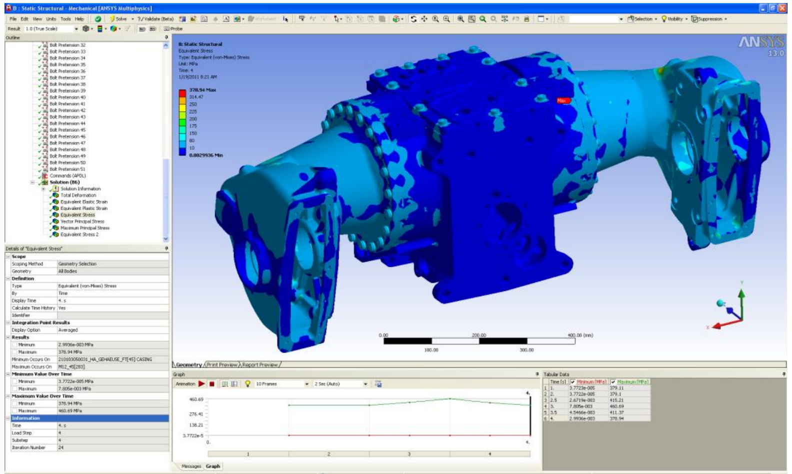

6 Mio Degrees of Freedom
Plasticity, Contact
Bolt pretension
4 load steps

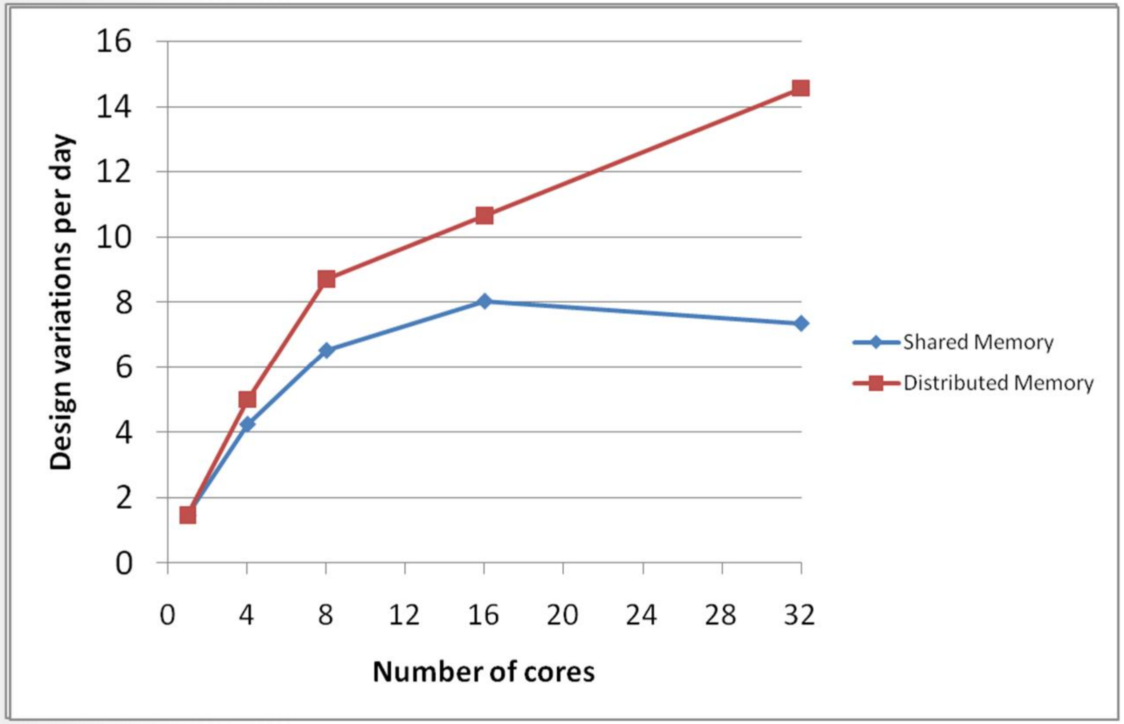

ANSYS

What about GPU Computing?

CPUs and GPUs work in a collaborative fashion

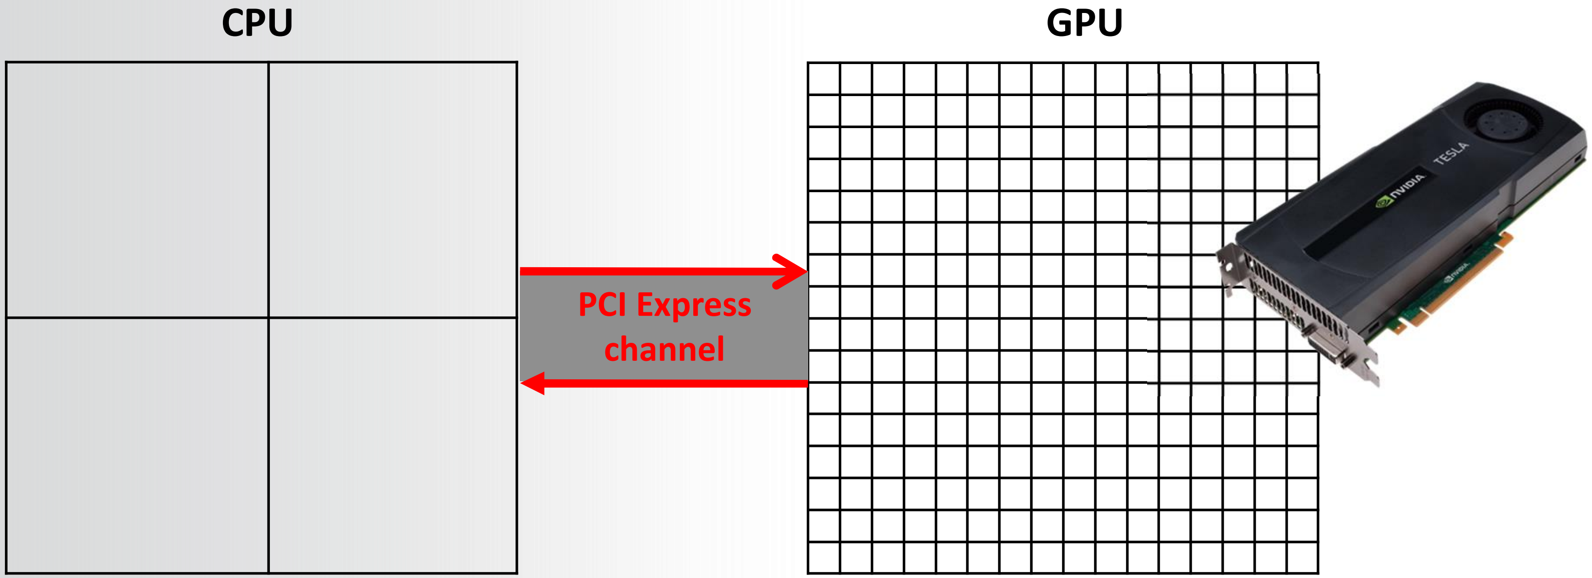

Multi-core processors

- Typically 4-6 cores
- Powerful, general purpose

Many-core processors

- Typically hundreds of cores
- Great for highly parallel code, within memory constraints

Optimized Solver Performance

- ANSYS Mechanical

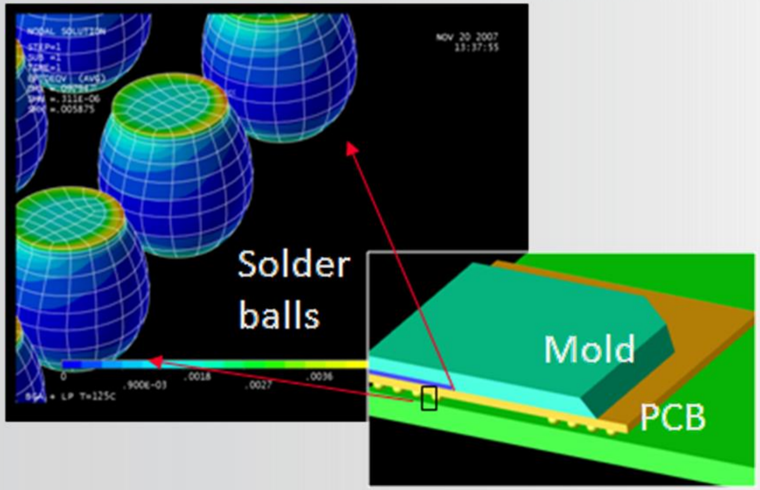

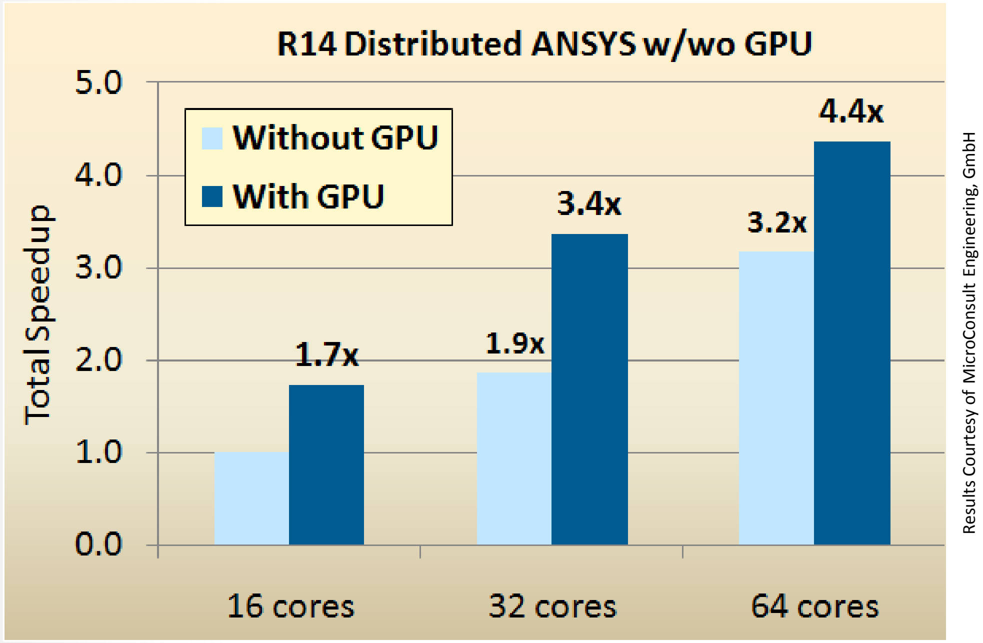

GPU Acceleration can be used with Distributed ANSYS to combine the advantage of GPU technology and the power of distributed ANSYS

ANSYS and NVIDIA Collaborations

Release	ANSYS Mechanical 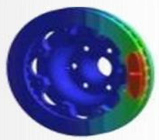	ANSYS Fluent 
13.0 Dec 2010	SMP, Single GPU, Sparse and PCG/JCG Solvers	
14.0 Dec 2011	+ Distributed ANSYS; + Multi-node Support	Radiation Heat Transfer (beta)
14.5 Oct 2012	+ Multi-GPU Support; + Hybrid PCG; + Kepler GPU Support	+ Radiation HT; + GPU AMG Solver (beta), Single GPU

ANSYS Mechanical SMP – GPU Speedup @ 14.0

Modal analysis of a radial impeller

Å Block Lanczos Eigensolver

Å Cyclic symmetry model with 2 million DOF:

- 337916 nodes
- 222725 elements
- 10-node tetrahedral solid element

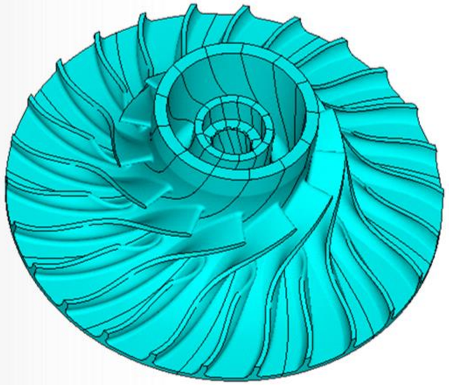

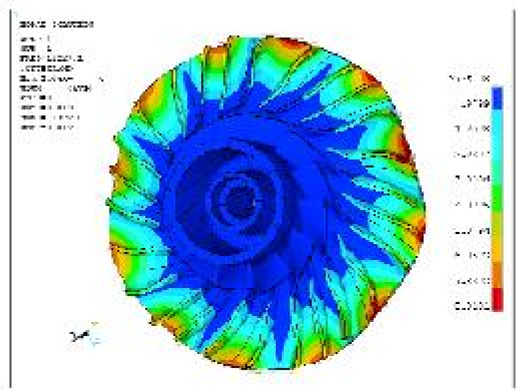

Results (baseline is 1 core):

Å With GPU, ~6x speedup on 1 core

Å ~8.5x speedup on 4 cores

Å If 2 cores is taken as baseline instead, 2 cores with GPU Accelerator results in 3.7x speedup!

Cores	GPU	Speedup
1	no	1.00
2	no	1.99
4	no	3.61
1	yes	5.92
2	yes	7.43
4	yes	8.52

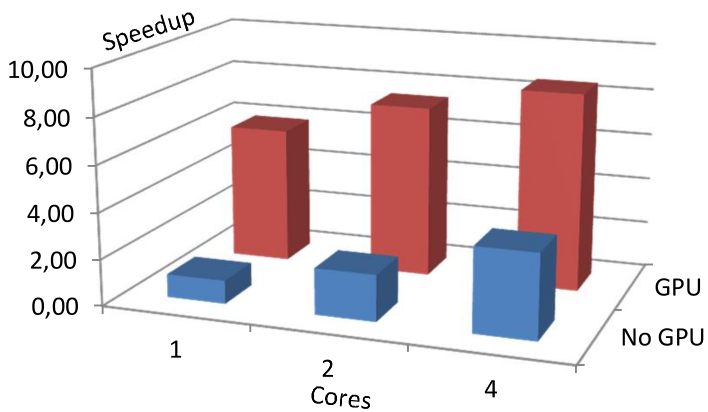

Windows workstation: Two Intel Xeon 5530 processors (2.4 GHz, 8 cores total), 48 GB RAM, NVIDIA Quadro 6000

Optimized Solver Performance
- ANSYS Fluent

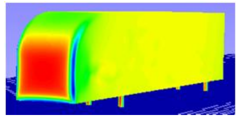

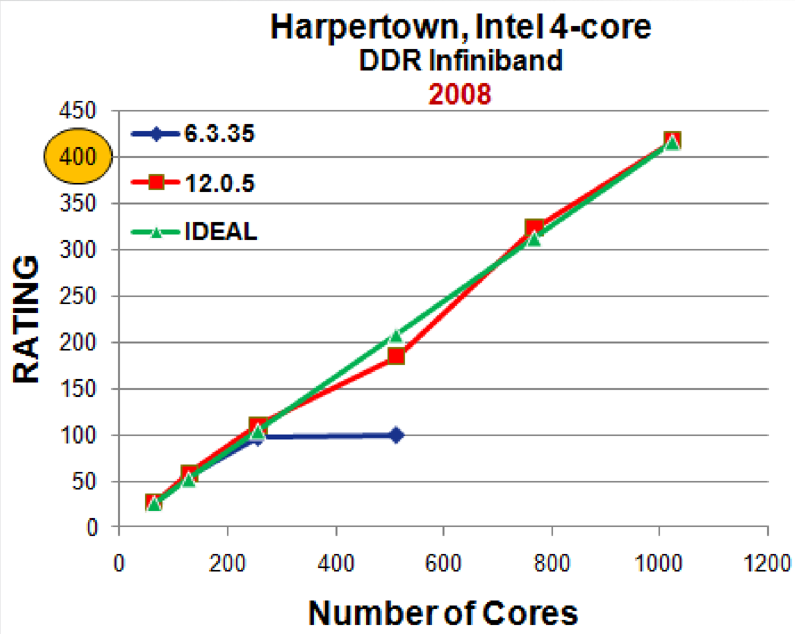

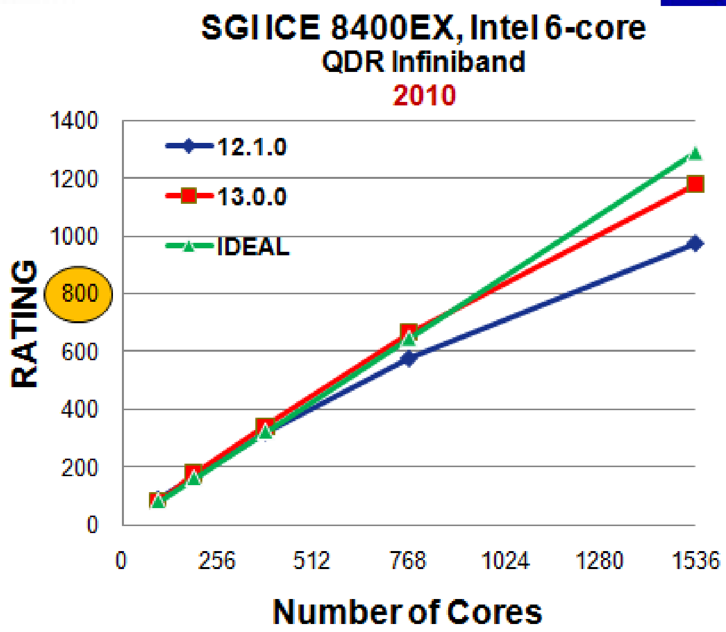

Continuous performance improvements version over version

Parallel scalability near ideal (98%+)!

