Petaflops Opportunities for the NASA Fundamental Aeronautics Program

Dimitri Mavriplis (University of Wyoming)
David Darmofal (MIT)
David Keyes (Columbia University)
Mark Turner (University of Cincinnati)

Overview

- Motivation:
  - NASA Aeronautics used to lead in High Performance Computing (HPC)
  - Why is NASA not present at the HPC table today ?
    - Is Science more important than Engineering ?
  - Do we have a vision of what leading-edge HPC could do for engineering applications in general and ~~aeronautics~~ aerospace in particular ?
- Illustrate the possibilities through sample Grand Challenge Problems
- Identify barriers to progress
- Discuss required areas of investment
- Conclude by examining actions of other communities


G5
G12
G2
G9
G6
G10
G7
G4
G11
G1
G1

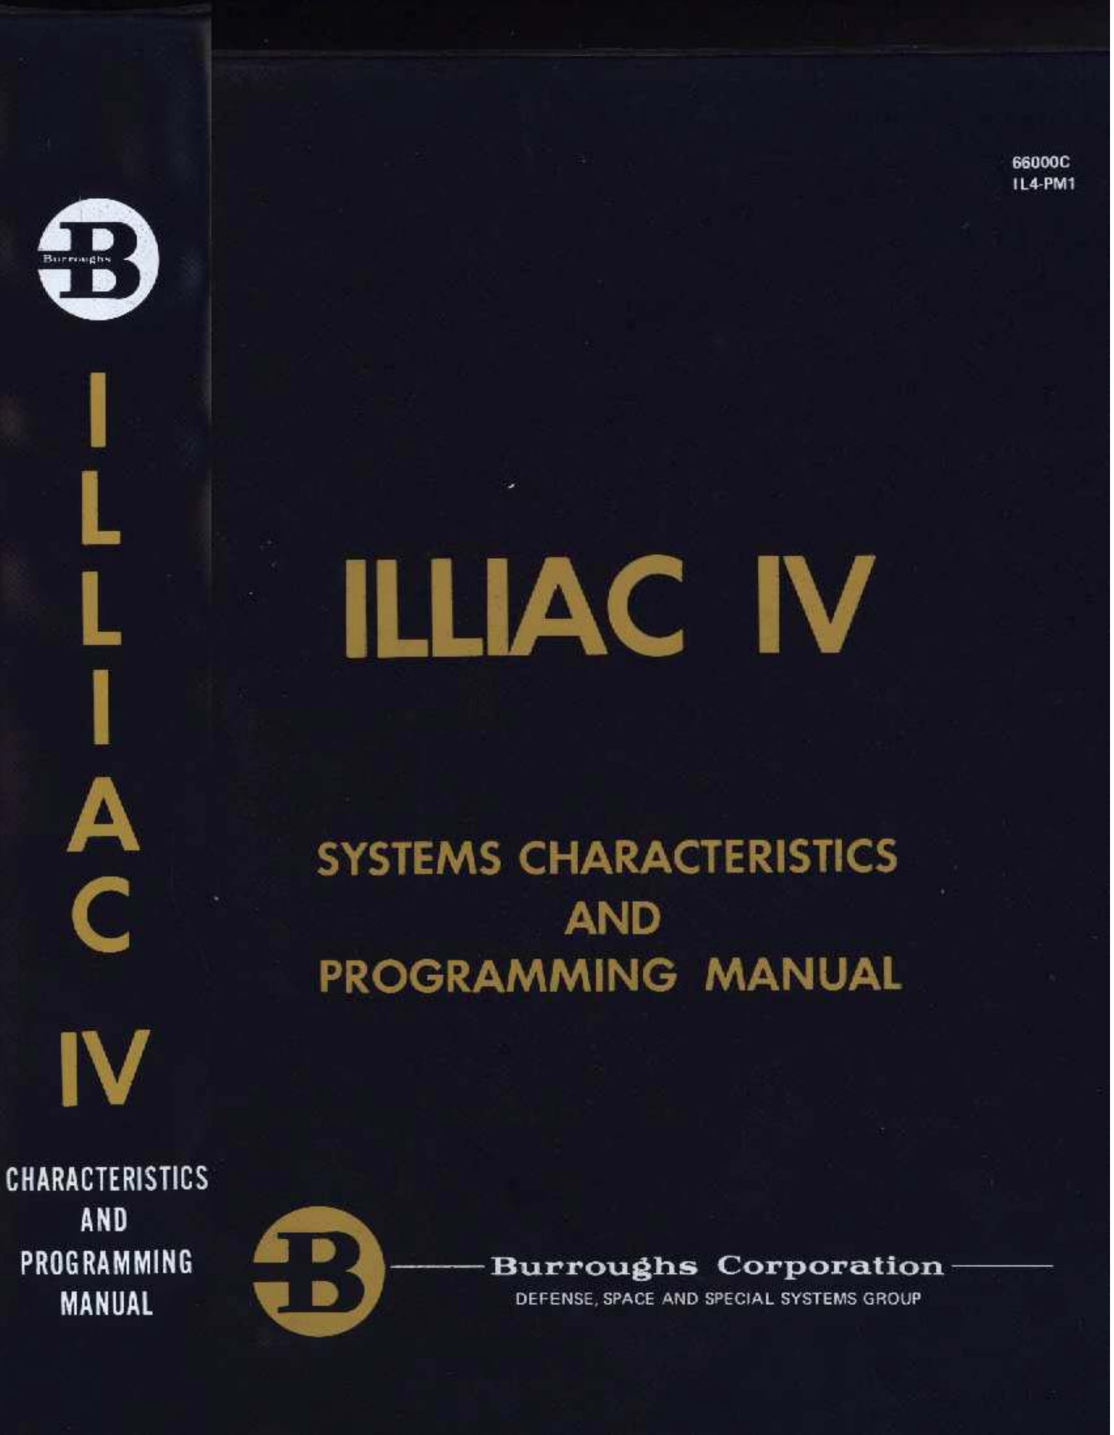

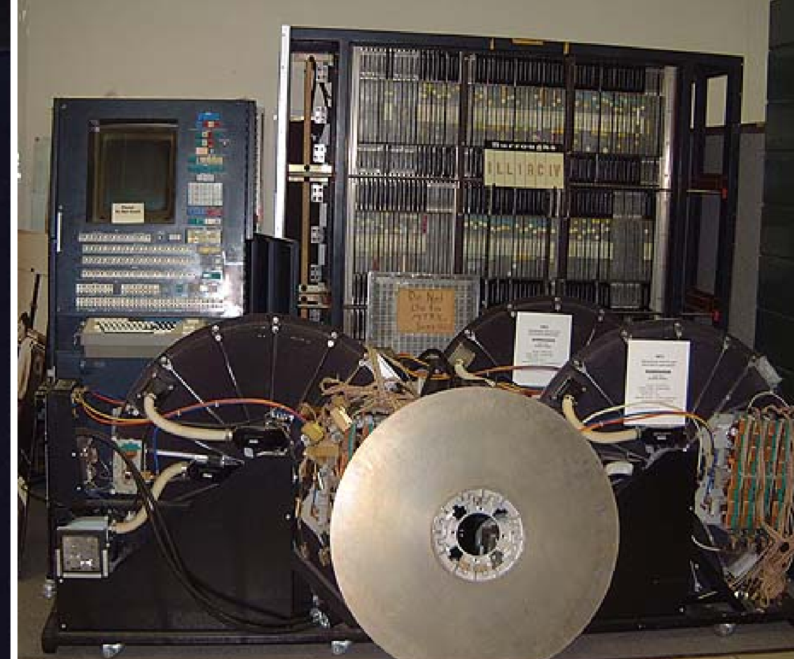

The most powerful computer in the world from 1976 to 1980

- NASA Ames Research Center
- Principal Applications: CFD
- Leading National HPC Driver
- A principal element of applied mathematics research

NASA HPC Leadership Continues …

- 1980’s: National Aerodynamic Simulator (NAS)
  - one of the first: Cray 2, Cray YMP, Cray C90
- 1990’s: High Performance Computing and Communication Program (HPCCP)
  - Transition from small numbers of vector processors to upcoming class of “massively” parallel microprocessors (O(100) cpus)

Some Statistics (circa 1992)

- 1992 HPCCP Budget:
  - $596M (Total)
    - $93M Department of Energy (DOE)
    - $71M NASA
      - Earth and Space Sciences (ESS)
      - Computational Aerosciences (CAS)

- CAS Objectives:
  - "…*integrated, multi-disciplinary simulations and design optimization of aerospace vehicles throughout their mission profiles*"
  - "… *develop algorithm and architectural testbeds … scalable to sustained teraflops performance*"

Fast Forward 2007

- NASA Columbia Supercluster:
  - 10,240 cpus
- Mostly used as capacity (not capability) facility
  - Many “small” jobs of order O(100) cpus
  - Not much progress since 1992
    - Published NASA code benchmarks stop at 512 cpus
    - 512 cpu runs on Intel Touchstone Delta Machine (ICASE/NASA at Supercomputing ’92)
    - Supercomputing’05: Only 1 NASA Paper
  - NASA is no longer a credible HPC player

Science vs. Engineering

- HPC advocacy has increasingly been taken up by the science community
  - Numerical simulation is now the third pillar of scientific discovery on an equal footing alongside theory and experiment
  - Increased investment in HPC will enable new scientific discoveries
- SciDAC, ScaLES, Geociences, NSF Office of Cyberinfrastructure (OCI)….

DOE SciDAC Program

- Scientific Discovery through Advanced Computing
  - Enable new scientific discoveries
  - Initial 5 year program
  - Renewed for 5 years