Demonstrable ability to solve large problems on large clusters very efficiently

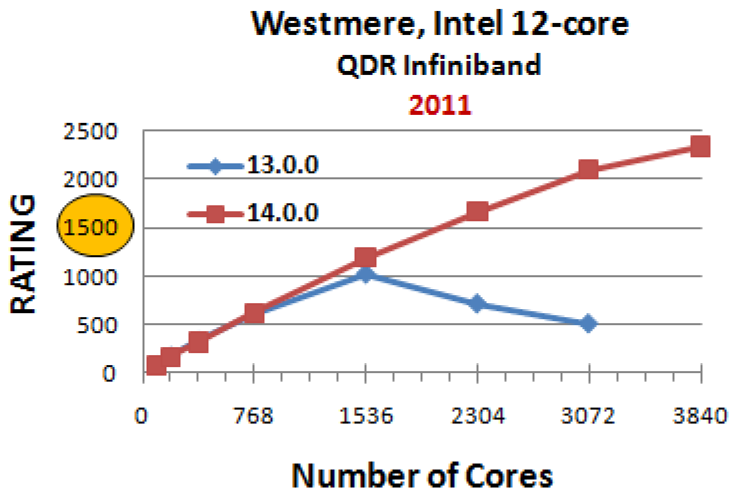

Optimized Solver Performance
- ANSYS Fluent

Hybrid parallelism for best performance on multi-core chips within clusters

Fast Parallel I/O

Architecture-aware partitioning

Good scalability for simulations with monitors enabled

Improved Scaling with Hybrid Parallelism - Nehalem EX

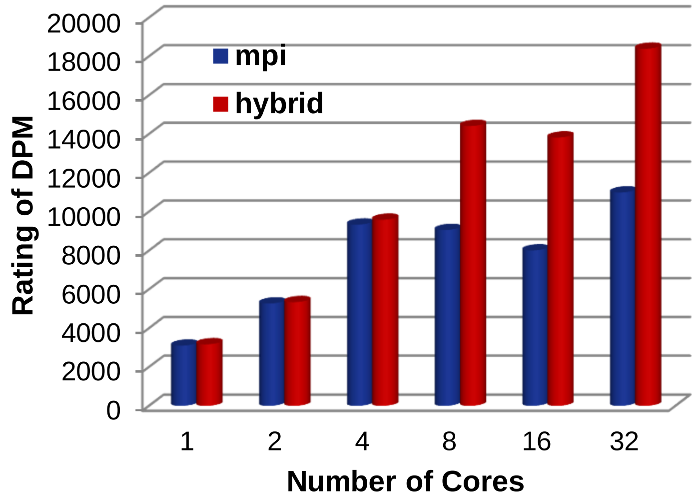

Minimizing Network Traffic

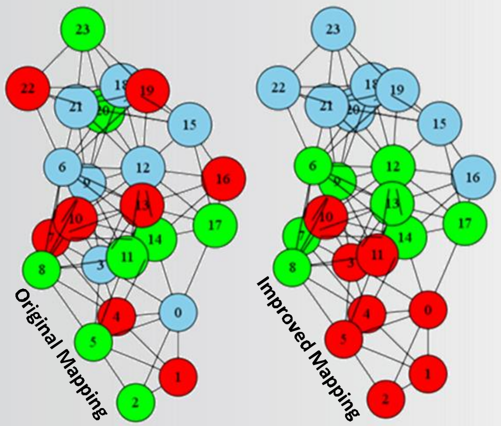

Improved Scaling with R14 Monitors

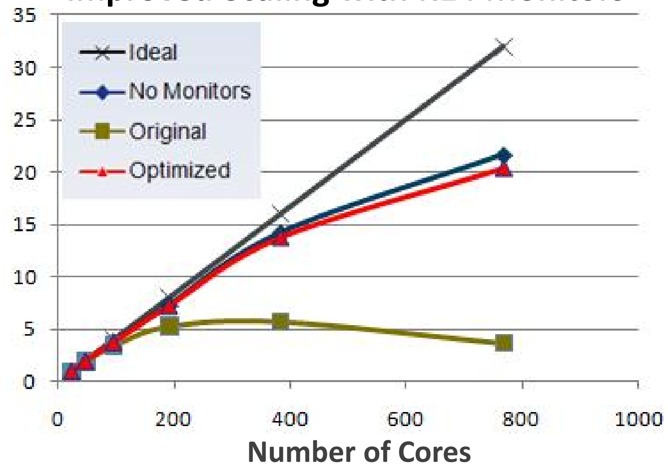

Optimized Solver Performance

- ANSYS HFSS

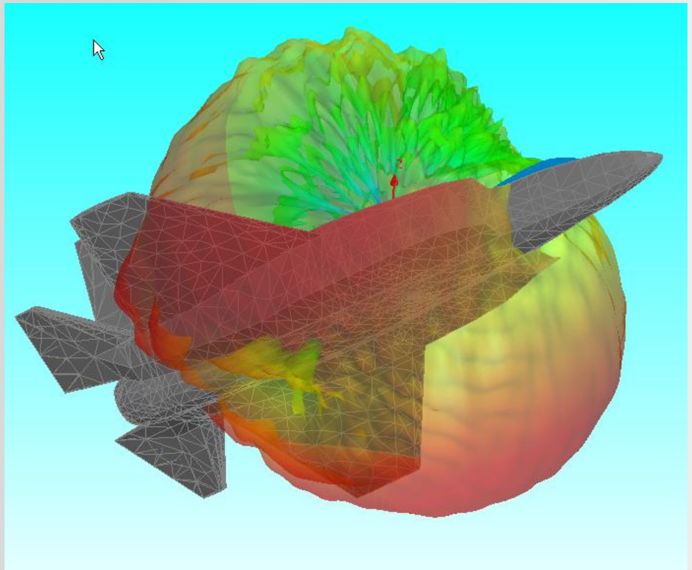

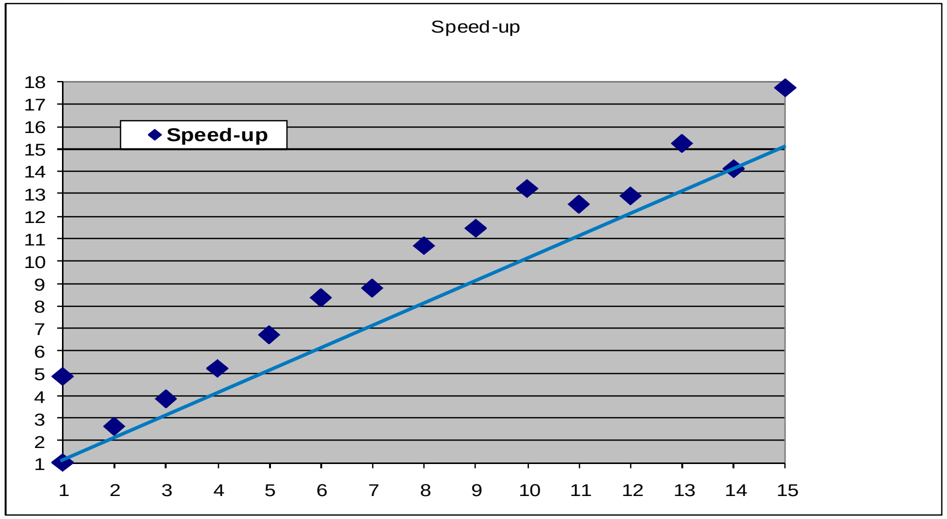

Domain Decomposition Method (DDM), incl. support of finite antenna arrays (R14)

Increased memory efficiency for large and very large problems allows super-scaling!

Faster solutions across multiple processors

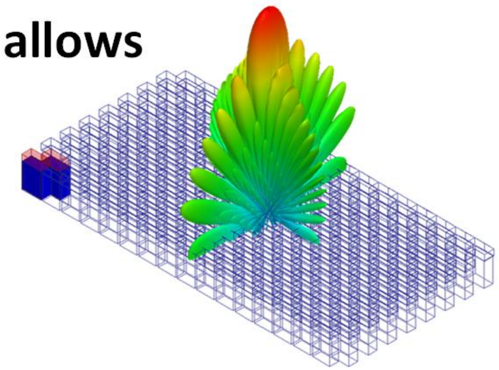

ANSYS

Realize Your Product Promise™

More Simulations

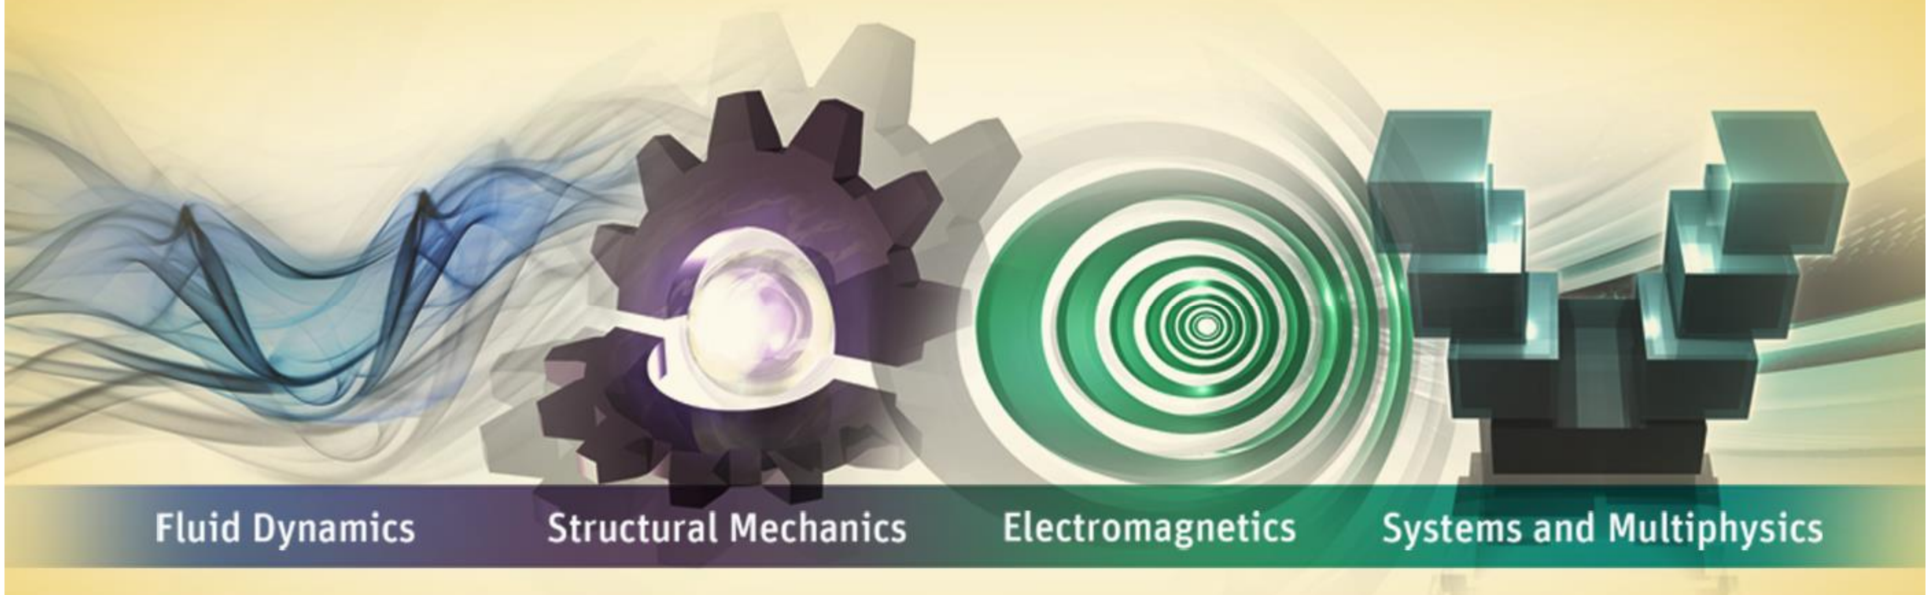

Introduction

- Levels of Parametric Simulation

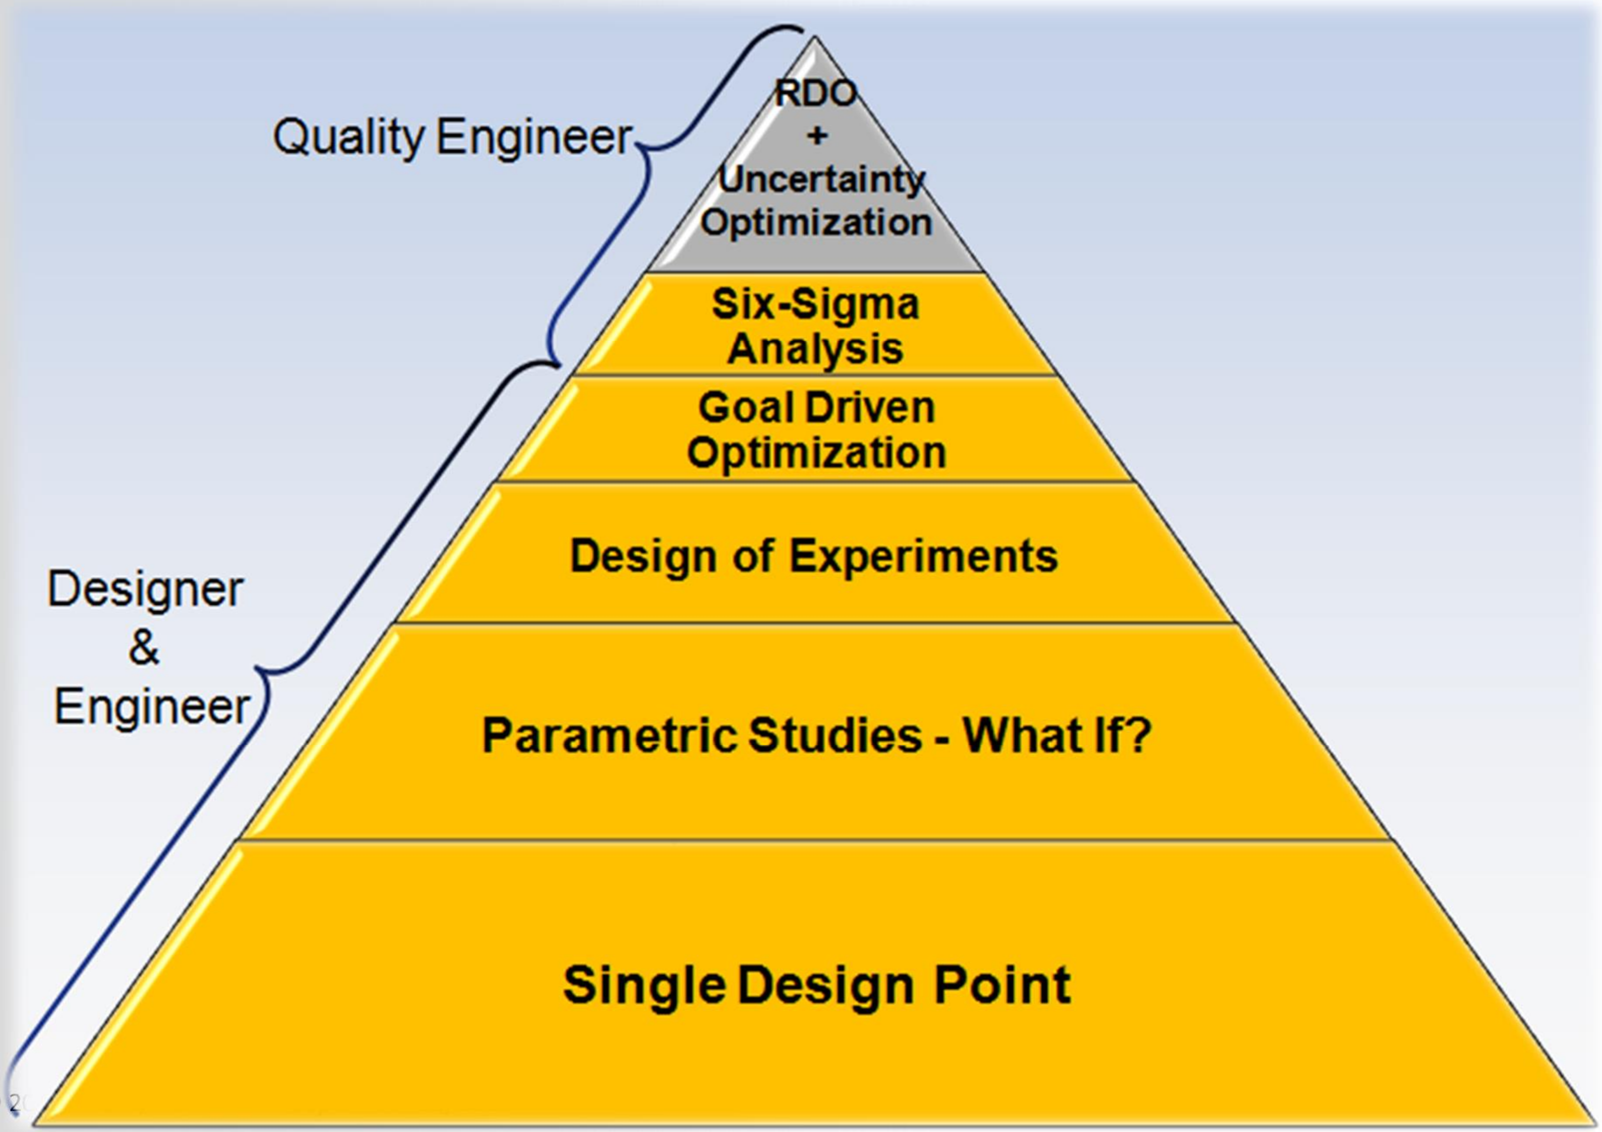

Introduction

- Evolution of Parametric Simulation

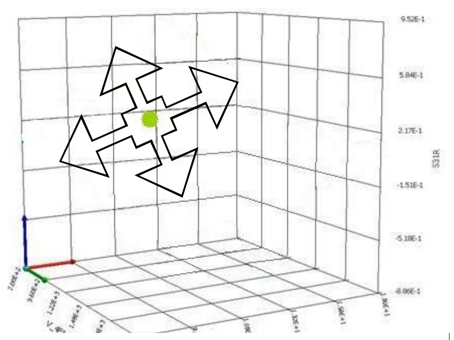

Doesn't provide direction for design improvement

Single Design Point

- Solves a single simulation involving single or multiple physics
- Users are interested in solution robustness, speed, accuracy, ease of use and engineering results
- And the ease and power of the physics coupling

Is this the best design? How can I improve performance? Can I reduce weight or cost? What is limiting performance? Is this a robust design?

Introduction
- Evolution of Parametric Simulation

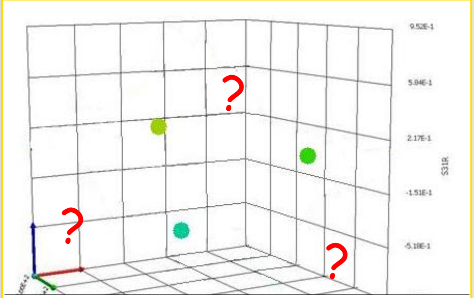

Solutions compared, but design is not well understood and no optimum is found

"What If" Study

- User adjusts inputs and investigates results
- Builds on previous expectations, adds requirement of easy and robust parametric updates and comparative reports

Need a more scientific and automated way to decide which points to solve

Need a way to interpolate between these points

Introduction
- Evolution of Parametric Simulation

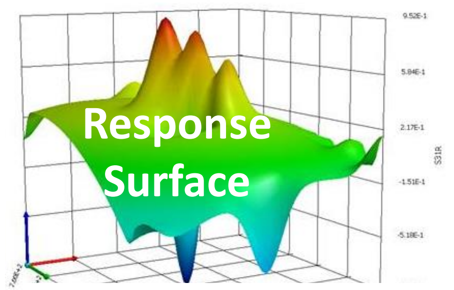

Provides design understanding, but optimization is not automated

Design Exploration

- Scientific methods to explore the design space fully
- Amplifies the importance of the previous technology
- Adds requirements for: robust efficient & affordable distributed solve, sensitivity and correlation, DOE and response surface technology, mesh morphing, charting and reporting

Difficult to optimize a design with many inputs and goals

Introduction
- Evolution of Parametric Simulation

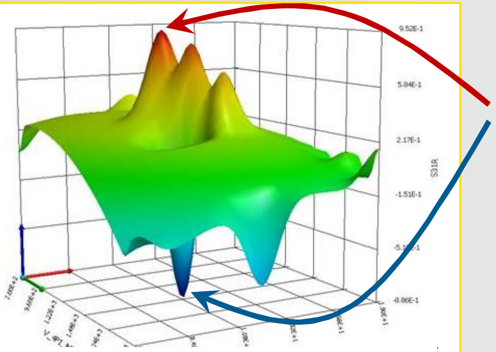

Solutions may be too sensitive to input variability

Optimization

- Searches the design space for optimal candidates, given user-defined goals and priorities
- Amplifies the importance of the previous technology
- Adds requirements for: advanced optimization algorithms to efficiently search for candidates, comparative reporting

Real-world inputs typically have some variation and may require a more "robust design" goal

Introduction

- Evolution of Parametric Simulation

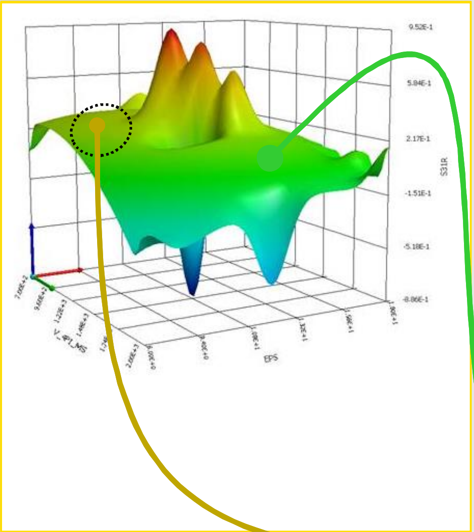

Robust Design

Å Taking the variation of inputs into account, and seeking a design with a probabilistic goal

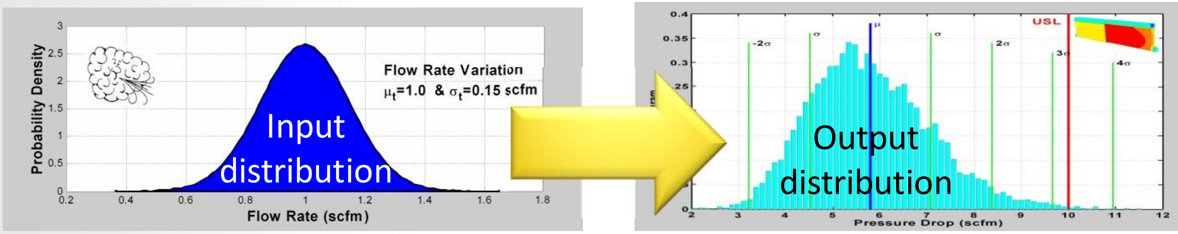

- Å RDO => Min standard deviation of the results
- Å Six Sigma => Optimal design within a safe domain
- Å There are other Robust Design methods/goals…

Å Amplifies the importance of the underlying Workbench and solver technologies

Å Adds requirements for: probabilistic parameters, specific probabilistic optimization algorithms

Our Solution

- From Single Physics to Robust Design

Robust Design is an ANSYS Advantage

Increasing understanding, innovation, ROI

Single Physics Solution

- Accuracy, robustness, speed…

Multiphysics Solution

- Integration Platform

"What if" Study

- Parametric Platform
- Simultaneous Solve

Design Exploration

- DOE, Response Surfaces, Correlation, Sensitivity, Unified reporting, etc.

Optimization

- Algorithms
- Published API

Robust Design

- Six Sigma Analysis
- Probabilistic Algorithms
- Adjoint solver methods

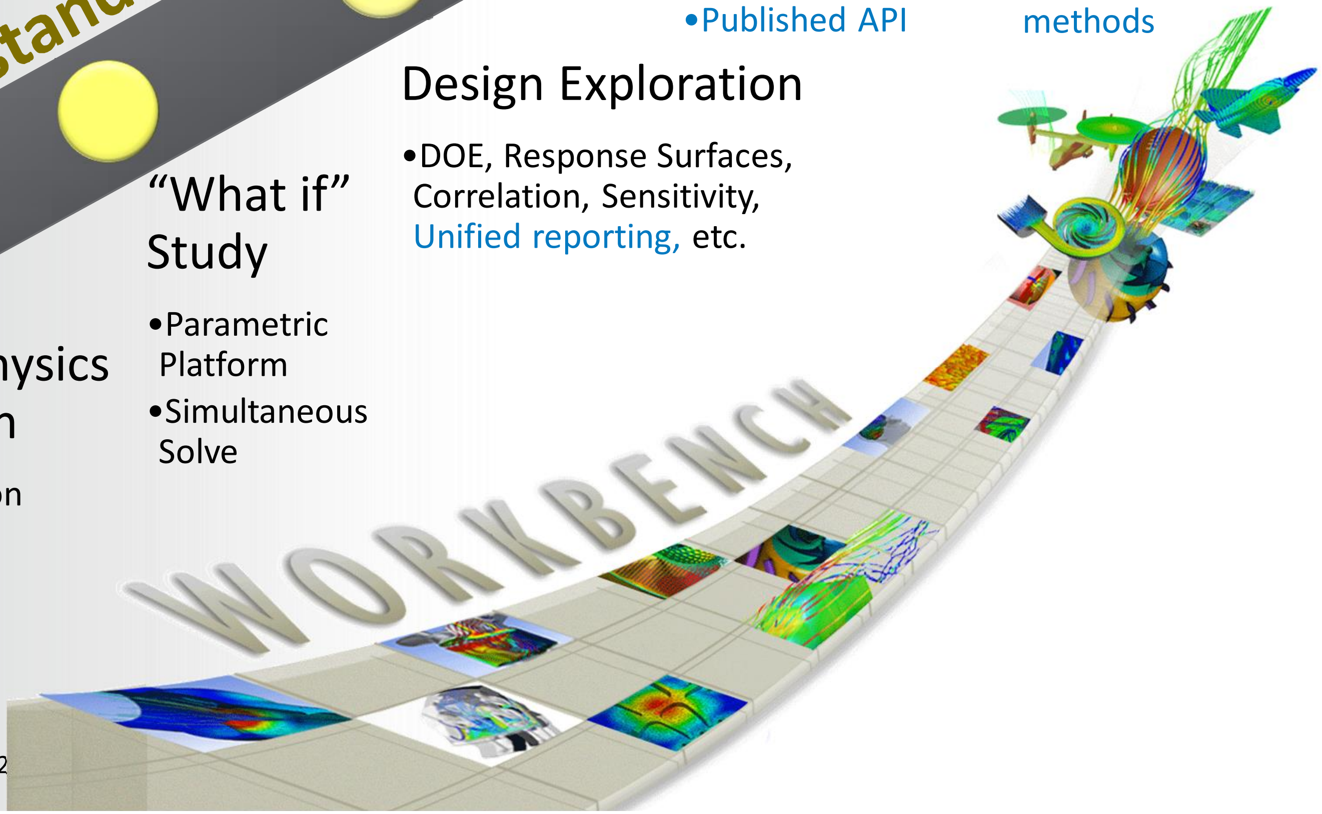

ANSYS Workbench Enables…
Sequential Design Point Update

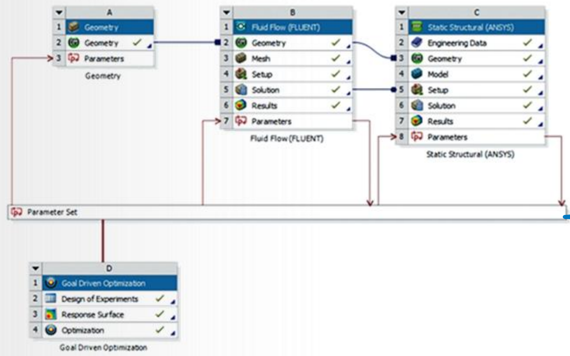

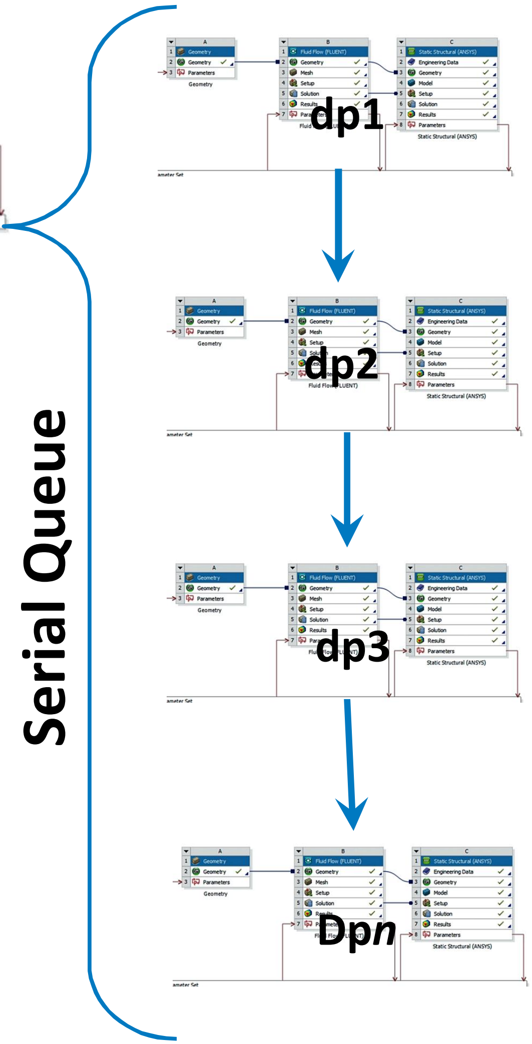

- **Until ANSYS 14.0, design points had to be solved sequentially**
- **That is run dp0 through to dp *n*.**
- **With potentially hundreds of long-running design points, this can be *time* prohibitive.**