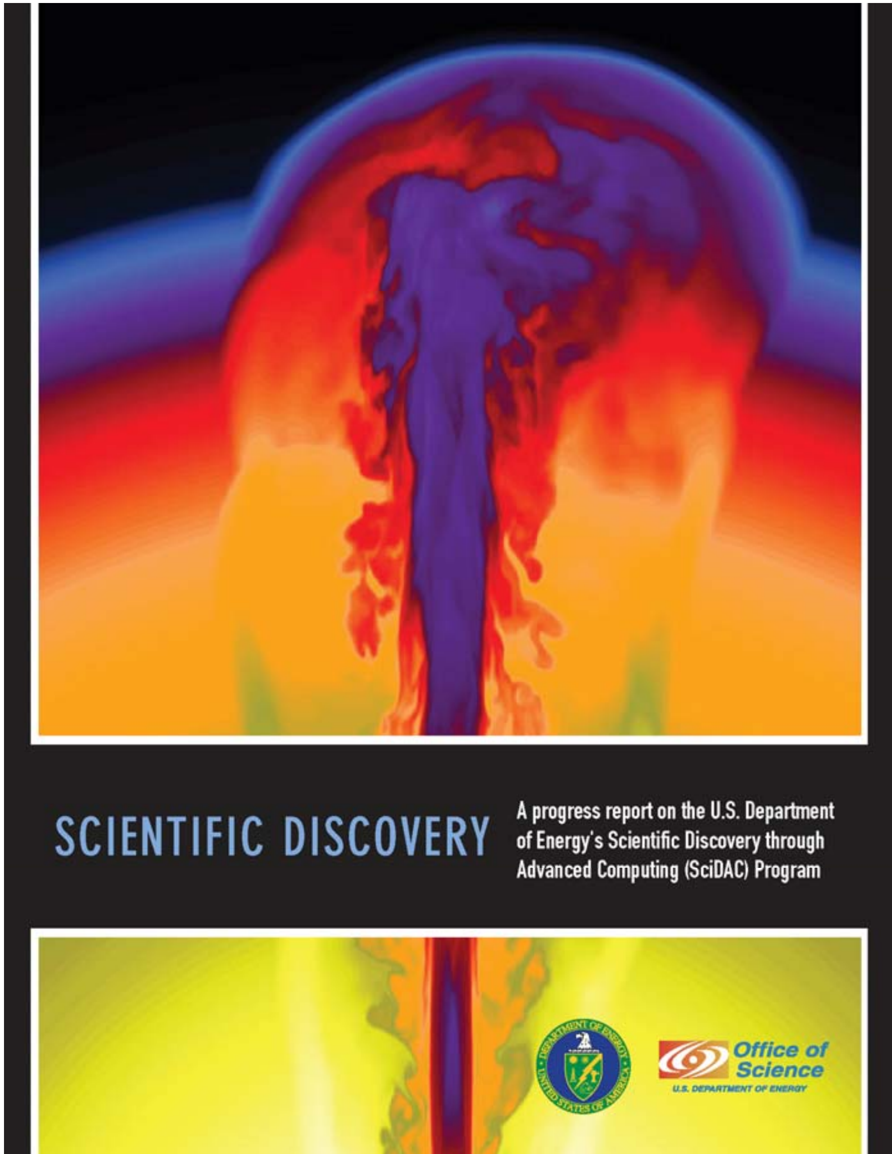

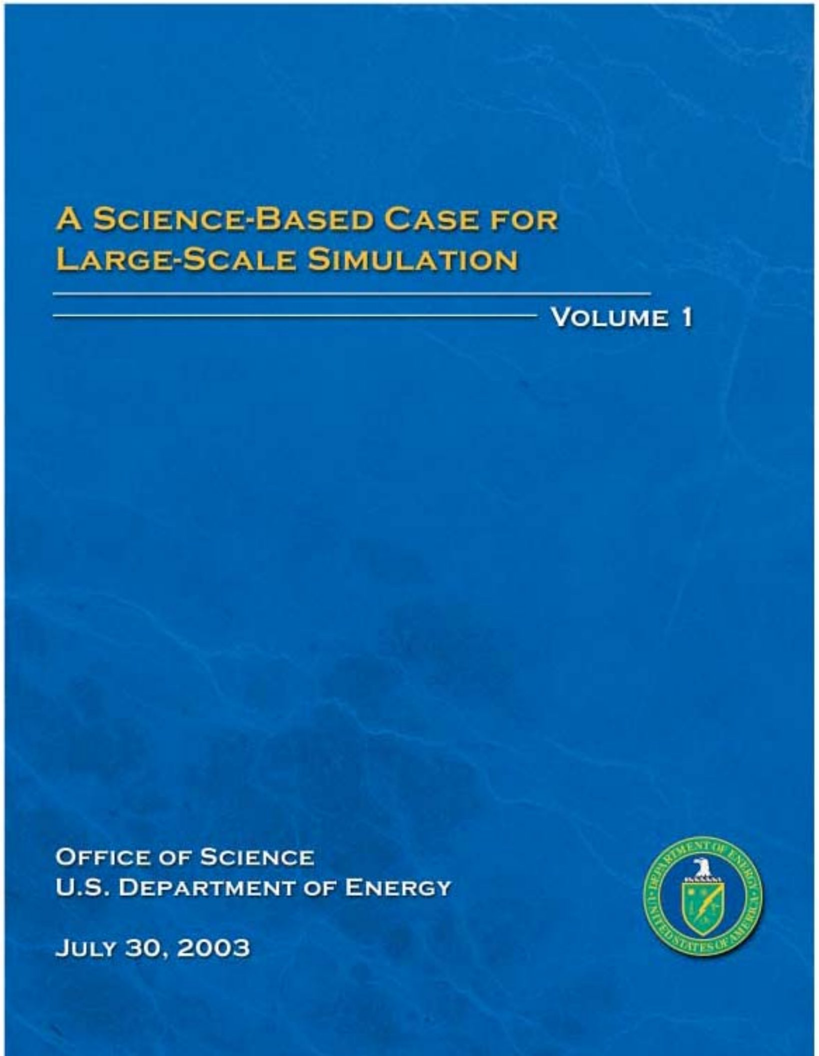

A SCIENCE-BASED CASE FOR LARGE-SCALE SIMULATION

Volume 2

Office of Science
U.S. Department of Energy

September 19, 2004

"There will be opened a gateway and a road to a large and excellent science, into which minds more piercing than mine shall penetrate to recesses still deeper." Galileo (1564-1642)

[*on the "experimental mathematical analysis of nature," appropriated here for "computational simulation"*]

ESTABLISHING A PETASCALE COLLABORATORY FOR THE GEOSCIENCES

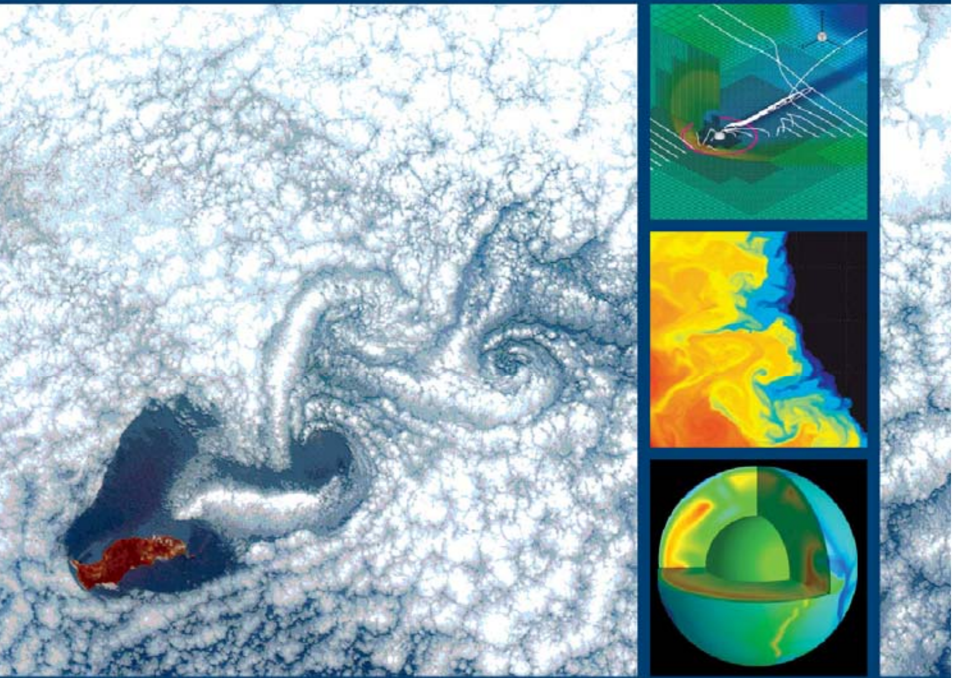

Scientific Frontiers

ESTABLISHING A PETASCALE COLLABORATORY FOR THE GEOSCIENCES

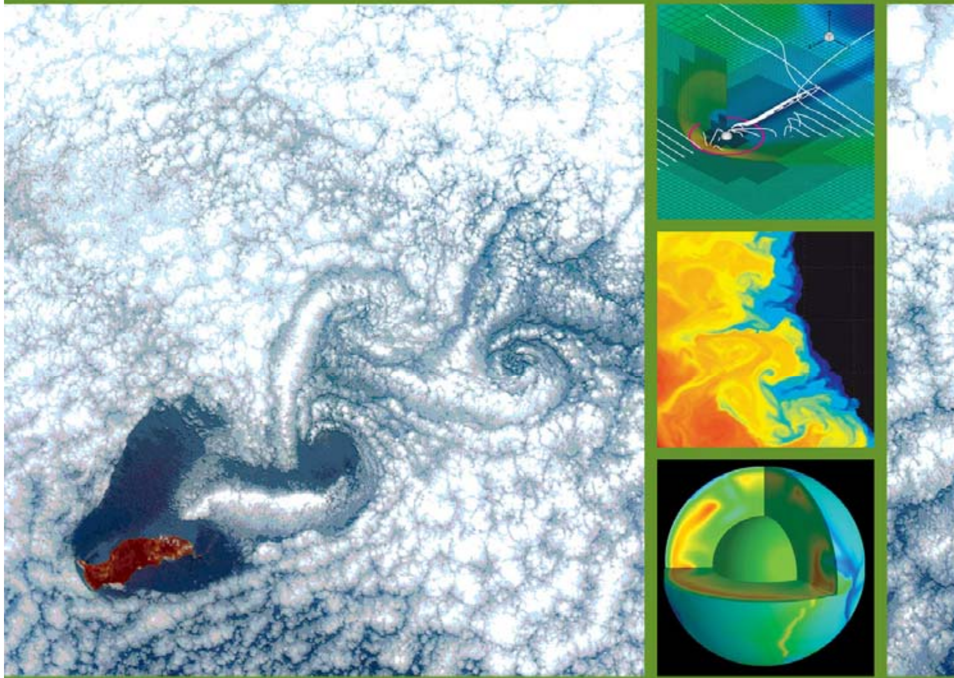

Technical and Budgetary Prospectus

Engineering Community

- Engineering in general and NASA Aero in particular:
  - Our problems are not complex enough to warrant such large scale simulations and hardware costs
  - Prefer to reduce cost of current simulation (i.e. move to a cluster) instead of increasing the simulation capability at fixed cost (on best available hardware)
  - That is intractable !
    - Doing time dependent MDO
      - Need to store entire time dependent solution history
    - Commonplace for large science applications today
      - Data asssimilation in atmospheric science (NCAR)
      - Inverse problems in earthquake simulation (San Diego Center)
- Commodity simulation on commodity hardware for commodity engineering
  - Our expertise is in systems integration (only!)…

Resurgence of HPC Nationally

- American Competitiveness Initiative (2006)
- Preceded by numerous studies and recommendations on the need for increased investment in HPC
  - NITRD (2005)
  - PITAC (2005)
  - NSF Simulation based Engineering Science (2006)

Report to the President

June 2005

Computational Science: Ensuring America's Competitiveness

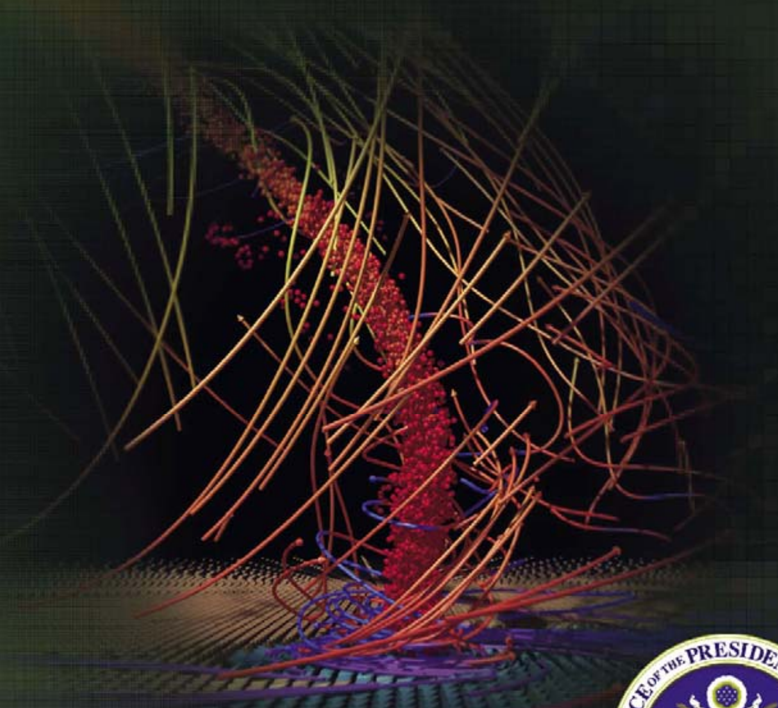

President's Information Technology Advisory Committee

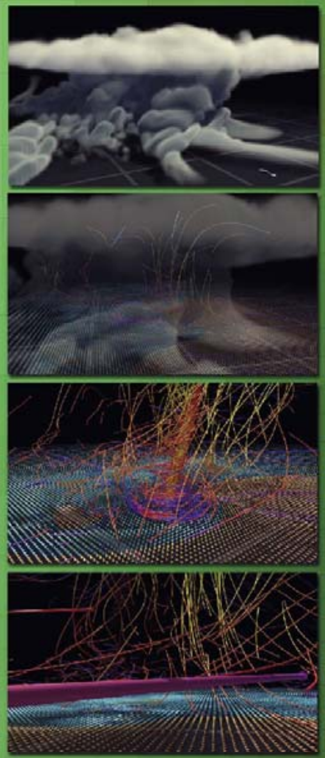

National Coordination Office for Information Technology Research and Development
Suite II-405
4201 Wilson Boulevard
Arlington, Virginia 22230
(703) 292-4873
Email address: nco@nitrd.gov
Web address: http://www.nitrd.gov/pitac