ANSYS

ANSYS Workbench Enables…
Simultaneous Design Point Update

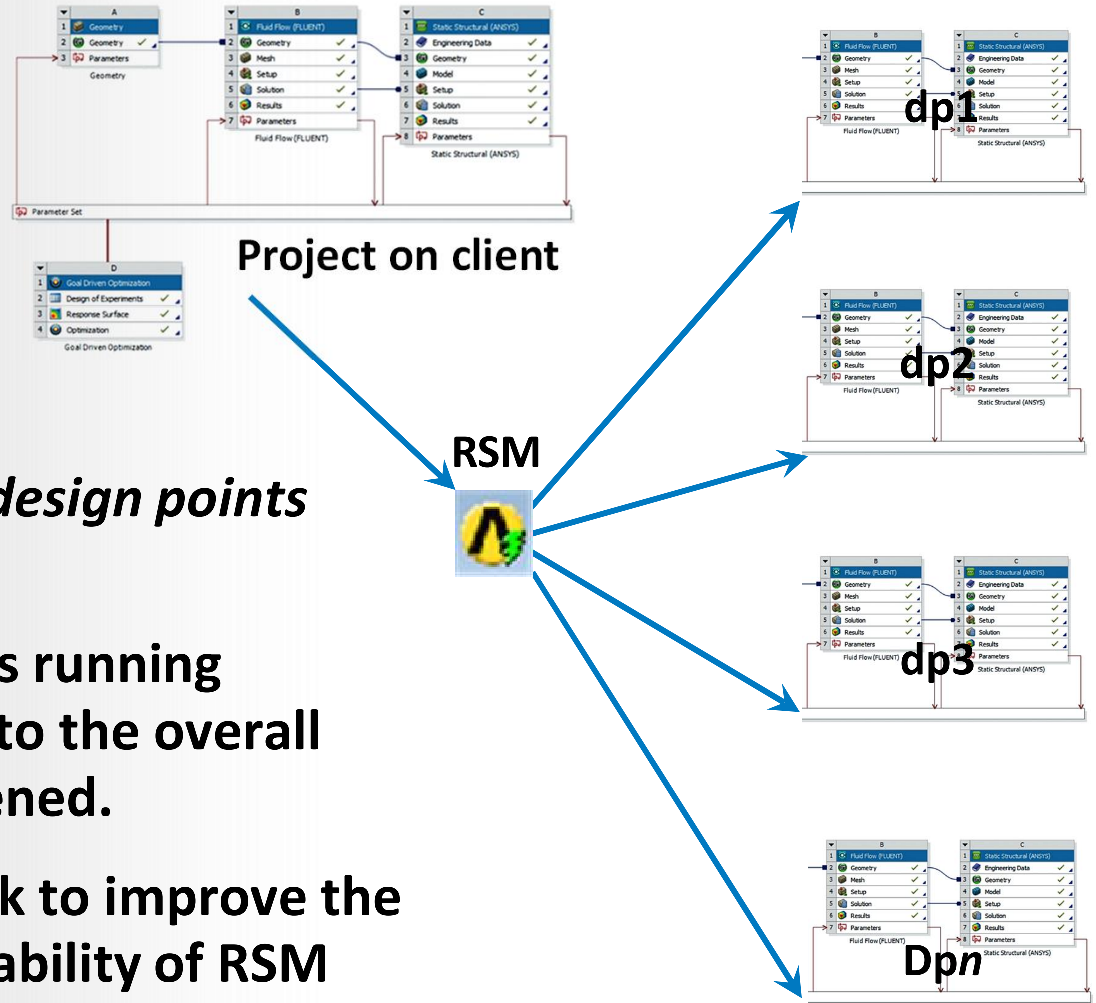

- **R14.0 supports updating *design points* simultaneously via RSM.**
- **With several design points running simultaneously, the time to the overall result can be greatly lessened.**
- **14.5 included a lot of work to improve the robustness, speed and usability of RSM**

ANSYS Workbench Enables…
RSM with 3rd Party Scheduler

RSM has two modes:

Å It can be used as a scheduler for local jobs, or

Å It can be used as a mechanism to access 3rd party schedulers for more advanced distributed solves…

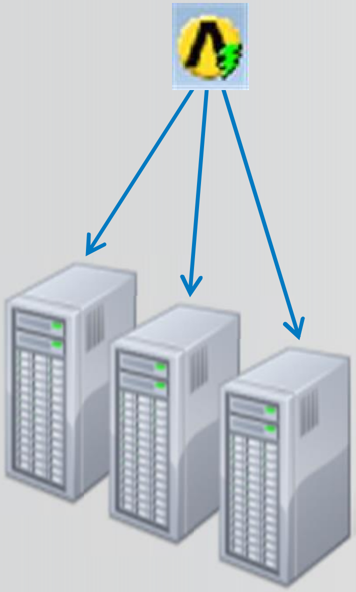

RSM as a scheduler
(Unit: Jobs)

You setup the compute servers and how many jobs run on each, the queues and which have priority

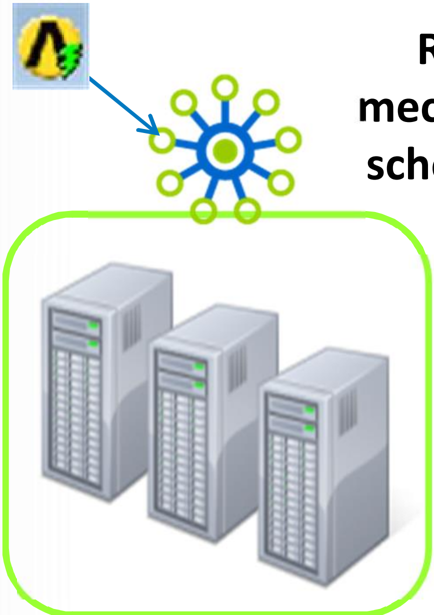

RSM as a transport mechanism to a 3rd party scheduler such as LSF or PBS (Unit: Cores)

Third party tools break up the jobs and can distribute them across a network

ANSYS Workbench Enables…
Simultaneous Design Point Update

License Usage

- ANSYS products “grab” licenses as each software component is executed
- To update n design points simultaneously you need *n* * the licenses.
- This makes running simultaneous design points *cost* prohibitive.

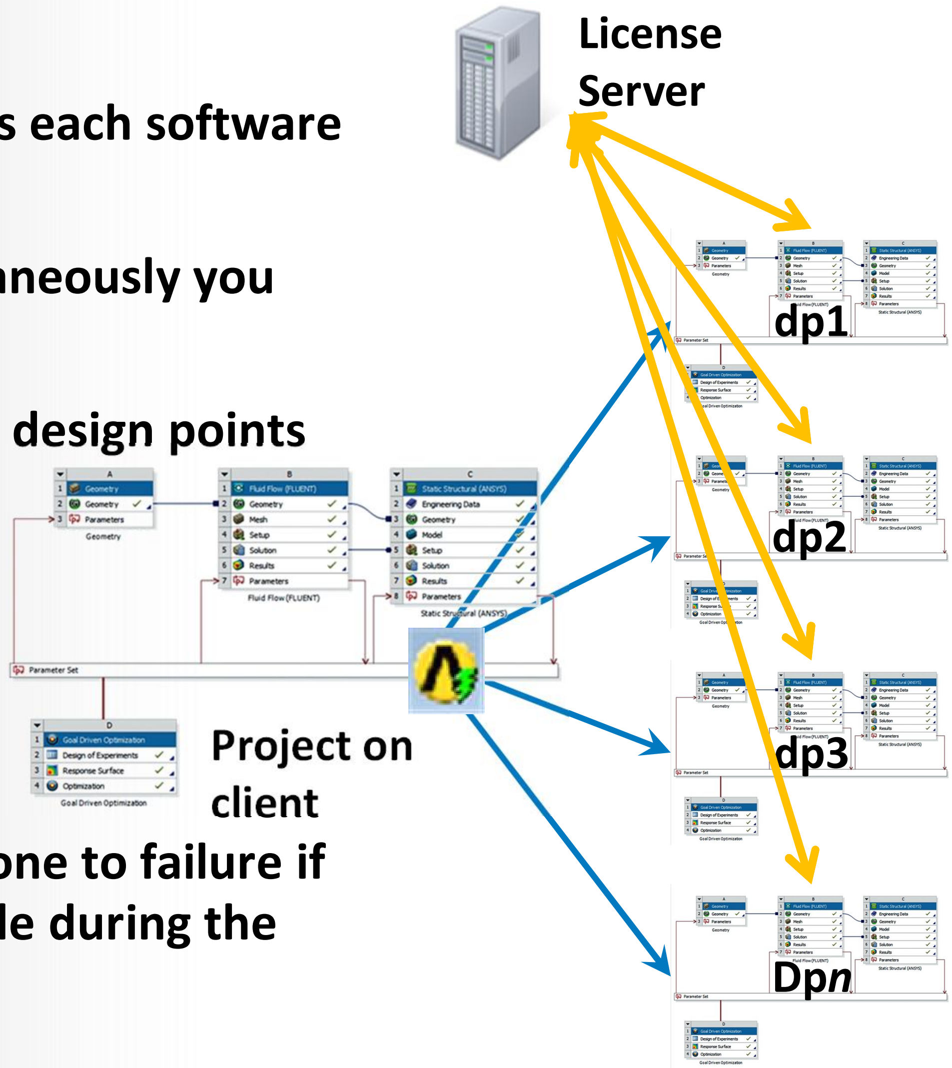

- It can also make design points prone to failure if not enough licenses were available during the update process.

ANSYS Workbench Enables…
ANSYS HPC Parametric Pack Licensing at R14.5

Scalable, like ANSYS HPC Packs

- Enhances the customer's ability to include many design points as part of a single study

Amplifies complete workflow

- Allow users to run n design points simultaneously, multiplying the "base" license(s)
- Design points can include execution of multiple products (pre, meshing, solve, HPC, post)

Requirements

- Parameters need to be in ANSYS Workbench
- Sequential execution of geometry updates