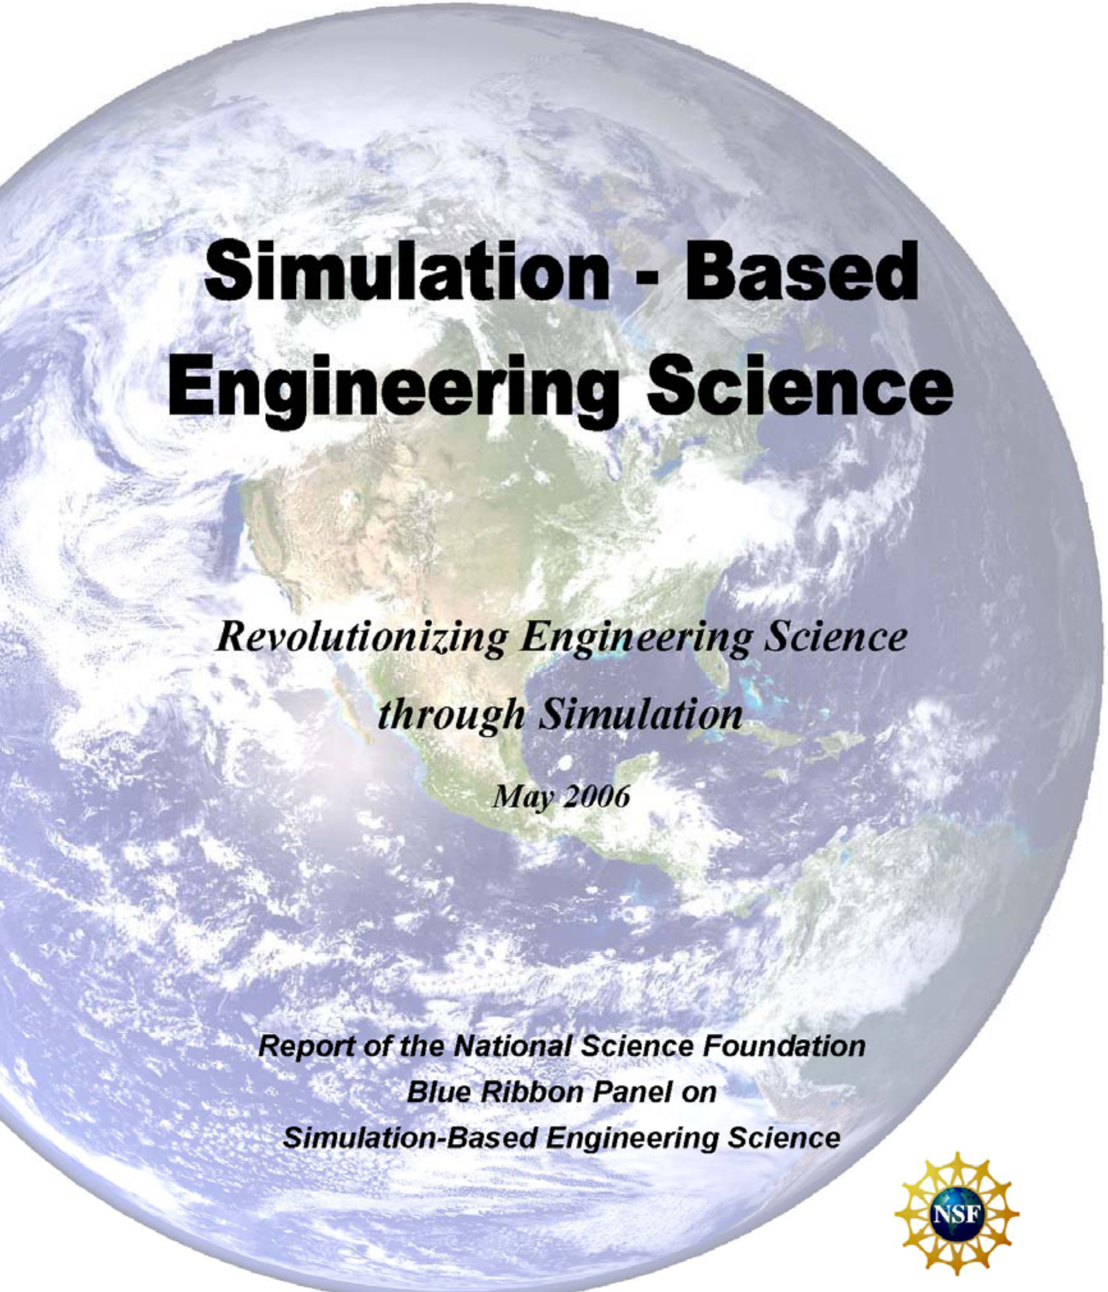

- Recent NSF Report
  - Engineering based simulation needs more attention
    - Science has been successful recently as advocate
  - Mainly structures, crash dynamics, materials
  - No mention of aeronautics activities

NASA's Missed Opportunity

- NITRD 2005:
  - No mention of NASA HPC at all
- PITAC 2005:
  - Aerospace HPC only mentioned briefly (and erroneously)

- Competitiveness Initiative Allocates $ for:
  - National *Science* Foundation
  - DOE Office of *Science*
  - NIST
  - Engineering small player, NASA not a player

- Isn't Engineering as important (or more) than Science with respect to National Competitiveness ?
  - Ask Louis Gallois or Jim McNerney

Reformulated NASA Aeronautics Program

- In their own words:
  - *"..long-term cutting-edge research in the core aeronautics disciplines across all flight regimes…"*
  - *"… aerospace research that benefits the community broadly…"*
    - *(L. Porter, Congressional Testimony Sept. 2006)*

- Decadal Survey of Civil Aeronautics (NAE):
  - *"…an important benefit of advances in physics-based analysis tools is the new technology and systems frontiers they open"*

- Perfectly aligned with a competitiveness initiative
  - Opportunity to re-engage HPC at national level
  - Opportunity to resume (broader) role as driver for engineering simulation at national level

Barriers and Challenges

- A long term vision is needed to:
  - Identify perceived and real barriers
    - Our problems don’t require more computing power
    - That is intractable
    - How to run on 100,000 cpus
    - How to solve bigger more difficult problems
  - Demonstrate the potential new frontiers to be opened by increased simulation capabilities
  - Identify required areas of investment
- Grand Challenges are a means, not an end

Selected Grand Challenges

- Digital Flight
  - Static (and dynamic) aerodynamic data-base generation using high-fidelity simulations
  - Time-dependent servo-aero-elastic maneuvering aircraft simulations
- Transient Full Turbofan Simulation
- New frontiers in multidisciplinary optimization
  - Time dependent MDO
  - MDO under uncertainty
- Examples only (not all inclusive)
  - e.g. Aeroacoustics not mentioned

Flight-Envelope Data-Base Generation (parametric analysis)

- Configuration space
  - Vary geometric parameters
    - Control surface deflection
    - Shape optimization
  - Requires remeshing
- Wind-Space Parameters
  - Vary wind vector
  - Mach, α:incidence, β:sideslip
  - No remeshing
- Completely Automated
  - Hierarchical Job launching, scheduling
  - Data retrieval
  - Failure recovery

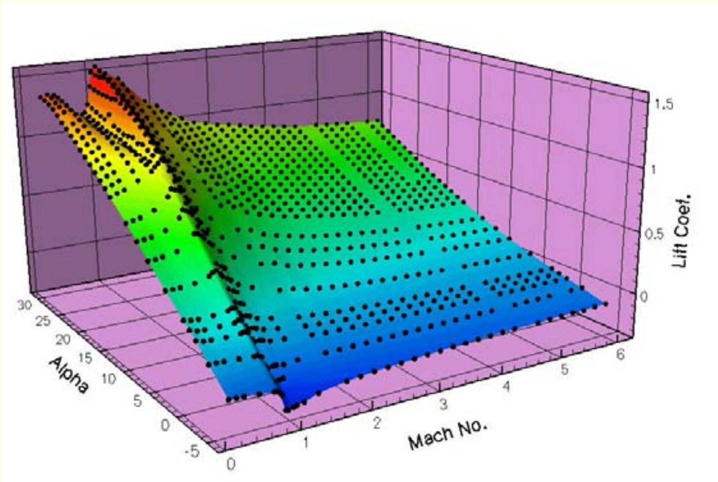

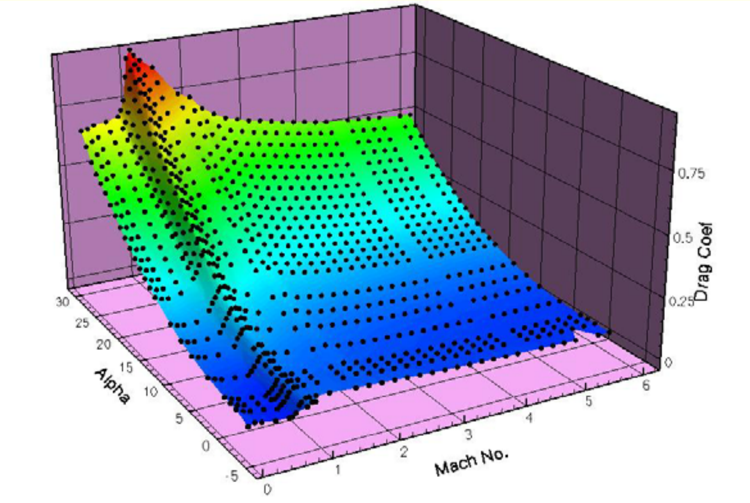

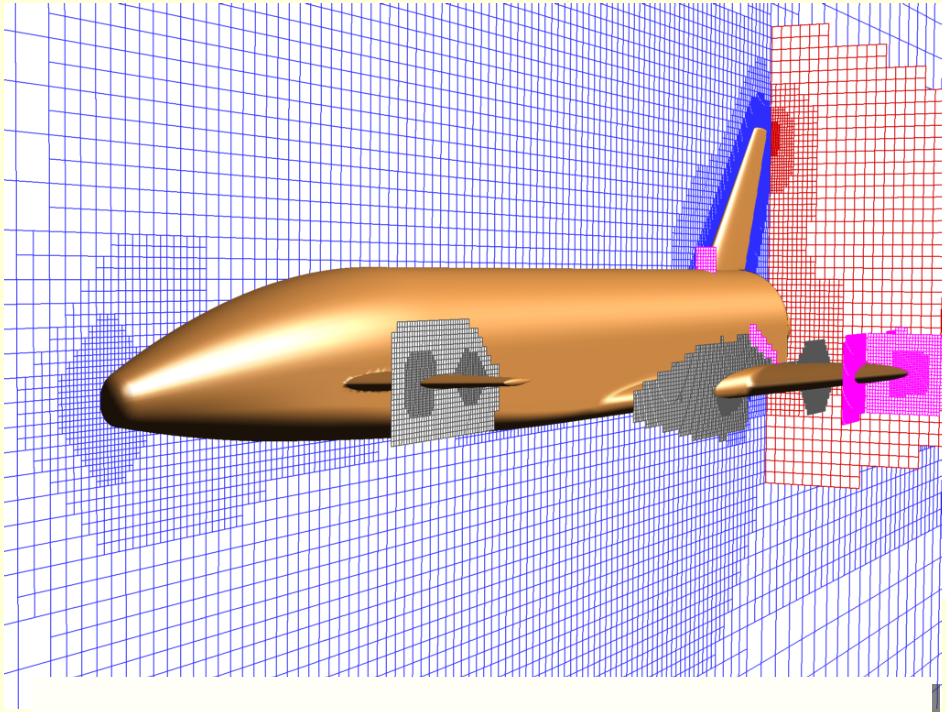

- Liquid glide-back booster
  - Crank delta wing, canards, tail
- Wind-space only

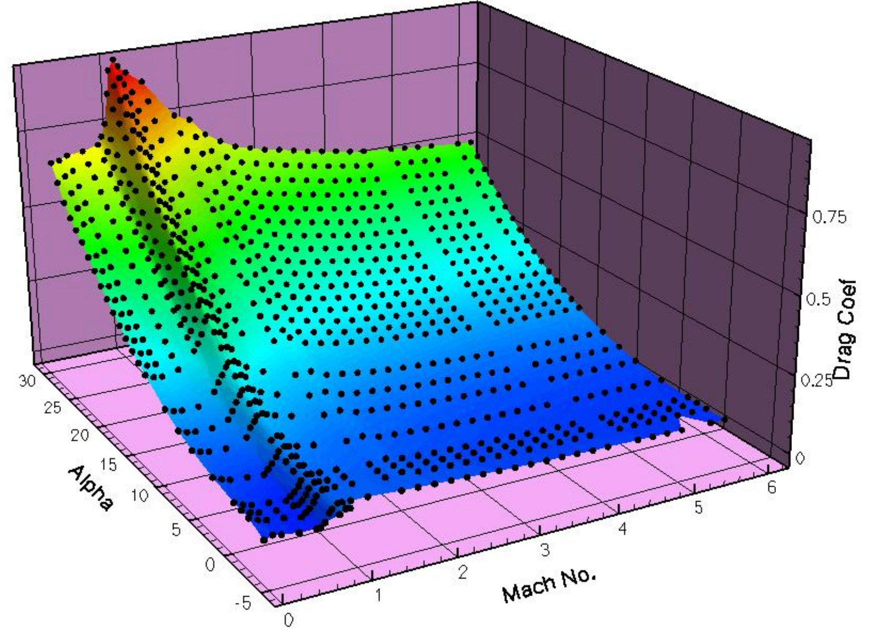

- Wind-Space:
  $M_\infty = \{0.2\text{-}6.0\}$, $\alpha = \{-5°\text{–}30°\}$, $\beta = \{0°\text{–}30°\}$
- P has dimensions (38 x 25 x 5)
- 2900 simulations