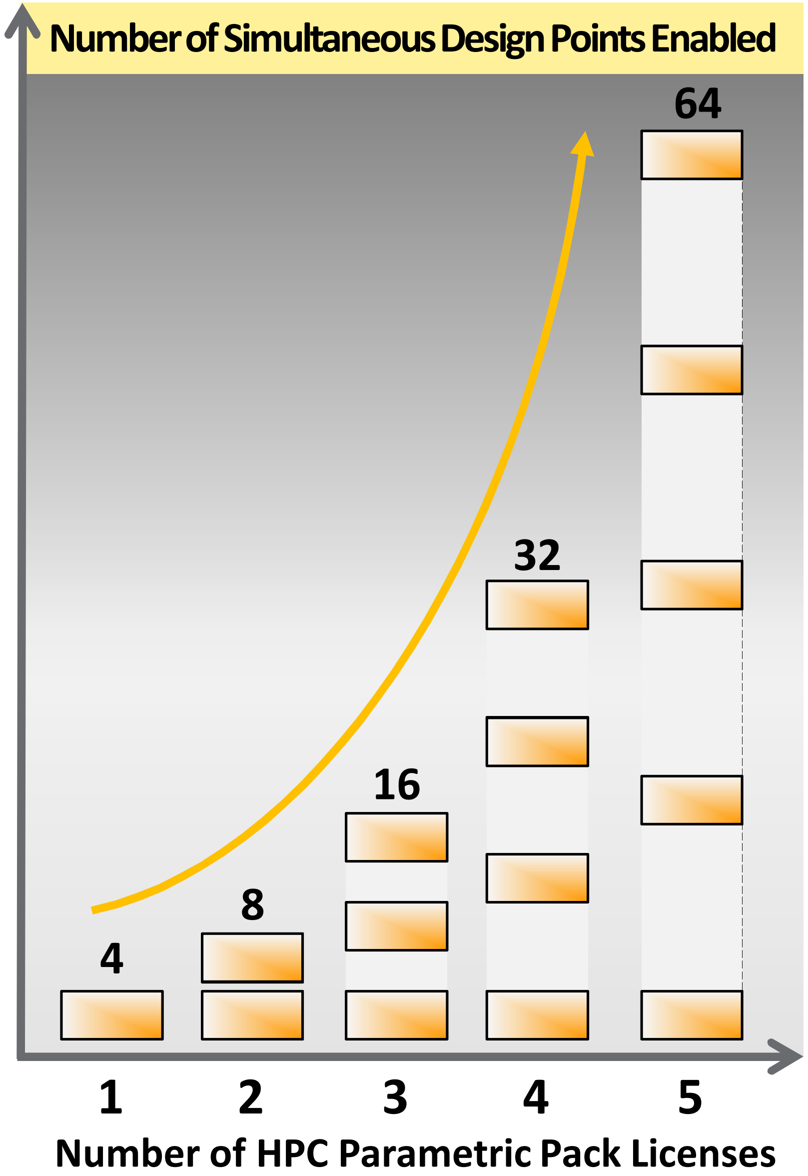

ANSYS Workbench Enables…
"Game Changing" Time to Design Insight

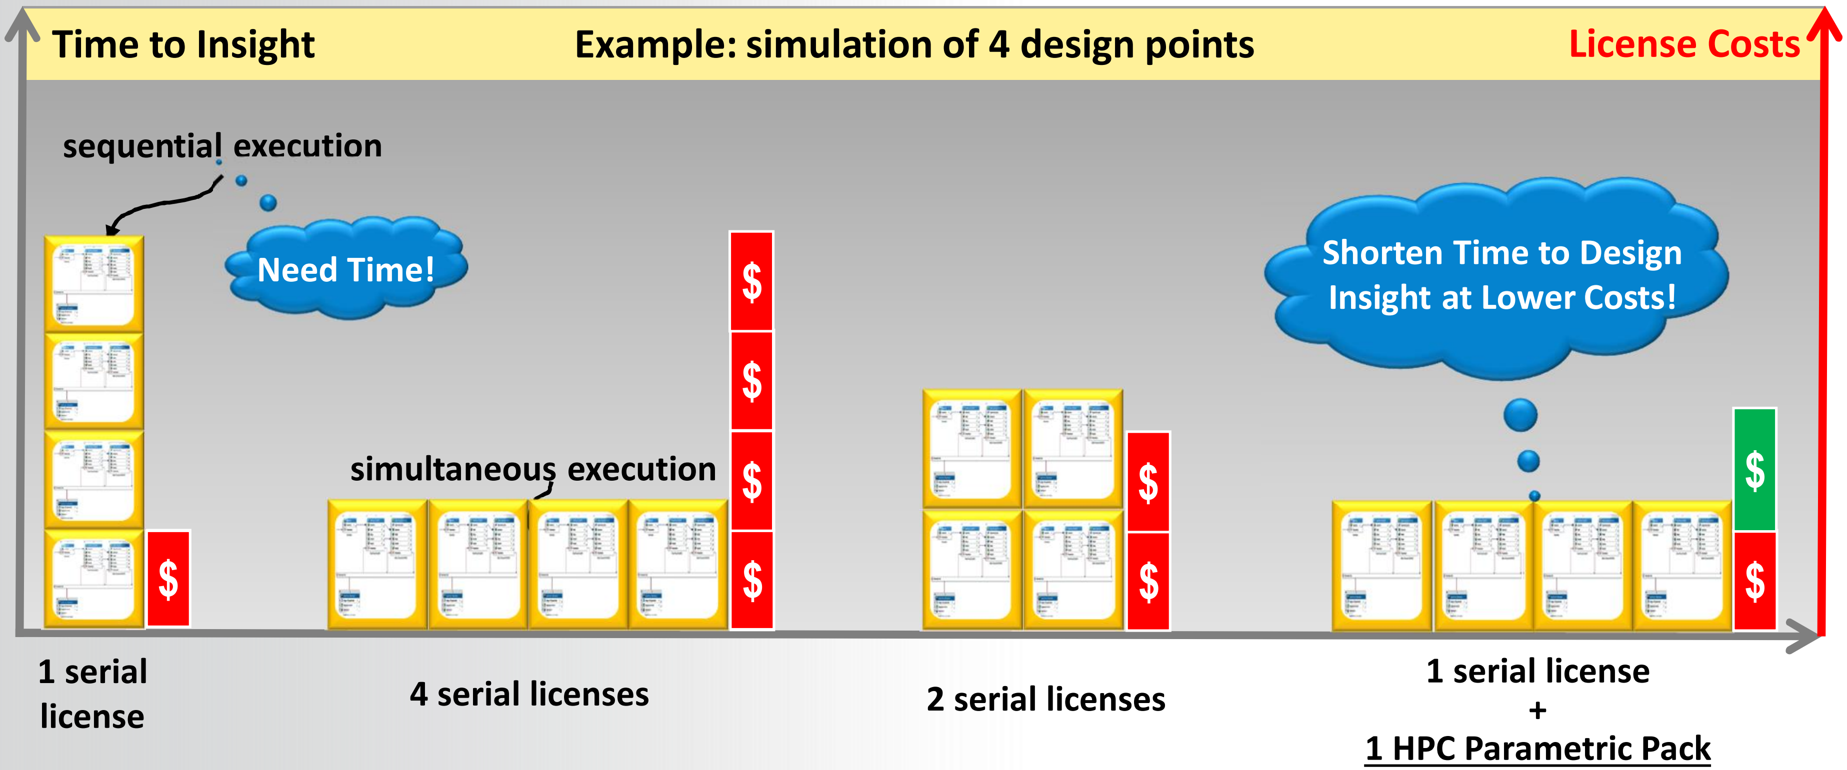

ANSYS Workbench Enables…
"Game Changing" Time to Design Insight

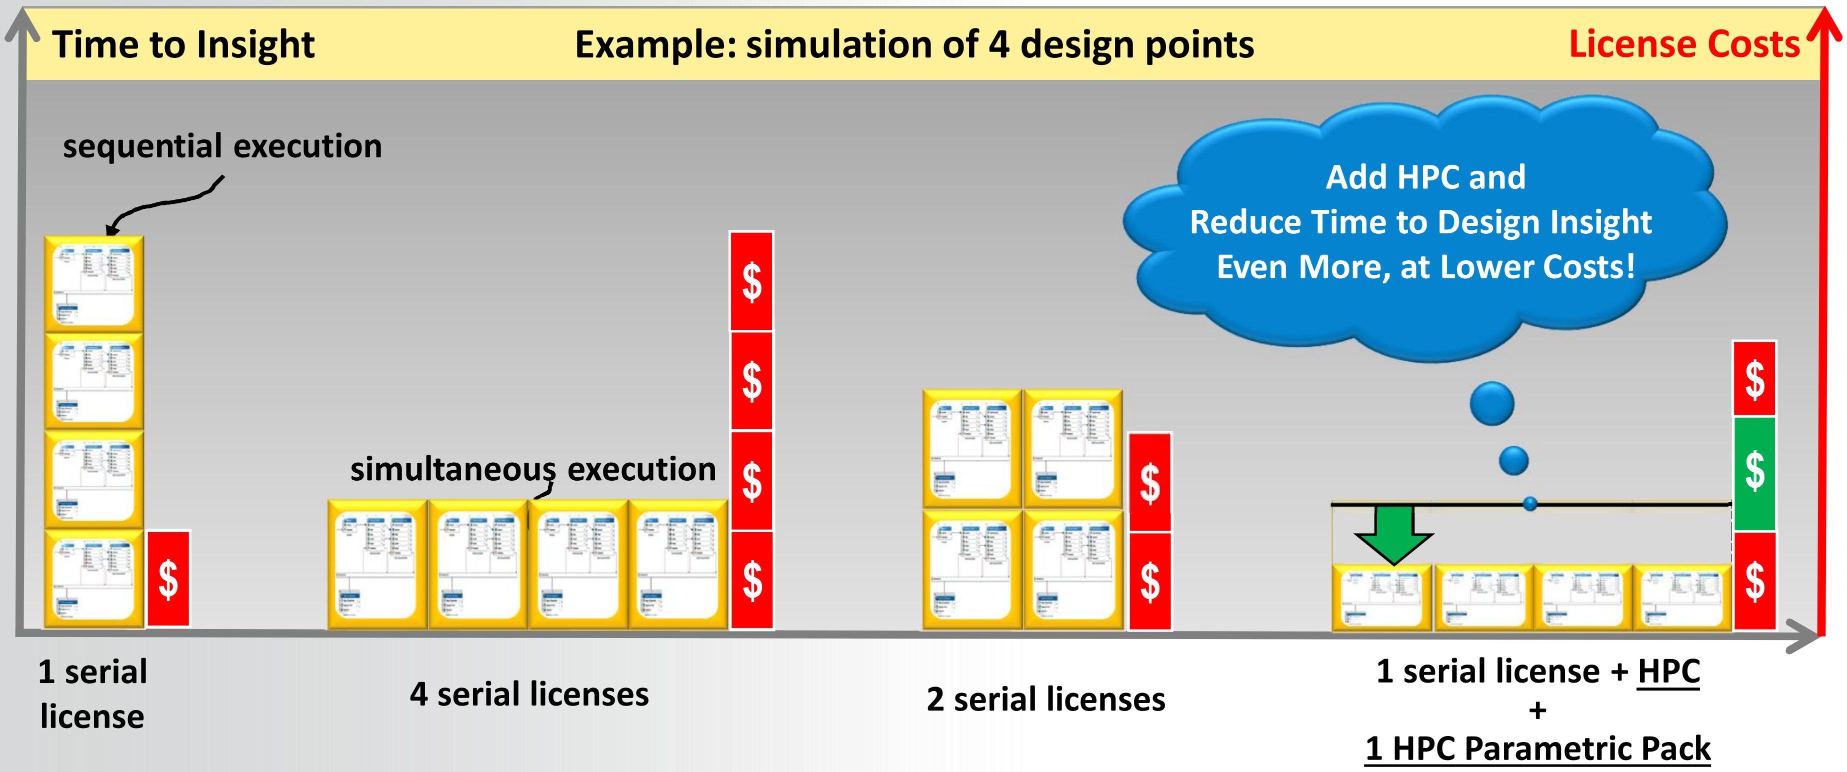

ANSYS Optimization Partners

ANSYS simulation software has been effectively used in concert with many optimization partners

Å MATLAB (Mathworks)

Å modeFRONTIER (Esteco)

Å optiSLang (Dynardo)

Å eArtius

Å Optimus (Noesis)

Å RBF-Morph

Å Sculptor (Optimal)

Å Sigma Technology (IOSO)

Å TOSCA (FE-DESIGN)

Å iSight (Dassault)

Å Qfin (Qfinsoft)

Å and more…

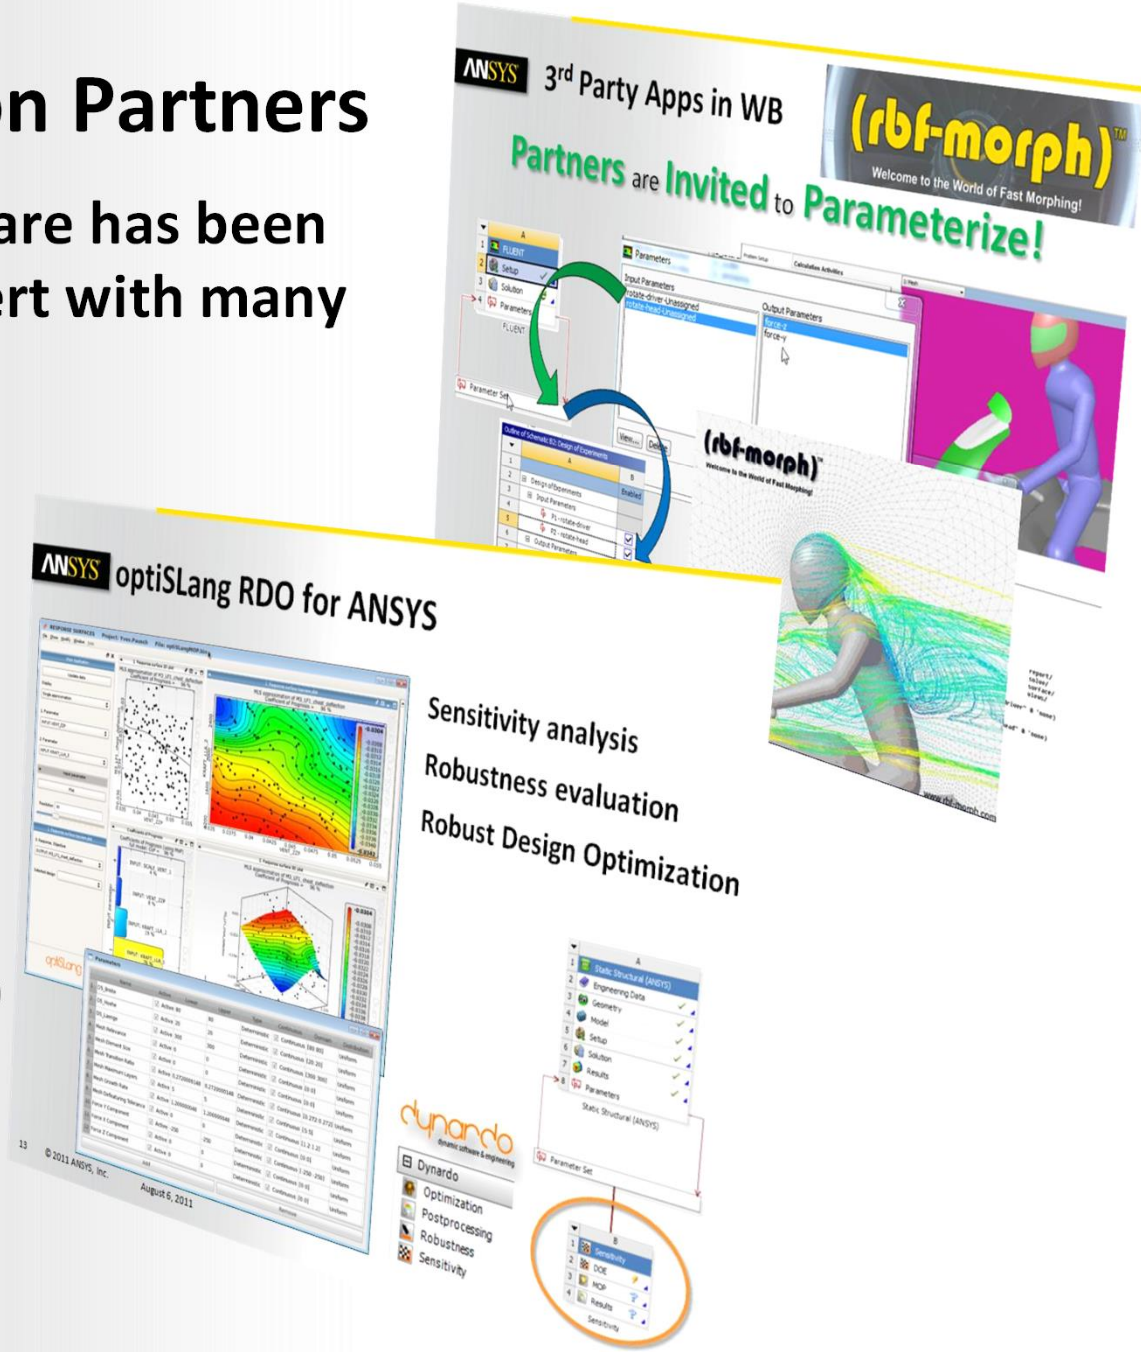

ANSYS

Realize Your Product Promise™

HPC Parametric Pack Example Applications

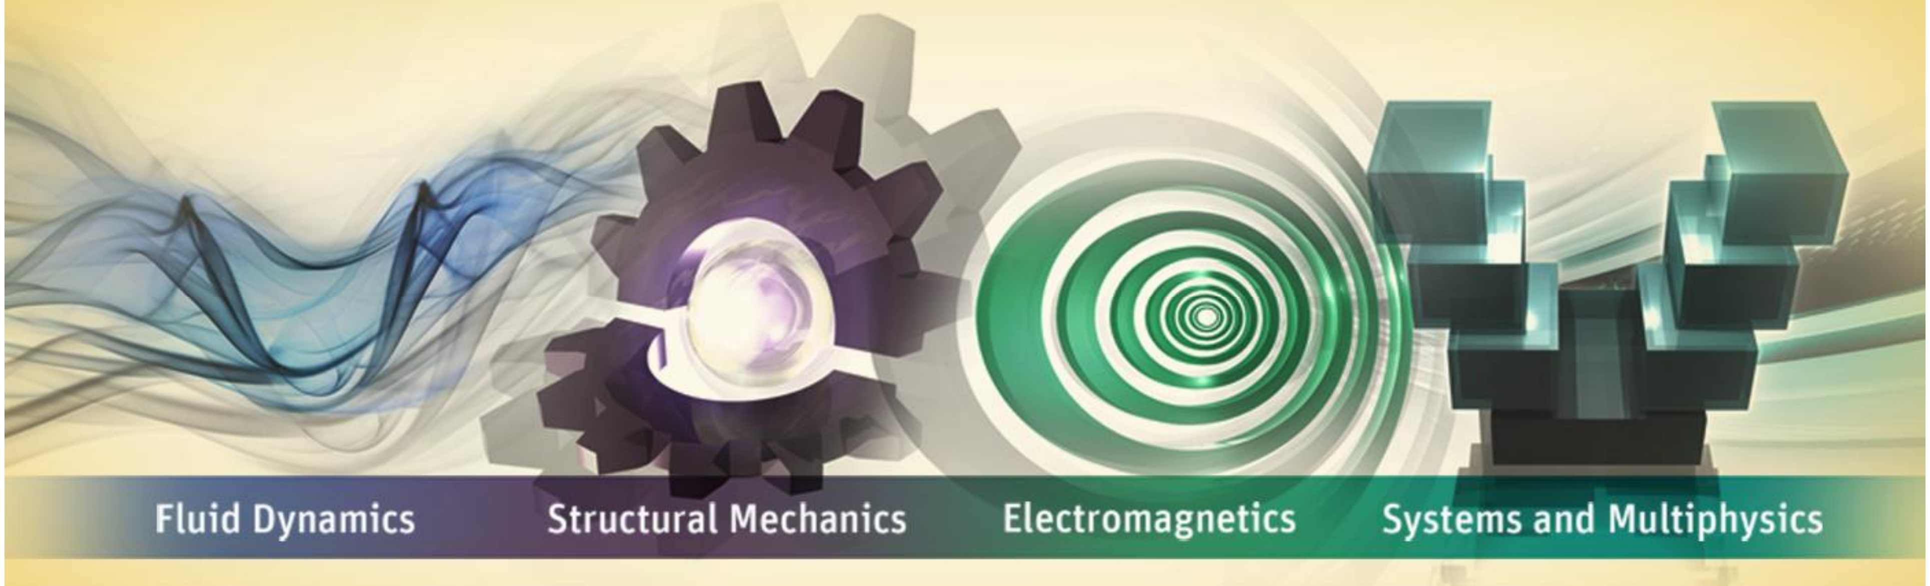

Rear Axle Model
- Evaluating Material Properties

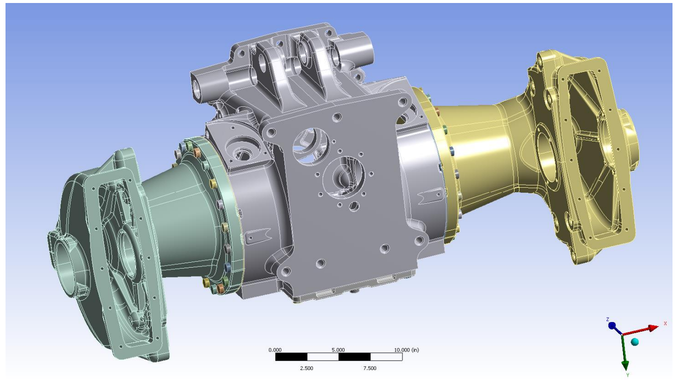

Problem Description

Å Large deflection non-linear static model investigating design sensitivity to material properties

Å Input parameter: material property (8 design points)

Å Detail:

- Sparse matrix solver running incore; 4 load steps
- 1,393,811 nodes, 829,701 elements (4,151,766 DOF)
- Hardware: Dell workstation with dual Intel Xeon E5-2690 (2.90 GHz, 16 cores), 256 GB memory, all jobs running 2 cores

Licensing Solution

Å 1 ANSYS Mechanical

Å 2 ANSYS HPC Parametric Packs

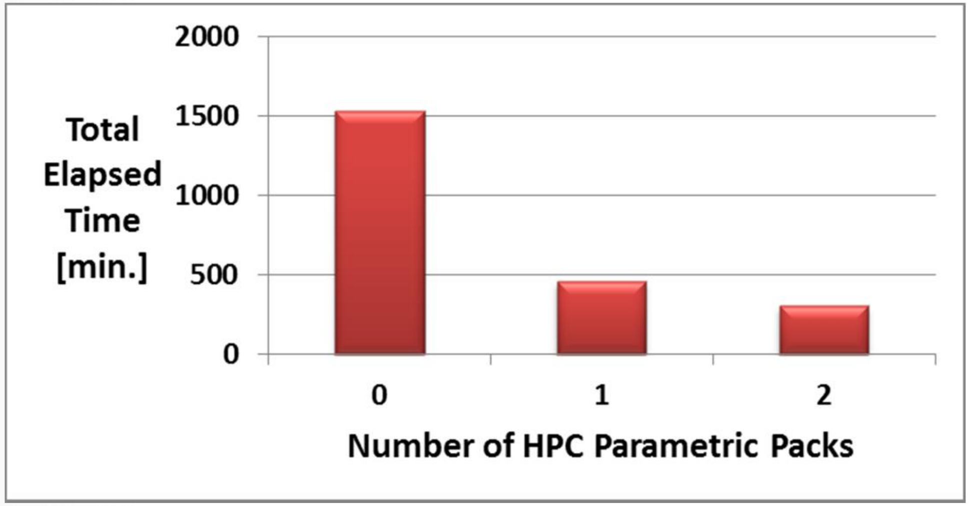

Result/Benefit

Å 5x speedup over sequential execution

Å Easier and fully automated workflow!

Acknowledgment: Paul Schofield and Jiaping Zhang, ANSYS Houston

Static Analysis of Semi-submersible
- Evaluating Shell Thicknesses

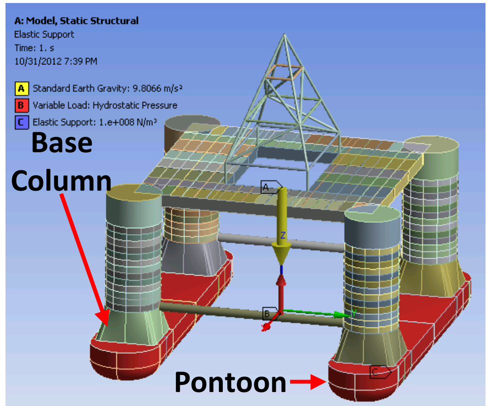

Problem

- Static Analysis of semi-submersible using beam & shell elements, subjected to hydrostatic pressure and gravity loading
- Design objective: minimize both total mass and equivalent stress
- Input parameters: pontoon thickness, base column thickness (16 design points)
- Detail:
  - 232,583 nodes, 230,770 elements
  - Hardware: Dell workstation with dual Intel® Xeon® E5-2690 (2.90 GHz, 16 cores), 256 GB memory, all jobs running 2 cores

Licensing Solution

- 1 ANSYS Mechanical
- 2 ANSYS HPC Parametric Packs

Result/Benefit

- ~6x speedup over sequential execution
- Easier and fully automated workflow

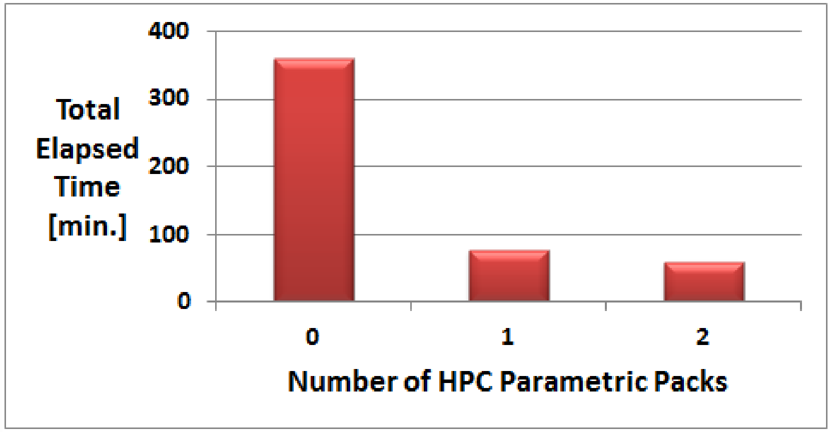

Acknowledgment: Paul Schofield and Jiaping Zhang, ANSYS Houston

Fatigue Analysis of Shaft
- Evaluating Geometries

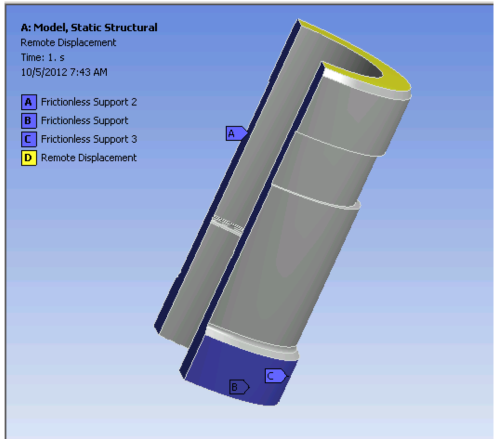

Problem

- Fatigue Analysis of steel shaft subjected to shear cyclic loading on top surface while being fixed on the bottom end
- Input parameters: base height, base thickness, groove height (15 design points)
- Detail:
  - Strain-life fatigue analysis of shaft subject to cyclic loading on the top surface
  - 364,959 nodes, 82,863 elements
  - Hardware: Dell workstation with dual Intel Xeon E5-2690 (2.90 GHz, 16 cores), 256 GB memory, all jobs running 2 cores

Licensing Solution

- 1 ANSYS Mechanical, 1 Fatigue Module, 1 ANSYS Design Modeler
- 2 ANSYS HPC Parametric Packs

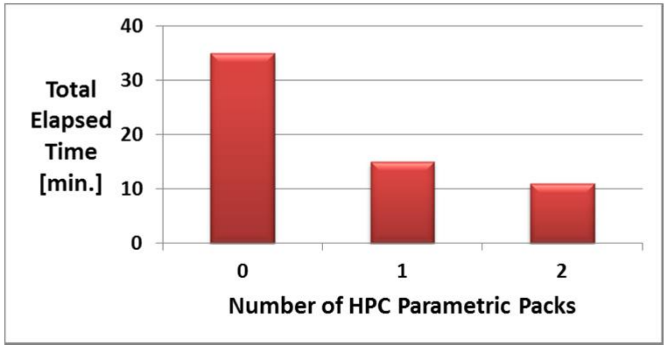

Result/Benefit

- 3.2x speedup over sequential execution
- Easier and fully automated workflow

Acknowledgment: Paul Schofield and Jiaping Zhang, ANSYS Houston

Response Spectrum of Pressure Vessel
- Evaluating Geometries

Problem

Å Pressure Vessel subjected to high internal pressure and subjected to acceleration in supports during earthquake

Å Input parameters: vessel thickness, vessel radius, vessel Height (16 design points)

Å Detail:

- "Static Structural" + "Modal Analysis" + "Response Spectrum"
- 62,439 nodes, 150,169 elements
- Hardware: Dell workstation with dual Intel Xeon E5-2690 (2.90 GHz, 16 cores), 256 GB memory, all jobs running 2 cores

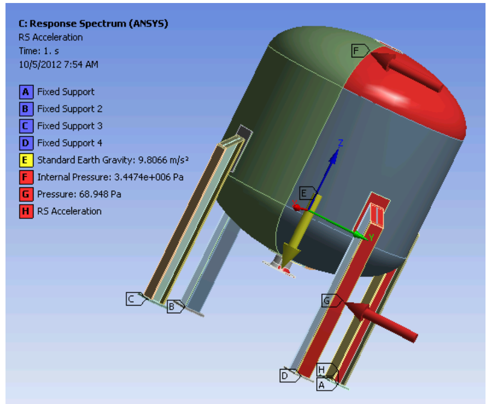

Licensing Solution

Å 1 ANSYS Mechanical, 1 ANSYS DesignModeler

Å 2 ANSYS HPC Parametric Packs

Result/Benefit

Å ~3x speedup over sequential execution

Å Easier and fully automated workflow

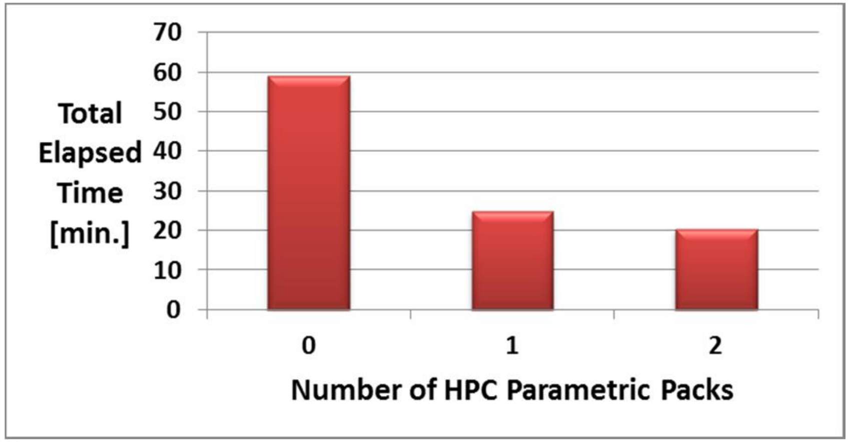

Acknowledgment: Paul Schofield and Jiaping Zhang, ANSYS Houston

Intake Manifold Fluid Analysis
- Evaluating Geometries

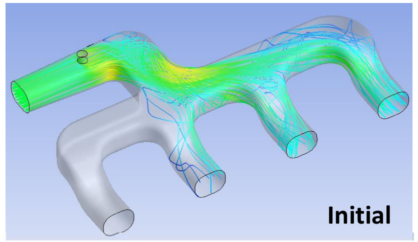

Problem Description

Å Non-homogenous air flow in intake manifold through the 4 outlets

Å Design objectives:

- Equal fresh and exhaust gas mass flow distribution to each cylinder
- To minimize the overall pressure drop

Å Input Parameters: radii of 3 fillets near inlet (16 design points)

Å Detail:

- Steady state pressure based solver, realizable k-epsilon model
- 57,790 nodes, 208,740 elements
- Hardware: Dell workstation with dual Intel Xeon E5-2690 (2.90 GHz, 16 cores), 256 GB memory

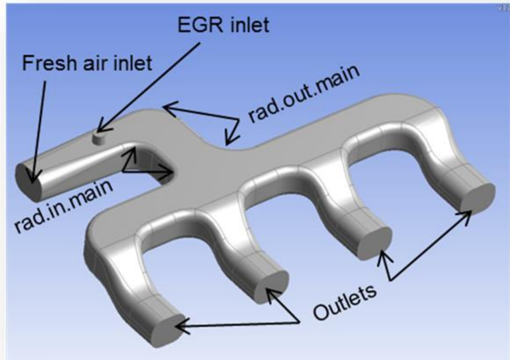

Licensing Solution

Å 1 ANSYS CFX, 1 ANSYS DesignModeler

Å 2 ANSYS HPC Parametric Packs

Result/Benefit

Å ~2.2x speedup over sequential execution

Å Easier and fully automated workflow

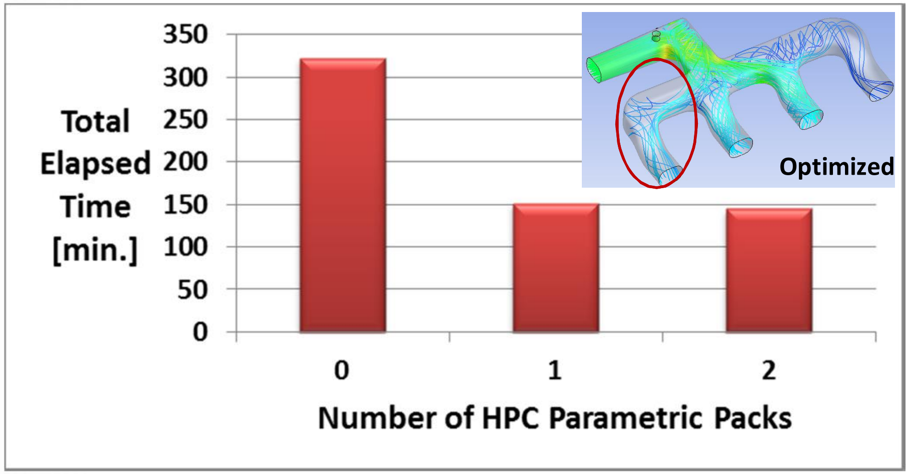

Acknowledgment: Paul Schofield and Jiaping Zhang, ANSYS Houston

ANSYS

Realize Your Product Promise™

Other
Example Applications

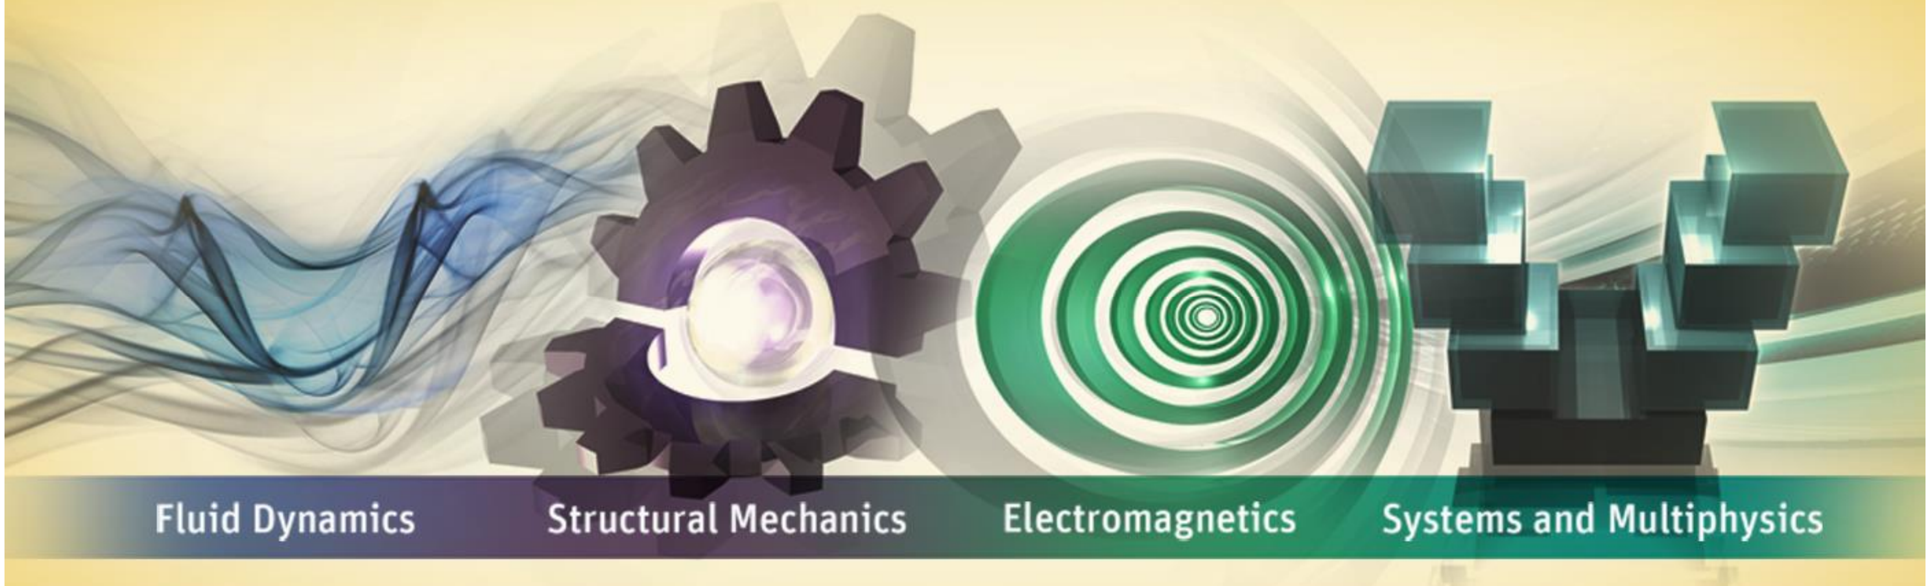

Turbine Blade Root

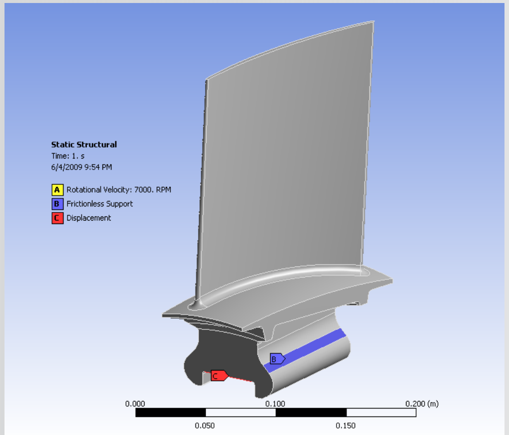

Input Parameters

Values			
Lower Bound	1.5		
Upper Bound	1.65	ds_xtilt	
Initial Value	1.5		
Values			
Lower Bound	0.9		
Upper Bound	1	ds_ytilt	
Initial Value	1		
Values			
Lower Bound	0.25		
Upper Bound	0.3	ds_rootrad	
Initial Value	0.25		