- Typically smaller resolution runs (Cart3d: Inviscid)
  - 32-64 cpus each
  - Farmed out simultaneously (PBS)
  - 2900 simulations

Computational Requirements

- Based on current NASA experience
  - Overflow: 8 million points, 211 simulations, 1 week
- Assuming:
  - 100 million grid points (RANS)
  - Additional parameters → $O(10^5)$ cases
  - Data-base generation in 1 week using 100,000 cpus
    - Available today (LLNL)
    - Wait 15 years for Moore's Law
    - Sooner using model reduction techniques
- Will this capability be ready when hardware is available at reasonable cost ?
  - Not if no investment is made today

Digital Flight Simulation Example
(c/o A Schutte, DLR)

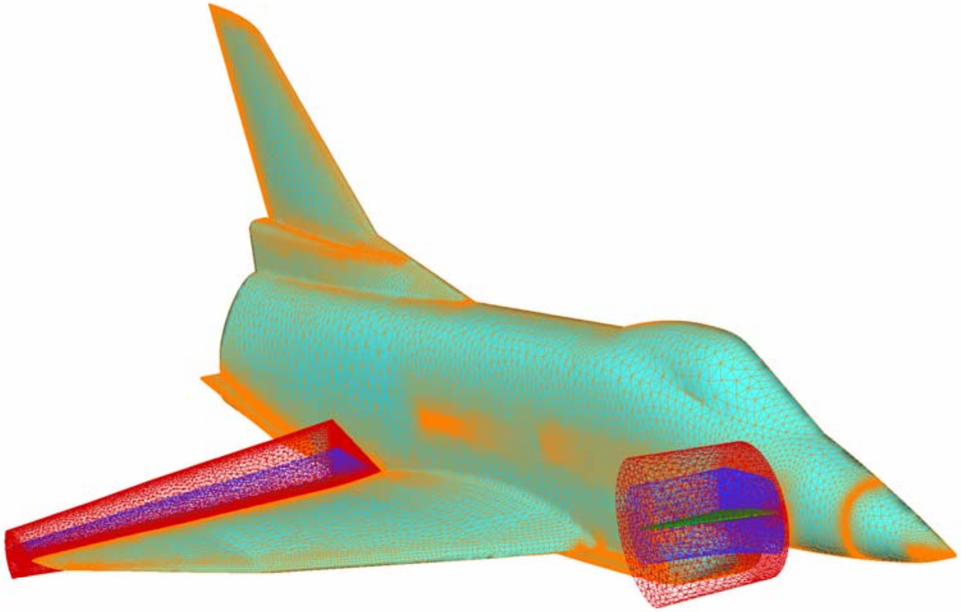

Time accurate multidisciplinary maneuvering aircraft simulation

- •Aero/structure/flight-control system
- •Requirements:
  - •Movable control surfaces
    - --Overset meshes
  - •Complex separated flows
    - --Adaptive meshes
  - •Strong/transient coupling
  - •Currently limited to inviscid flow simulations

- A. Schutte, G. Einarsson, A. Raichle, B. Schoning, M. Orlt, J. Neumann, J. Arnold, W. Monnich, and T. Forkert. *Numerical simulation of maneuvering aircraft by aerodynamic, flight mechanics and structural mechanics coupling*. **AIAA Paper 2007-1070**, presented at the 45th AIAA Meeting, Reno NV., January 2007.

Digital Flight Simulation Example
(c/o A Schutte, DLR)

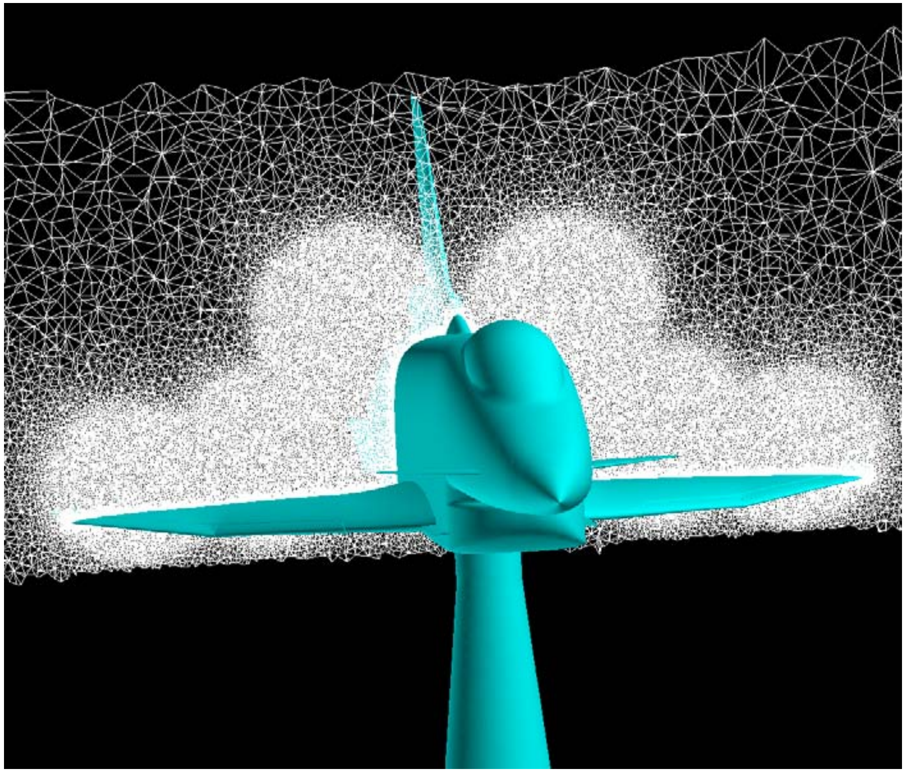

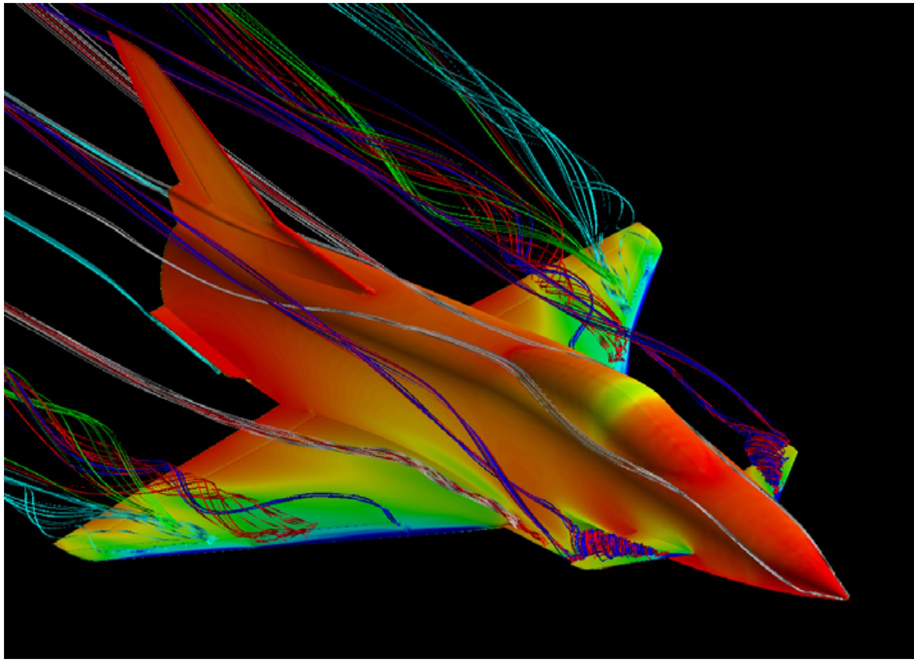

- A. Schutte, G. Einarsson, A. Raichle, B. Schoning, M. Orlt, J. Neumann, J. Arnold, W. Monnich, and T. Forkert. *Numerical simulation of maneuvering aircraft by aerodynamic, flight mechanics and structural mechanics coupling*. **AIAA Paper 2007-1070**, presented at the 45th AIAA ASM Meeting, Reno NV., January 2007.

Digital Flight: Trimming Example
(c/o A Schutte, DLR)

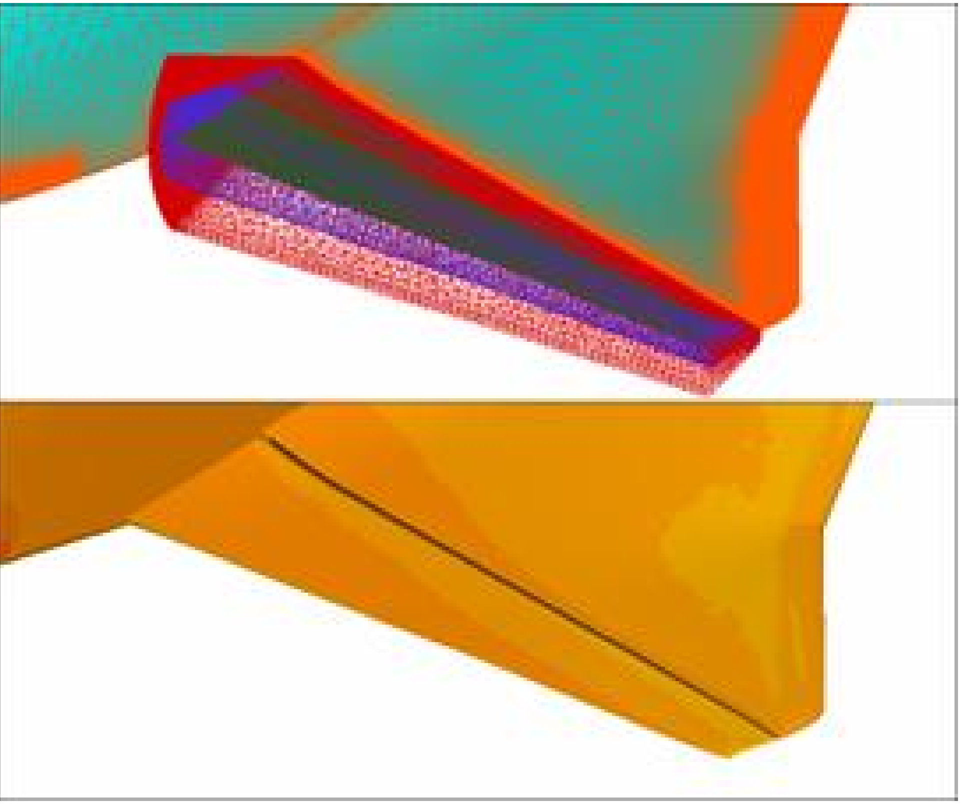

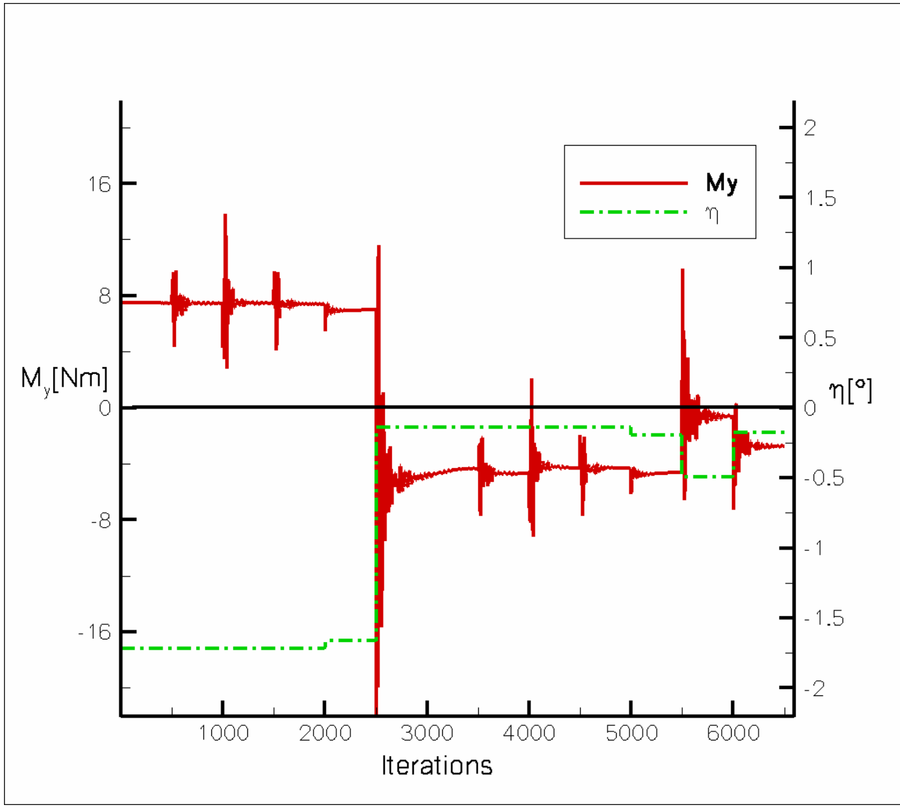

- A. Schutte, G. Einarsson, A. Raichle, B. Schoning, M. Orlt, J. Neumann, J. Arnold, W. Monnich, and T. Forkert. *Numerical simulation of maneuvering aircraft by aerodynamic, flight mechanics and structural mechanics coupling*. **AIAA Paper 2007-1070**, presented at the 45th AIAA ASM Meeting, Reno NV., January 2007.

Digital Flight: Free to Roll Maneuver
(c/o A Schutte, DLR)

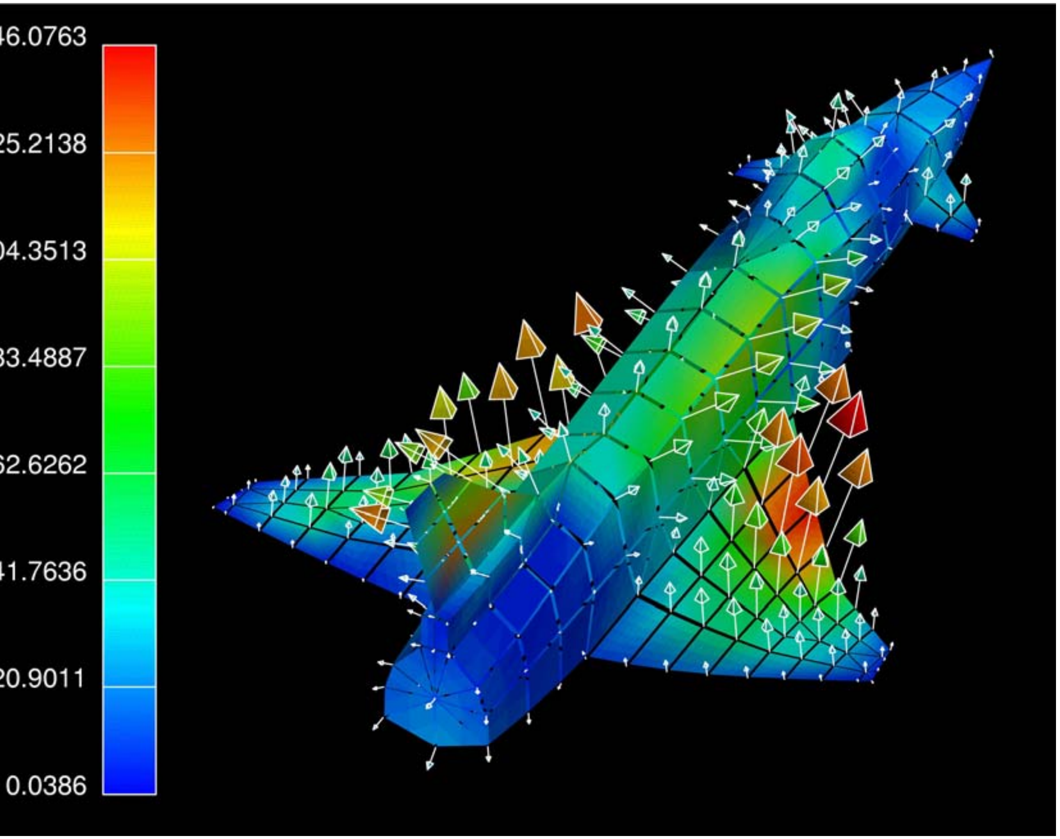

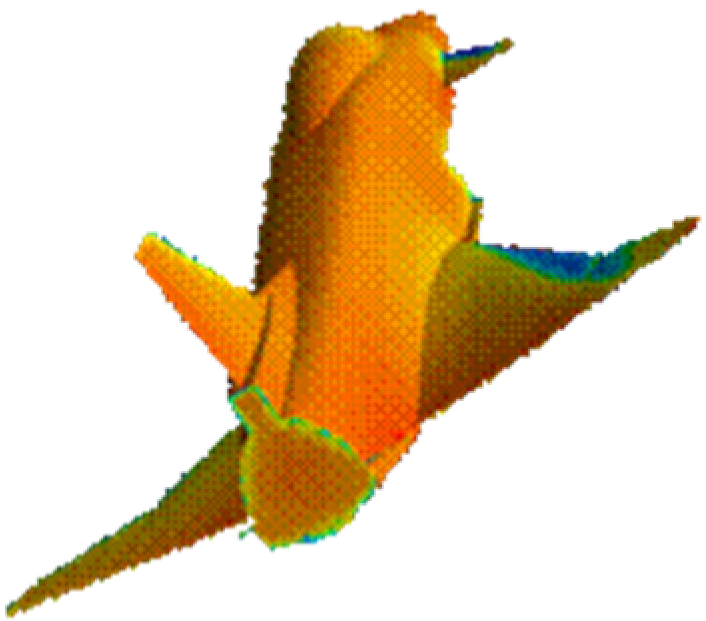

- A. Schutte, G. Einarsson, A. Raichle, B. Schoning, M. Orlt, J. Neumann, J. Arnold, W. Monnich, and T. Forkert. *Numerical simulation of maneuvering aircraft by aerodynamic, flight mechanics and structural mechanics coupling*. **AIAA Paper 2007-1070**, presented at the 45th AIAA ASM Meeting, Reno NV., January 2007.

Digital Flight: Free to Roll Maneuver
(c/o A Schutte, DLR)

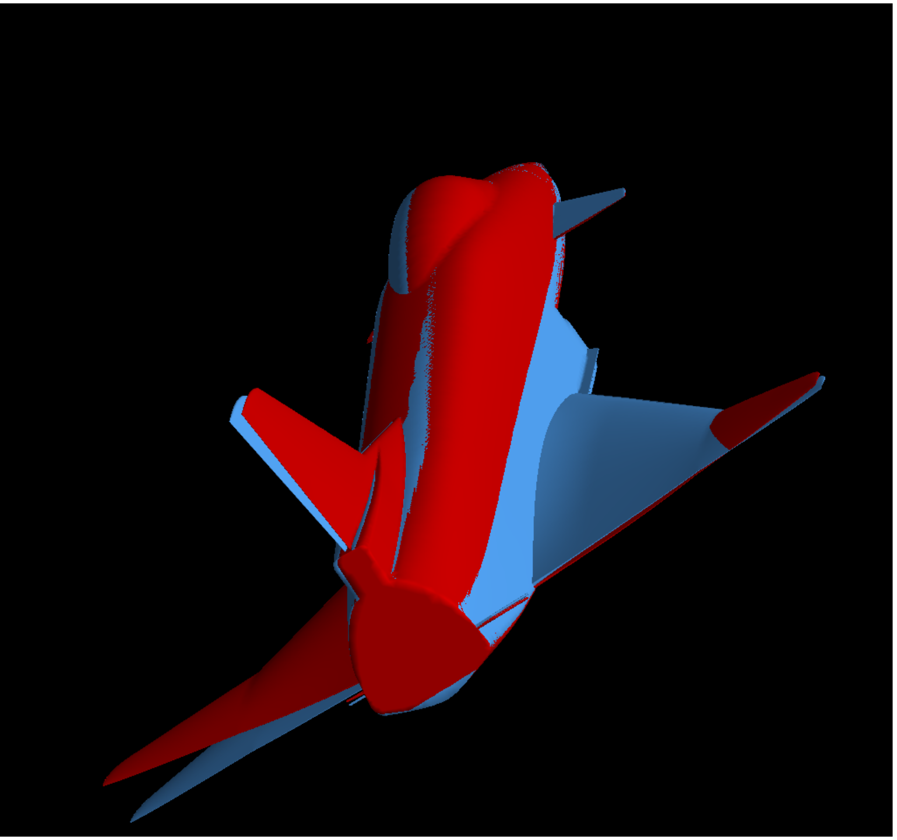

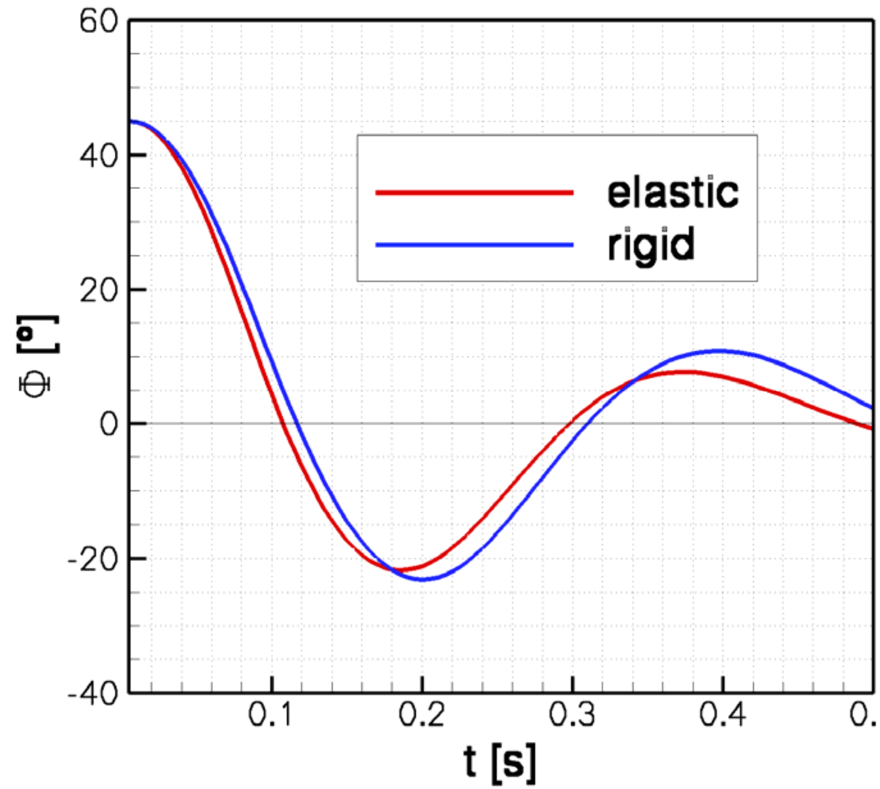

- A. Schutte, G. Einarsson, A. Raichle, B. Schoning, M. Orlt, J. Neumann, J. Arnold, W. Monnich, and T. Forkert. *Numerical simulation of maneuvering aircraft by aerodynamic, flight mechanics and structural mechanics coupling*. **AIAA Paper 2007-1070**, presented at the 45th AIAA ASM Meeting, Reno NV., January 2007.

Computational Requirements

- From M. D. Salas (2006): *Digital Flight: The last CFD Aeronautical Grand Challenge*
  - 60 seconds of flight = 1.5 days on 512 cpus
    - NASA codes, 50 million grid points, 50Hz time stepping
- Easily add:
  - Order of magnitude in grid resolution
  - Order of magnitude in time resolution
  - Multidisciplinary:
    - Structures, Heating, Flight control system
  - Overnight turnaround on 10,000 cpus