Output Parameter: Minimum Life

Design Objective:

- To determine the optimal parameters for maximum fatigue life of a blade root

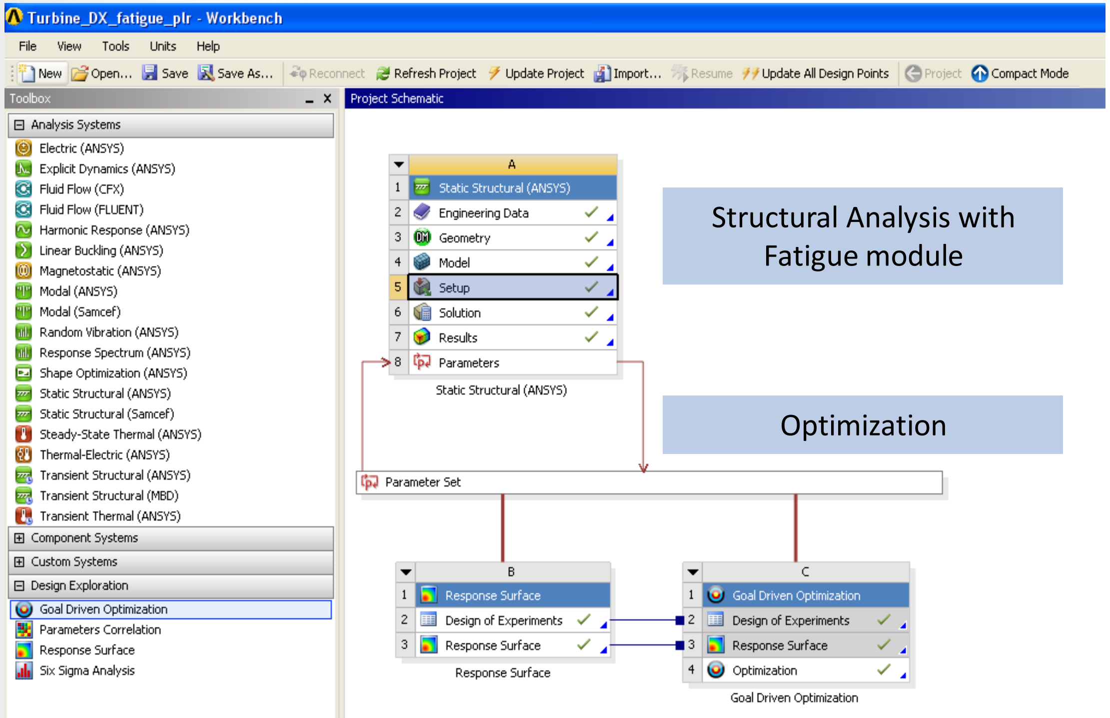

Turbine Blade Root

Response Surface

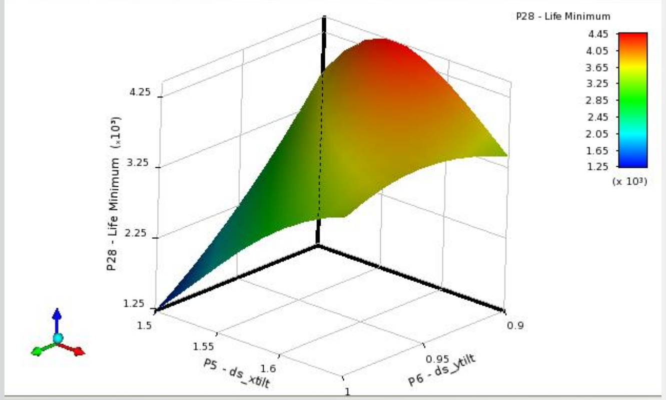

Optimization

Table of Schematic C4: Optimization

	A	B	C	D	E
1		P5 - ds_xtilt	P6 - ds_ytilt	P7 - ds_rootrad	P28 - Life Minimum
2	⊟ Optimization Study				
3	Objective	No Objective	No Objective	No Objective	Maximize
4	Target Value				
5	Importance	Default	Default	Default	Default
6	⊟ GDO Sample Set 1				
7	Candidate A	1.6457	0.98276	0.29984	★★★ 4801.6
8	Candidate B	1.6376	0.96636	0.29861	★★★ 4655.1
9	Candidate C	1.5553	0.90571	0.29721	★★★ 4639.1

Objective is to maximize fatigue life

Design point for best candidate

Initial Design

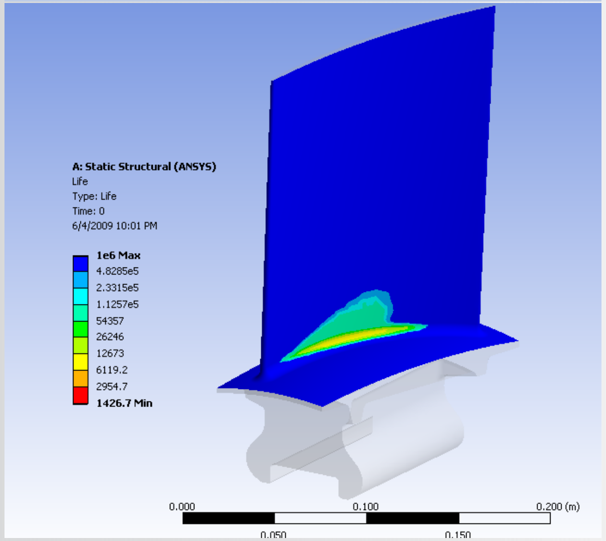

Optimized Design

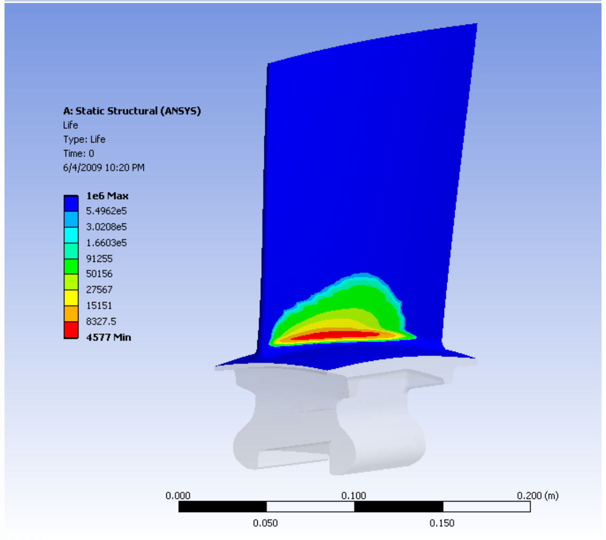

Hip Joint Implant

Design Objective:

Å To optimize the implant for minimum human discomfort

Å Constraint: the relative sliding between bone marrow and implant should be less than 120 µm but greater than 30 µm

Implant Geometry

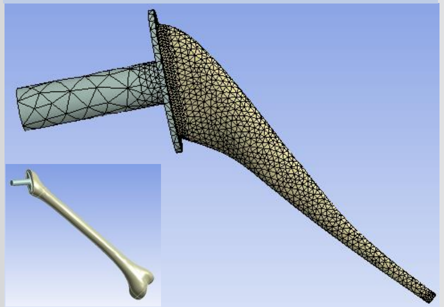

Static Structural Analysis

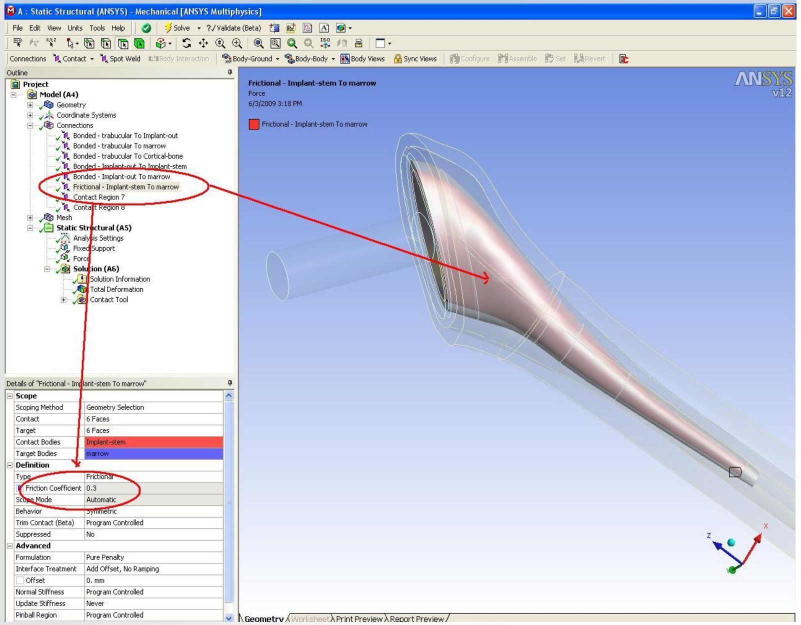

Contours of sliding distance

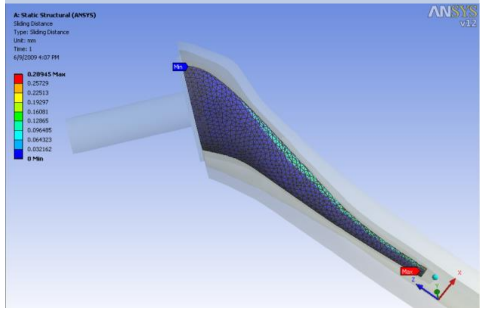

Hip Joint Implant

Response Surface

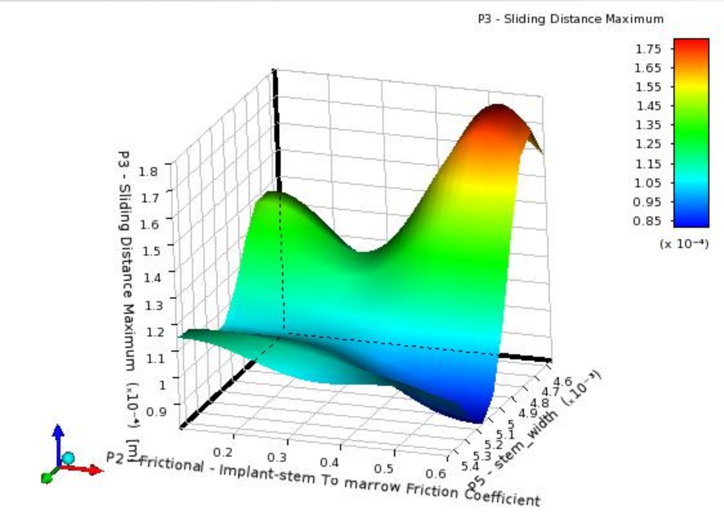

Sensitivity curve indicates sliding distance is more sensitive to steam width

Optimization

Table of Schematic C4: Optimization

	A	B	C	D	E
1		P2 - Frictional - Implant-stem To marrow Friction Coefficient	P5 - stem_width	P7 - Force Z Component (N)	P3 - Sliding Distance Maximum (m)
2	Optimization Study				
3	Objective	No Objective	No Objective	No Objec…	Seek Target
4	Target Value				9E-05
5	Importance	Default	Default	Default	Default
6	GDO Sample Set 1				
7	Candidate A	0.51988	0.0053988	-525.25	9.0006E-05
8	Candidate B	0.21238	0.0050125	-470.75	9.0022E-05
9	Candidate C	0.45288	0.0052309	-618.17	9.0707E-05

Optimized Geometry

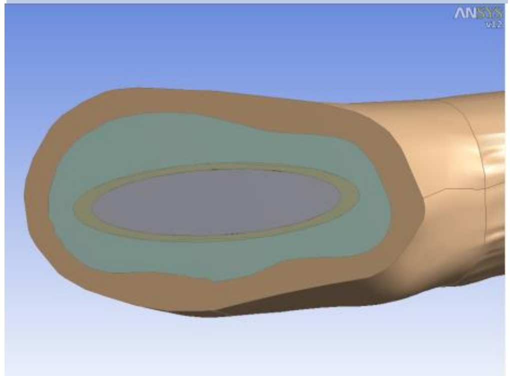

Stem width changed by 18 %

Dyson Air Multiplier™ Fan

Design objective:

- Maximize amplification ratio for a given size and power consumption
- 3 main design parameters, i.e. gap in annular ring, internal profile of ring, profile of external ramp

Customer benefits include:

- Explored 10-fold of design variations than would otherwise have been possible (each day 10 instead of 1)
- Improved performance 250% over original design

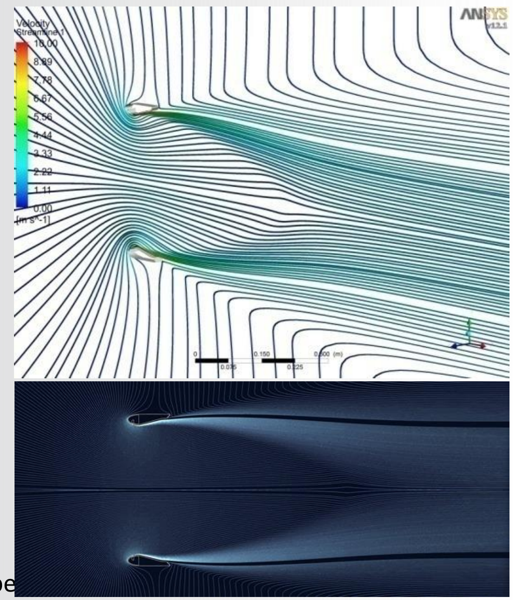

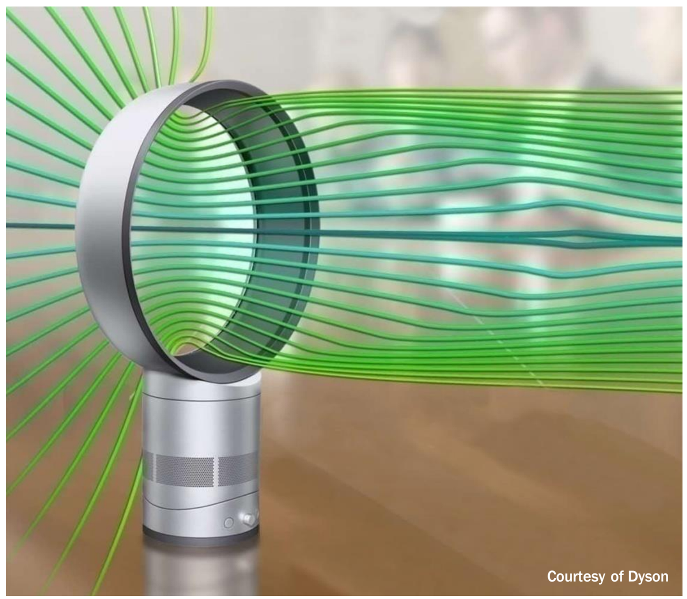

IC Engine Intake Port

Design Objective:

- Maximize Effective Flow Area of a gasoline engine within a specified range of input design parameters