Full Turbofan Simulation

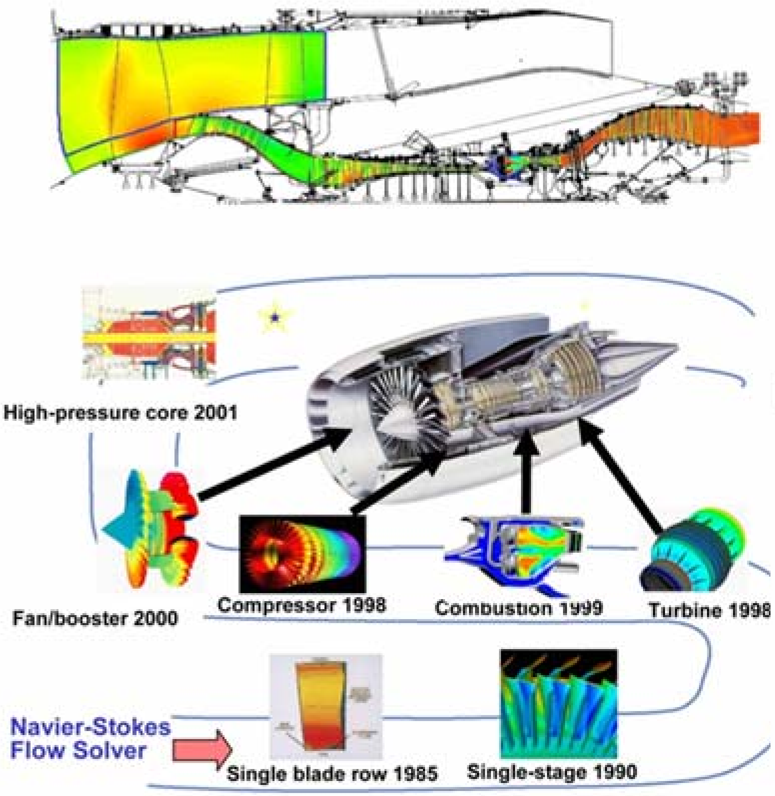

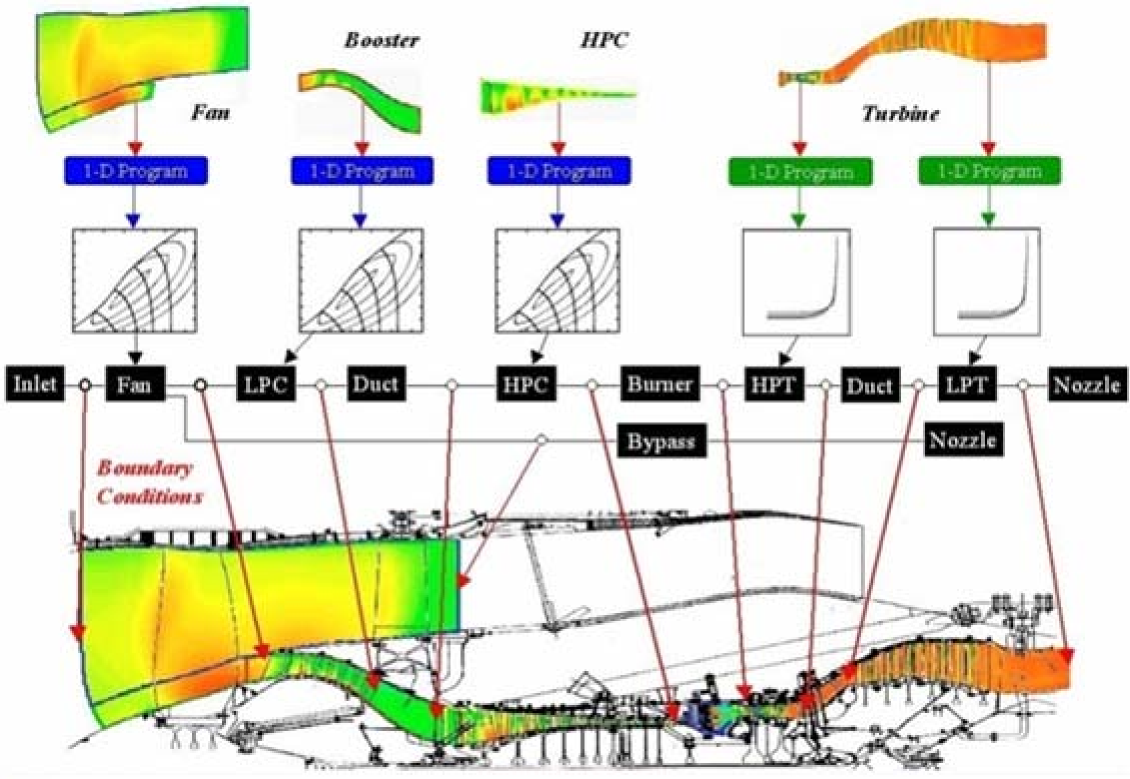

DOE ASC/Stanford Effort

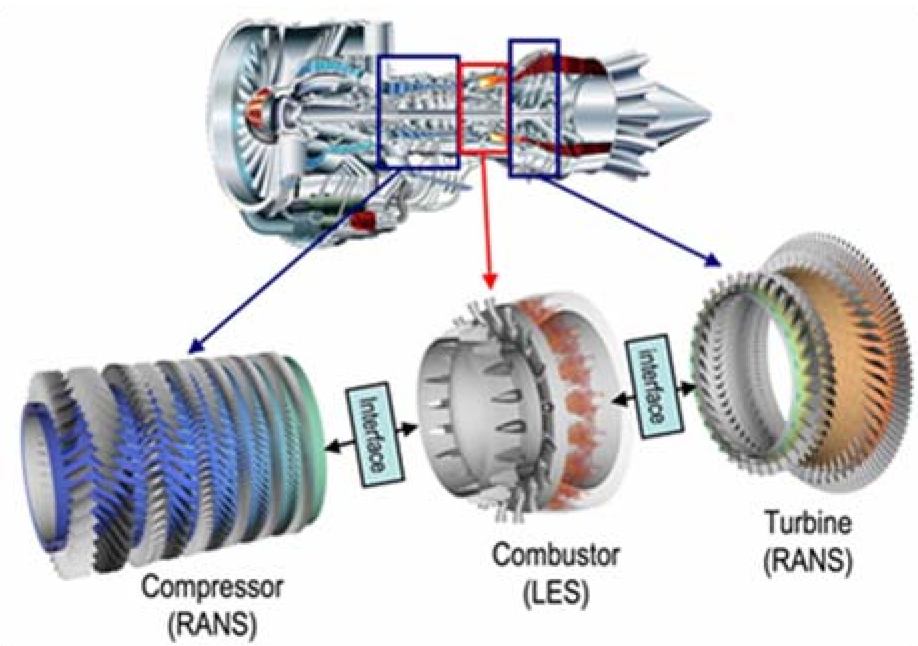

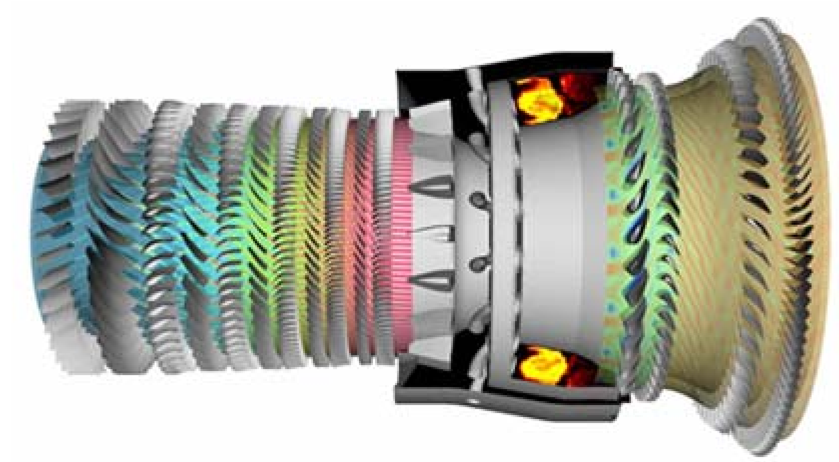

DOE ASC/Stanford Effort

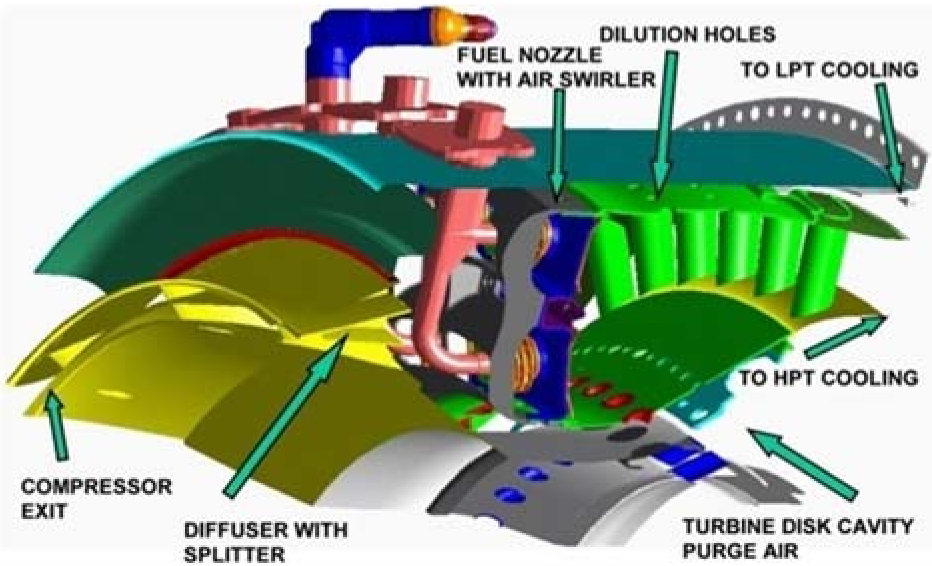

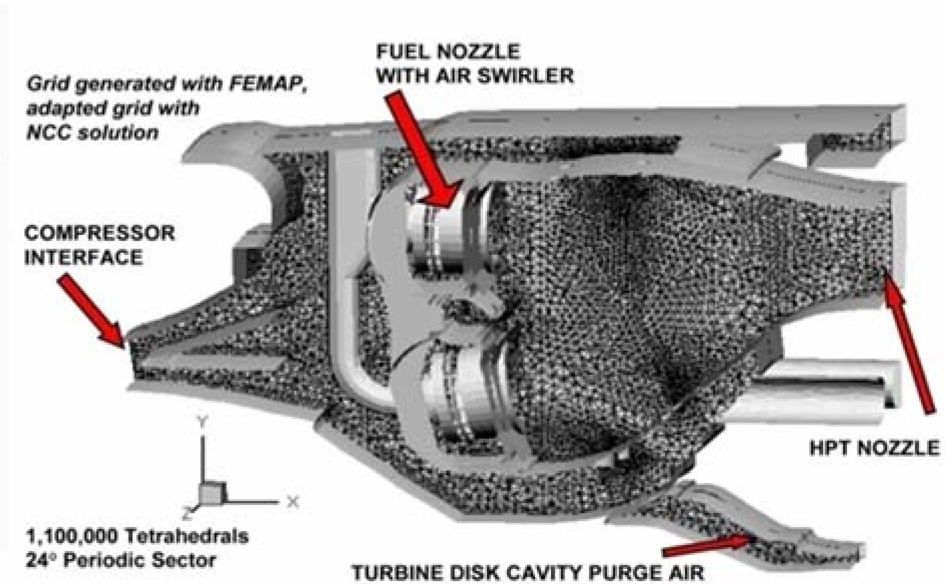

Full Turbofan Simulation

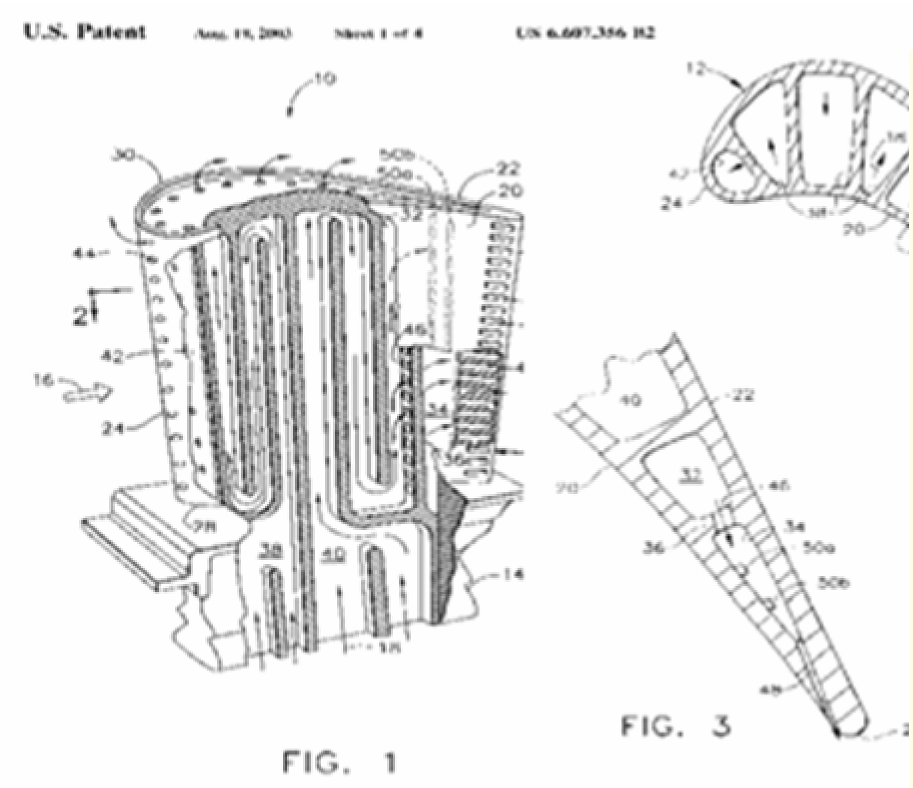

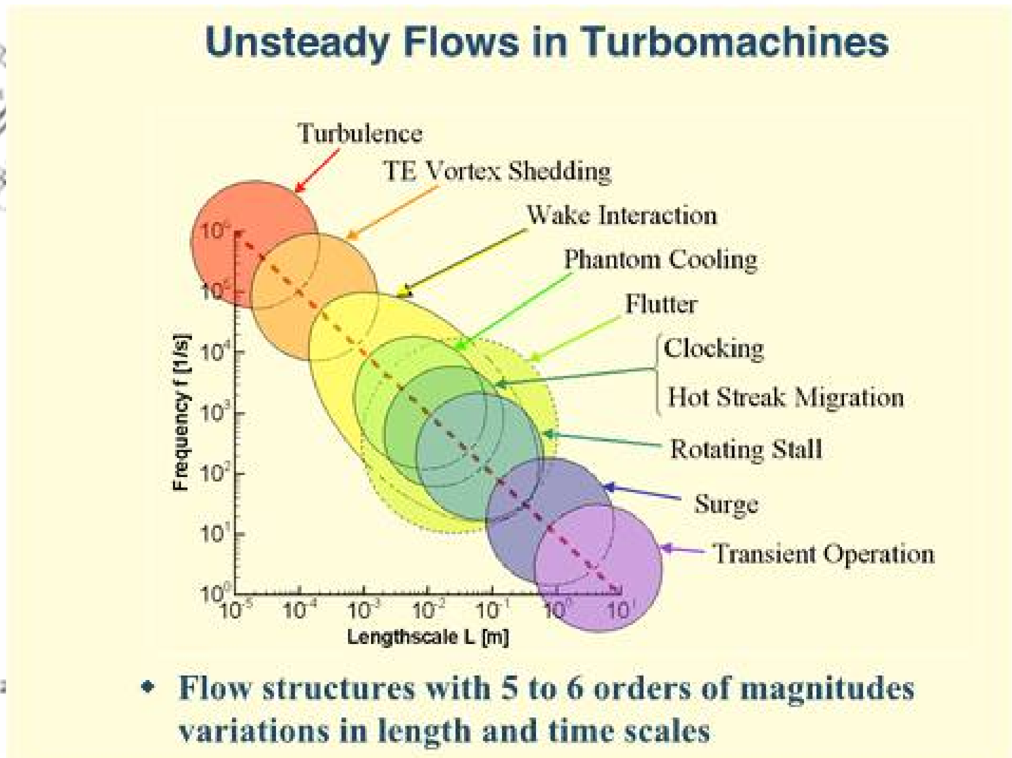

Full Turbofan Simulation

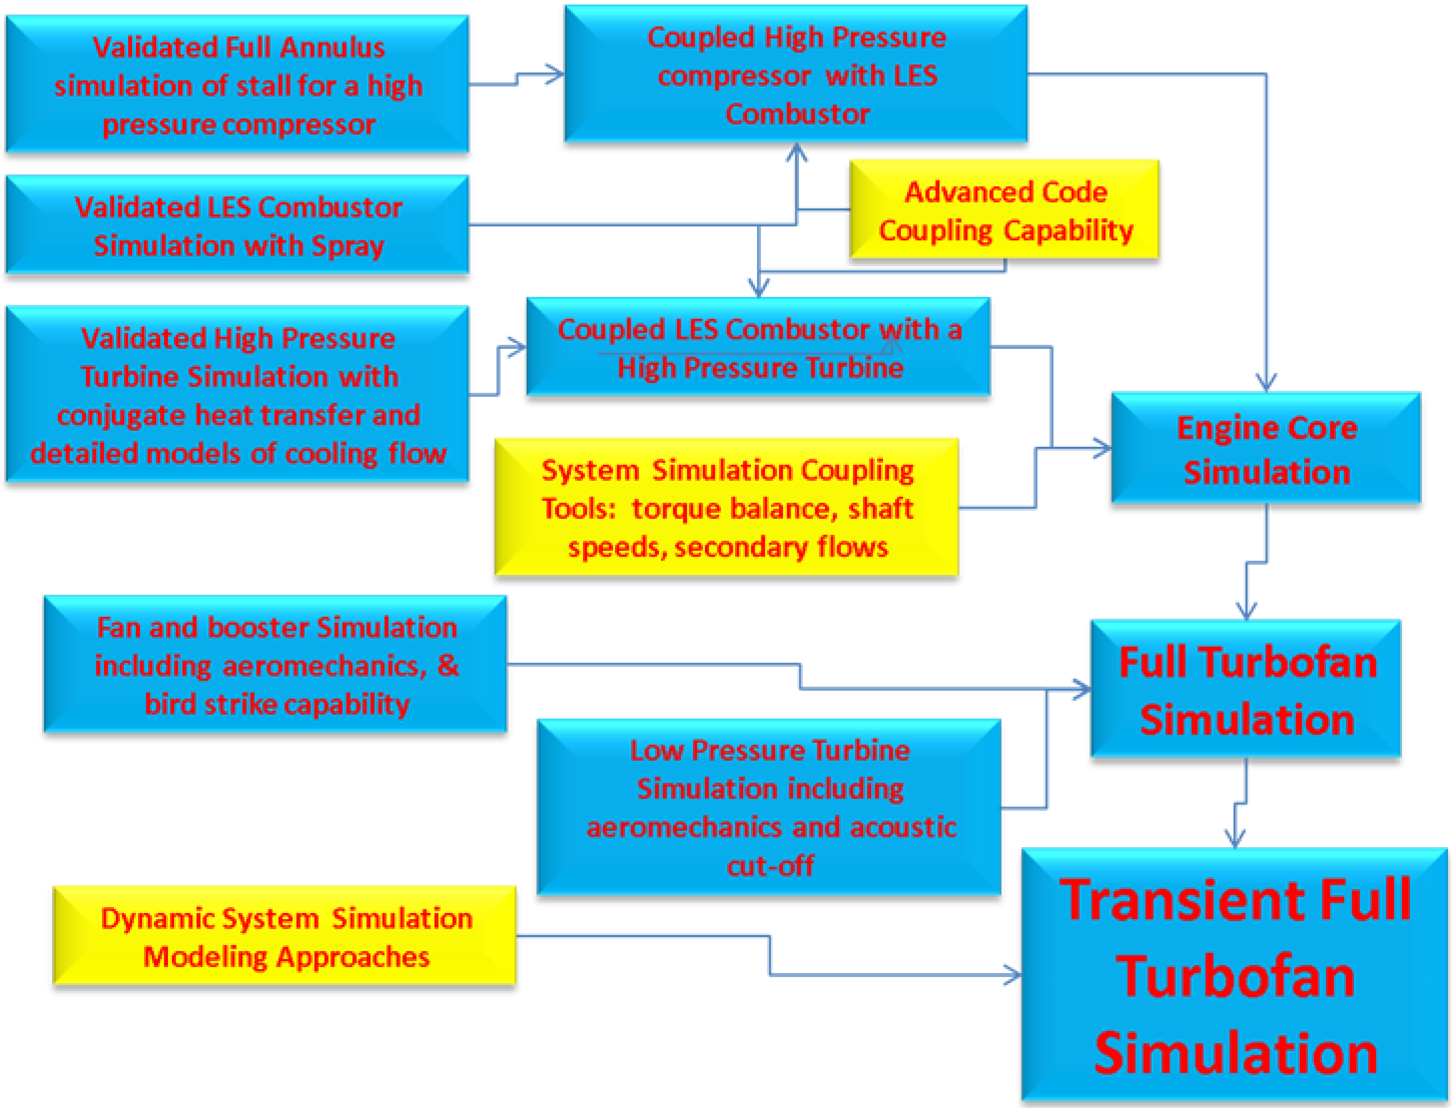

Grand Challenge: Computational Design & Optimization

- Computational engineering ultimately concerned with design (as opposed to computational science)
- Previous examples have focused on improving design through higher fidelity modeling