Parametric CAD model created in CATIA

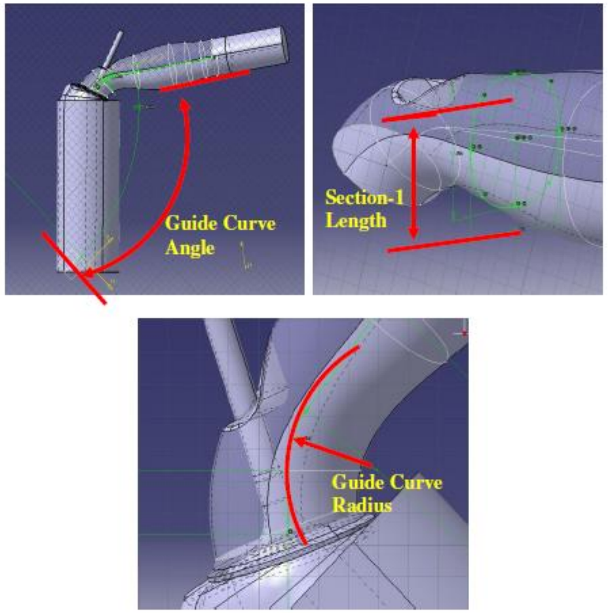

Imported geometry in WB

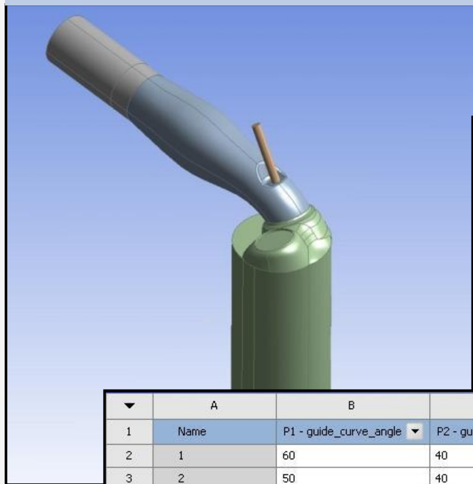

Tetrahedral meshing using AMP:
Mesh Count = 800K
Curvature and Proximity based sizing functions

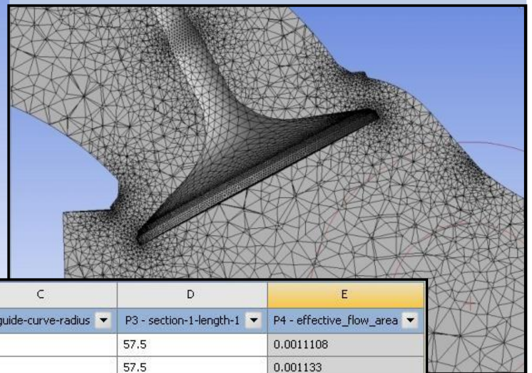

	A	B	C	D	E
1	Name	P1 - guide_curve_angle	P2 - guide-curve-radius	P3 - section-1-length-1	P4 - effective_flow_area
2	1	60	40	57.5	0.0011108
3	2	50	40	57.5	0.001133
4	3	70	40	57.5	0.0011038
5	4	60	30	57.5	0.0011562
6	5	60	50	57.5	0.0011055
7	6	60	40	50	0.0011002
8	8	52	32	51	0.0011309
9	9	68	32	51	0.0011115
10	10	52	48	51	0.0011141
11	11	68	48	51	0.0011136
12	12	52	32	64	0.0011578
13	13	68	32	63	0.0011322

Custom DOE generated with 13 design points for 3 input parameters

IC Engine Intake Port

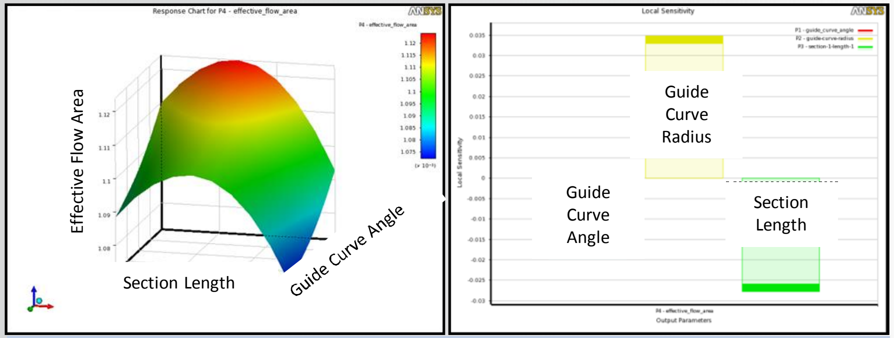

Response Surface and Sensitivity Chart

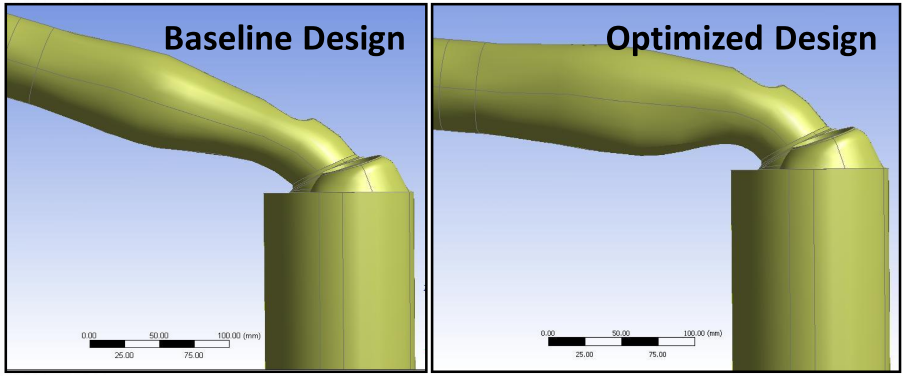

	Guide Curve Angle (Deg)	Guide Curve Radius (mm)	Section-1-Length (mm)	EFA (mm^2)
Baseline	**63**	**41**	**51**	**1100.2**
Optimized	**50**	**30**	**60.5**	**1180.4**

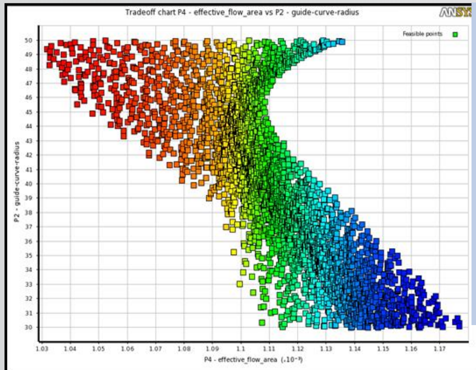

Statistical Analysis using 10000 points:
(A) Trade-off plot
(B) Multiple Goal Driven Optimization

Table of Schematic B4: Optimization

	A	B	C	D	E
1		P1 - guide_curve_angle	P2 - guide-curve-radius	P3 - section-1-length-1	P4 - effective_flow_area
2	Optimization Study				
3	Objective	Minimize	Maximize	No Objective	Maximize
4	Target Value				
5	Importance	Default	Default	Default	Higher
6	GDO Sample Set 1				
7	Candidate A	50.058	30.028	60.502	0.0011801
8	Candidate B	50.063	30.027	60.087	0.0011801
9	Candidate C	50.063	30.027	59.897	0.00118

Customer Benefits:

- Able to quickly identify the key parameters the design is most sensitive to
- Considerable reduction of labor time and chances of human error by automating the whole process

Fan-Heat Sink

Design Objective:

- Optimize the fan-heat sink geometry such that the temperature on the 2 chips is lower than the baseline design (with fixed fan design)

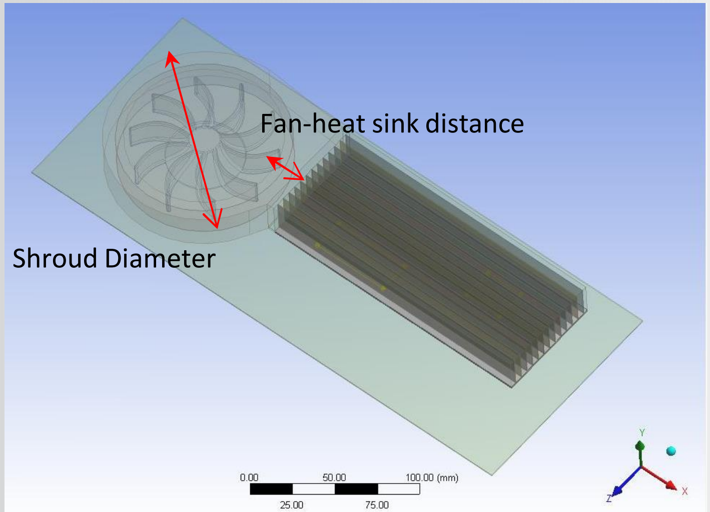

Fan RPM = 2000
Chip 1 = 35000 W/m^2
Chip 2 = 40000 W/m^2

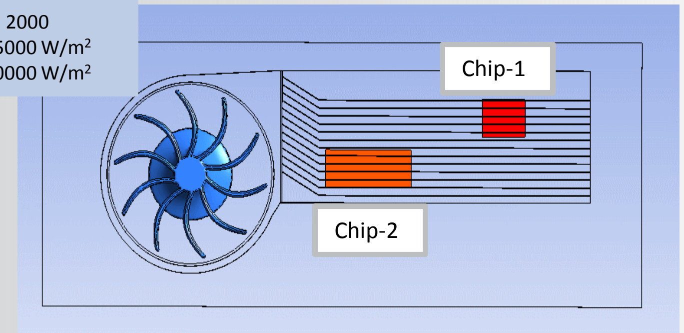

Parameter Manager

```
Shroud_HS_gap = 1
Fin_Angle = 30
Fin_front_length = 25
Shroud_Diameter = 120
Number_of_Fins = 12
```

Input Geometric Parameters

Design Parameters | Parameter/Dimension Assignments | Check

Output Parameters

Output Parameters			
P14	surface-chip1	325.57	K
P15	surface-chip2	323.36	K
P17	area-wt-avg-vel-blow-exit	3.5312	m s^-1
P18	vol-flow-out-blower	-0.0035515	m^3 s^-1
P19	surface-integral-Pressure-hs-inlet	9.9385	Pa
P21	surface-integral-pressure-hs-exit	-0.47845	Pa
P22	surface-integral-press-hs-inlet-exit	4.73	Pa

Fan-Heat Sink

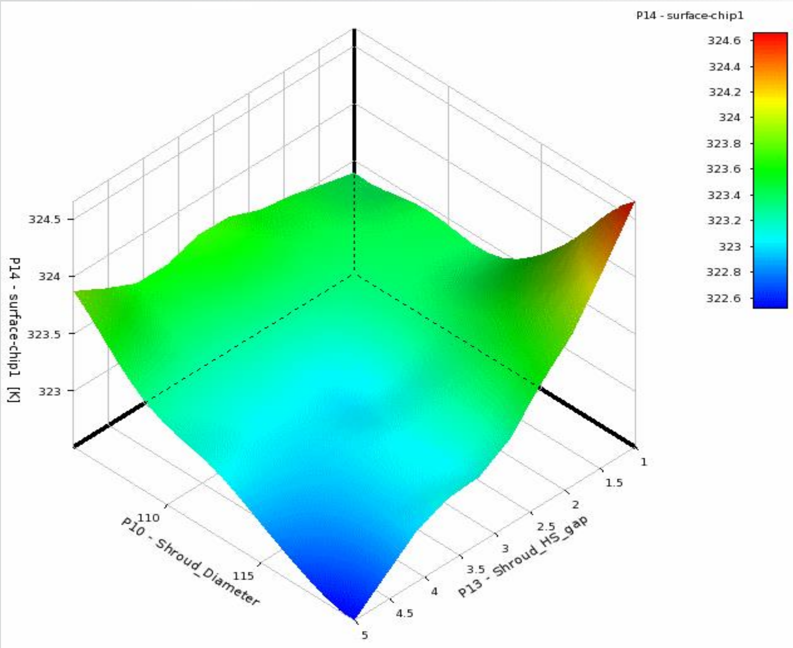

Chip temperature vs.
(Shroud Diameter & Fan Heat Sink gap)

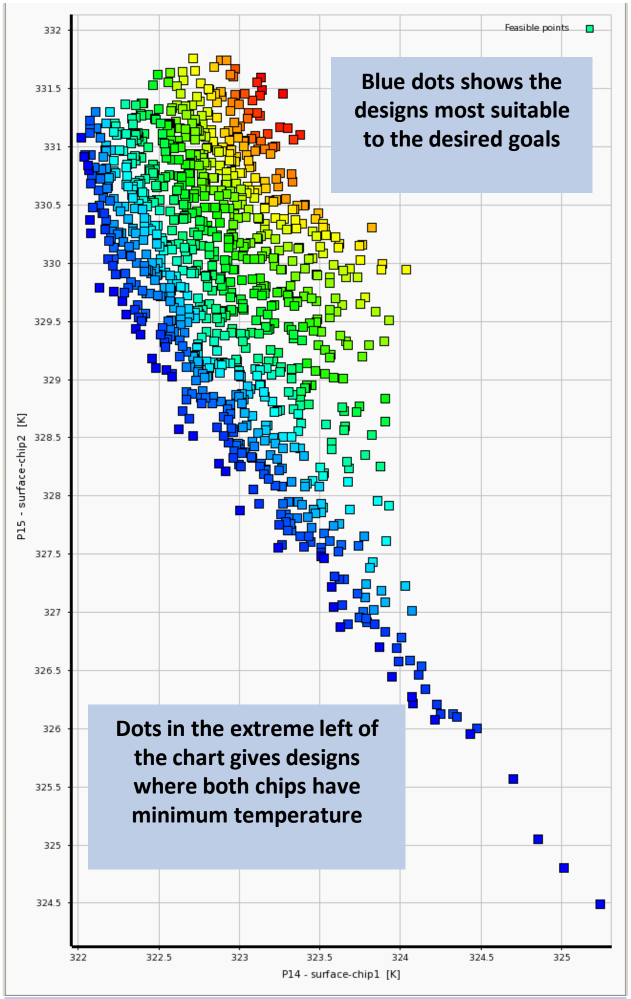

Tradeoff chart between Temperature of Chip1 and Chip2

Customer Benefits:

- Quick understanding of relationship between <u>many</u> design variables and performance
- Easy exploration of a large number of 'optimal' designs (by using trade-off charts)

Ongoing HPC Initiatives

ANSYS focus on HPC is ongoingé

- Å Ongoing optimization and performance tuning
  - Å Dynamic load balancing; optimized resource mapping, compiler evaluation
- – Architecting for next level scalability
  - Å Performance at 10,000 cores or more; increased core density and GP-GPUs
  - Å Innovative mechanical solvers: Multilevel PCG, 2D parallel DSPARSE fronts
  - Å Hybrid distributed/shared memory and vector processing paradigms
- – Scalability across all components and full simulation process
  - Å Meshing, setup, solver, I/O, visualization, optimizationé
  - Å Parallel for linear dynamics, including mode superposition-based analyses
  - Å Distributed domain solver, especially for contact nonlinearities
  - Å Partial factorization (in-core substructuring) for localized nonlinearities
- – Usability
  - Å Multi-component parallel execution environment, job scheduler support
  - Å Hardware fault tolerance, system performance tracking and debugging

All to achieve next-generation capability / performance!

HPC Partnerships

ANSYS maintains close technical collaboration with the leaders in HPC

This mutual commitment ensures that you get the most possible value from your overall HPC investment

Some current examples:

- **Optimized performance on multicore processors from Intel, with R&D focused on Intel's Many Integrated Core (MIC)**
  - **Over <u>60% performance boost</u> for the latest Intel® Xeon® E5-2600 processor (Sandy Bridge) family compared to previous Intel (Westmere) generation**
- **GPU computing accelerates ANSYS Mechanical today, with very active R&D engagement with NVIDIA across full portfolio**
- **ANSYS and IBM – Optimized cluster and storage architectures for ANSYS**
- **ANSYS and Cray – Support for extreme scalability of ANSYS CFD on the Cray XE, up to 1000's of cores**

CADFEM (Austria) GmbH
+43 1 587 70 73 0
info@cadfem.at
www.cadfem.at