- Computation can be used to improve design in other manners, e.g.
  - Optimization methods
  - Design under uncertainty

Design Optimization Challenges

- Unsteady Multidisciplinary Design Optimization:
  - Adjoint methods require backwards integration in time
  - Requires entire time-dependent solution set to be stored (to disk)
- Design under uncertainty
  - Ensemble averages for uncertainty estimation
  - Stochastic methods

Rotorcraft Applications

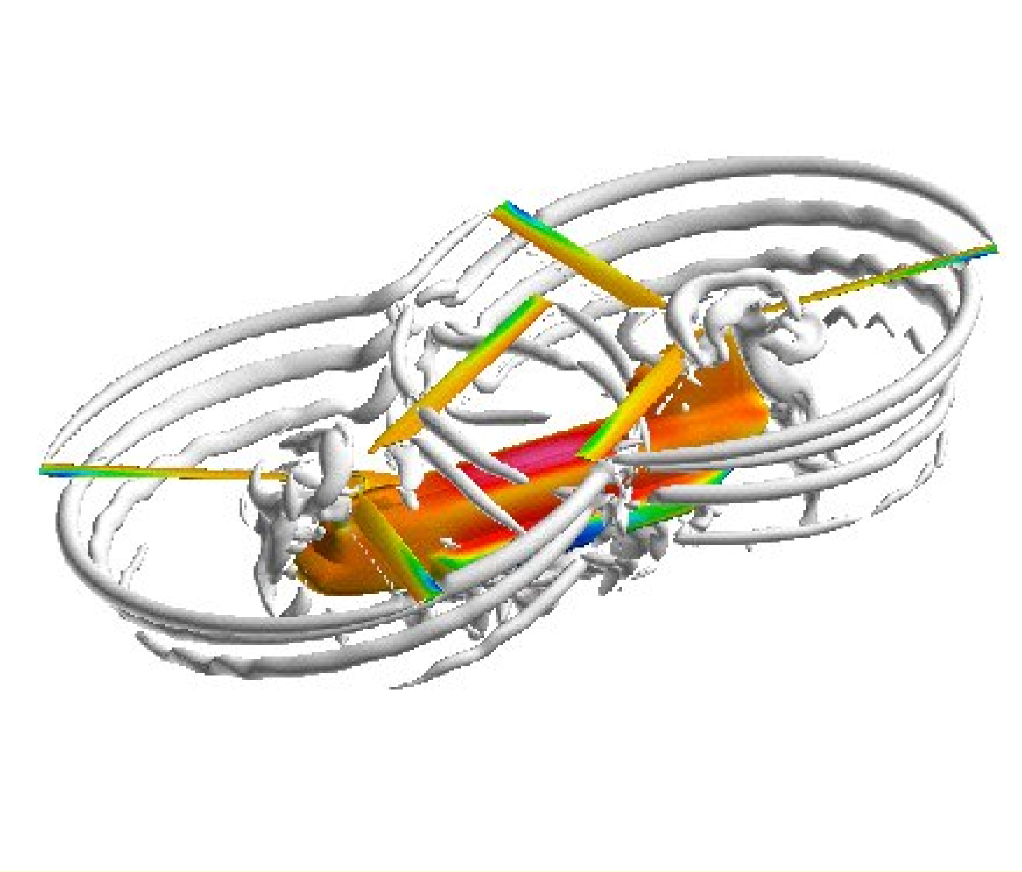

- Complex geometry
- Inherently unsteady
- 78 million grid points
- 12.5 hours on 256 cpus (IBM Power 5) for one revolution
- Runs up to 100 million grid points

Dimanlig, Meadowcroft Strawn and Potsdam: “Computational modeling of the CH-47 helicopter in hover” May 2007.

Time Dependent Design Optimization

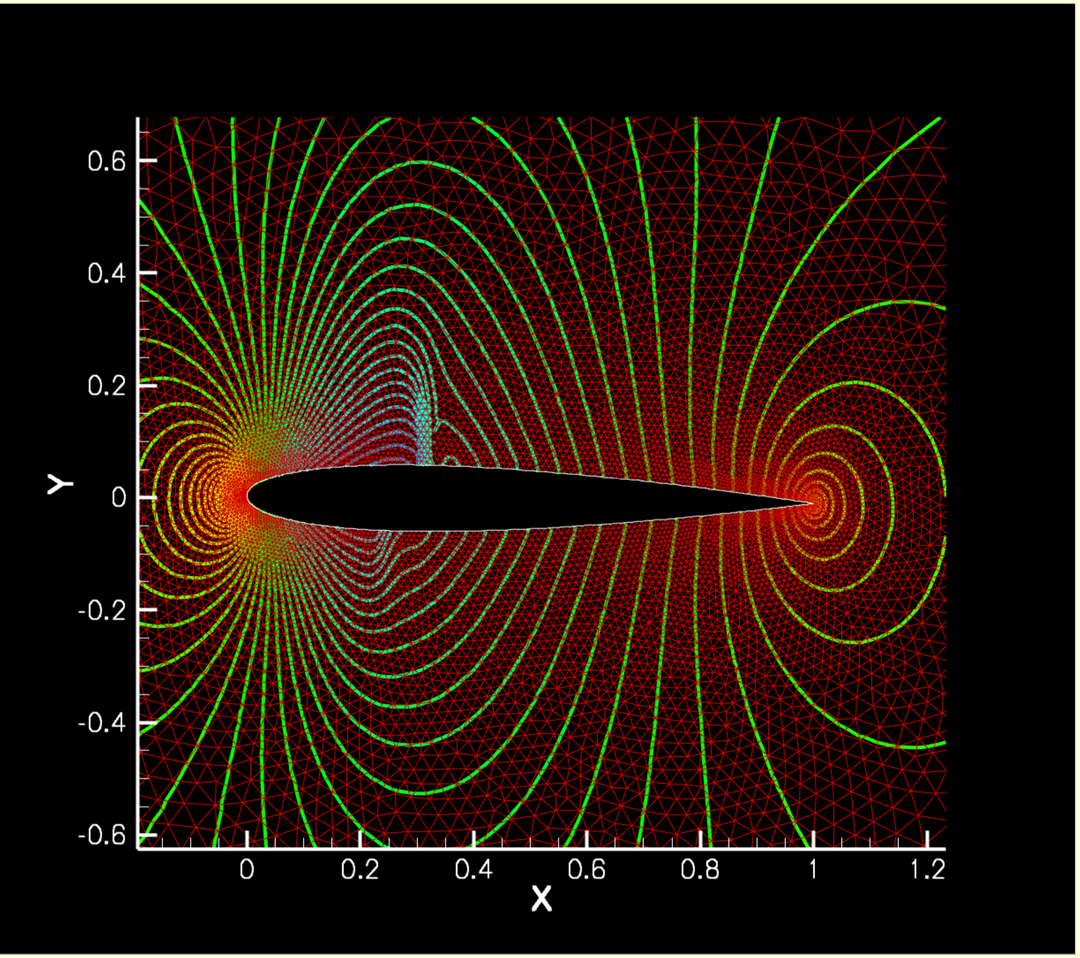

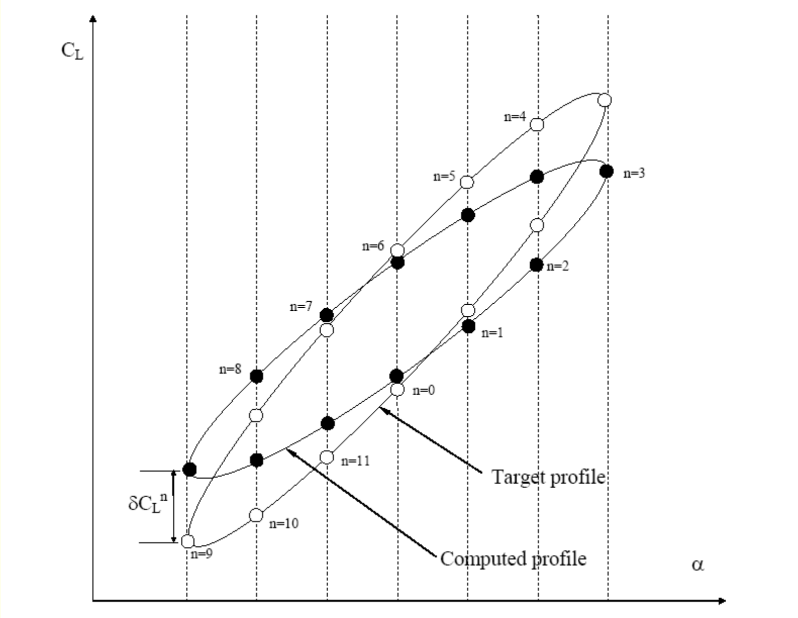

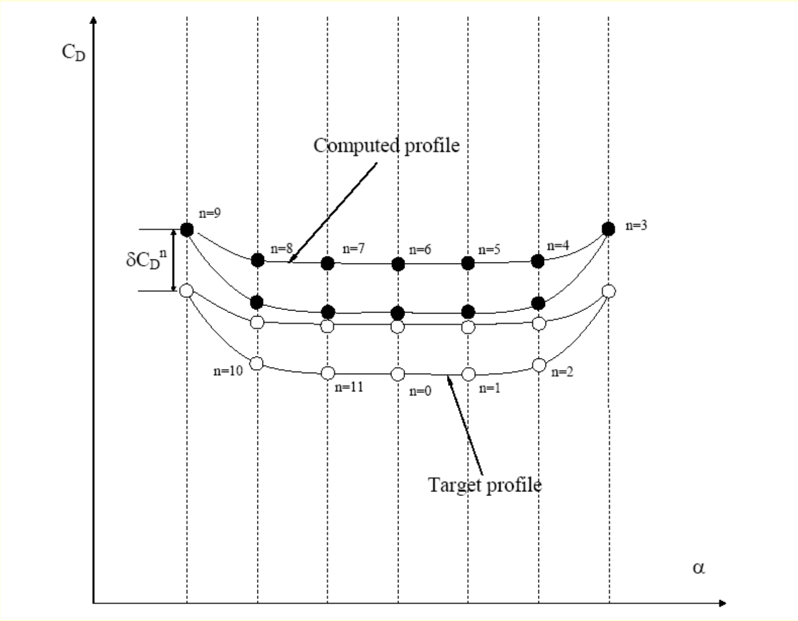

K. Mani and D. J. Mavriplis. *An unsteady discrete adjoint formulation for two-dimensional flow problems with deforming meshes*.
AIAA Paper 2007-0060

Time-Dependent Load Convergence/Comparison

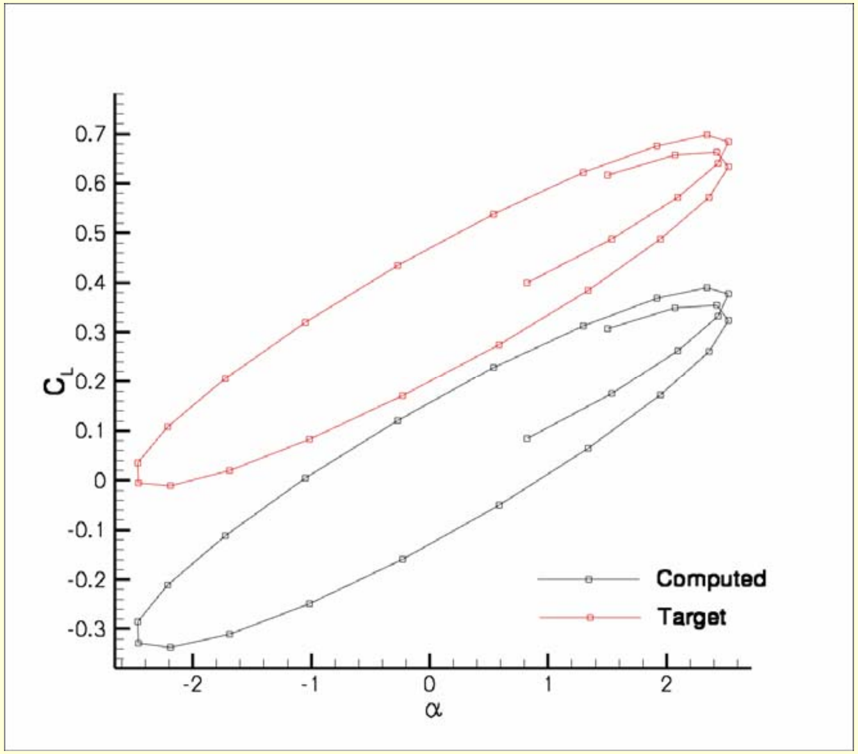

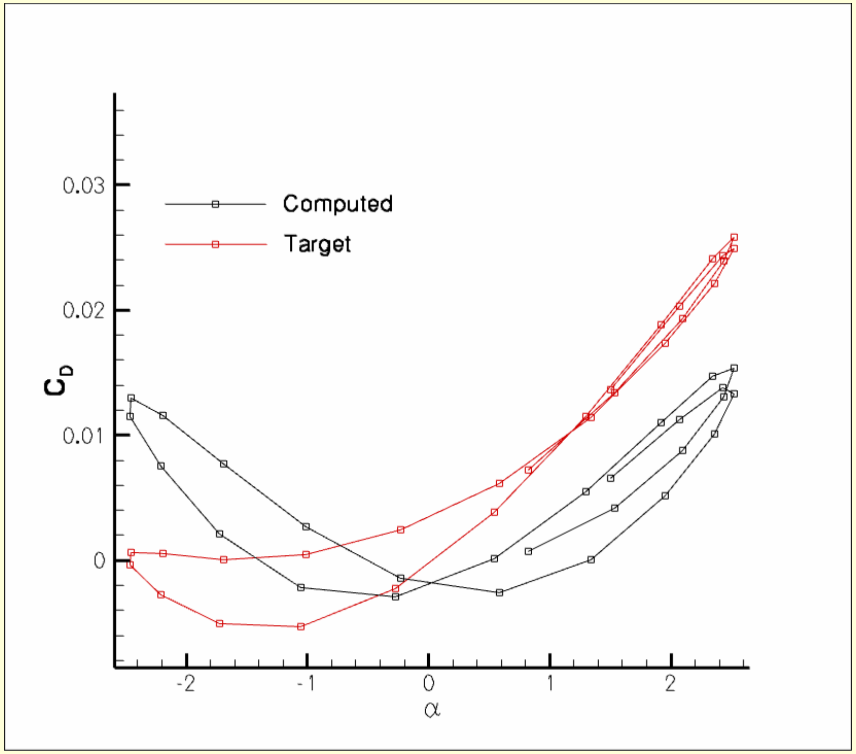

Shape Comparison/Convergence Exaggerated Scale

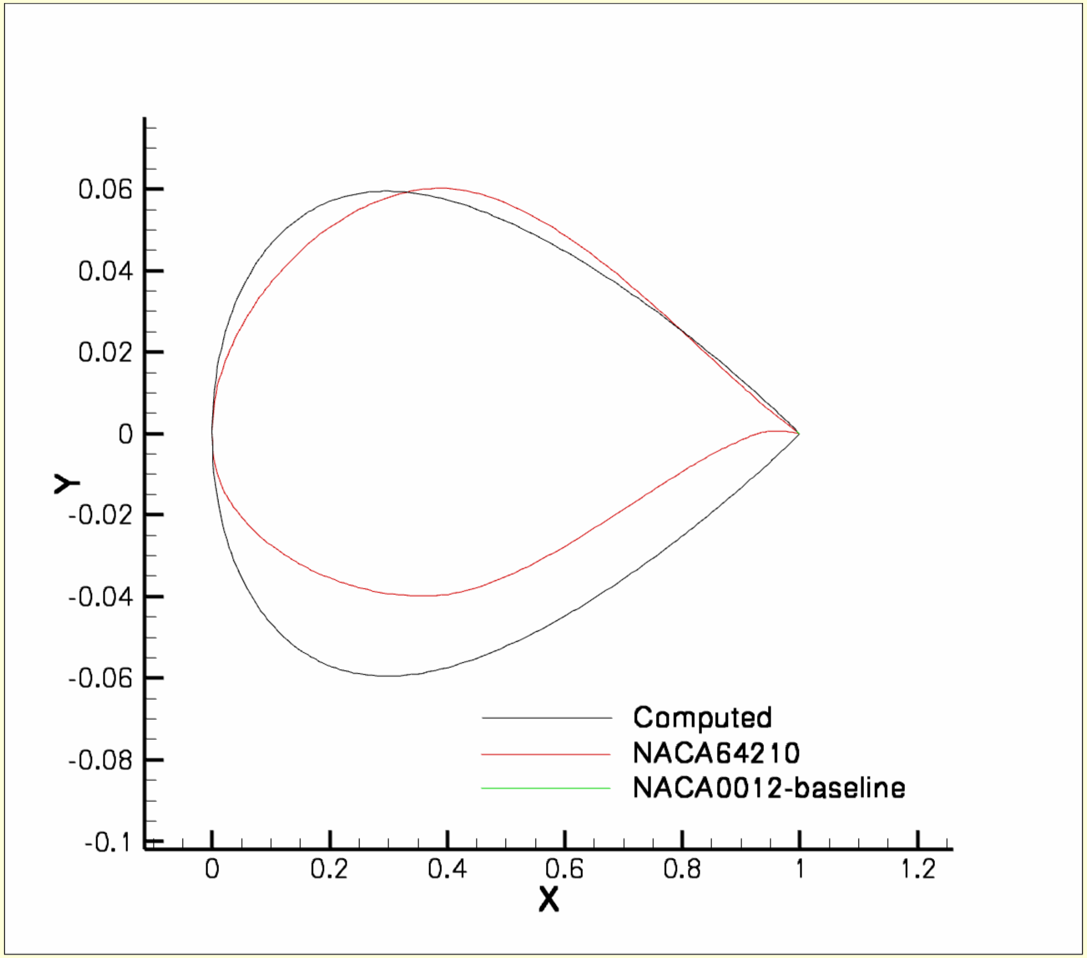

Objective Convergence

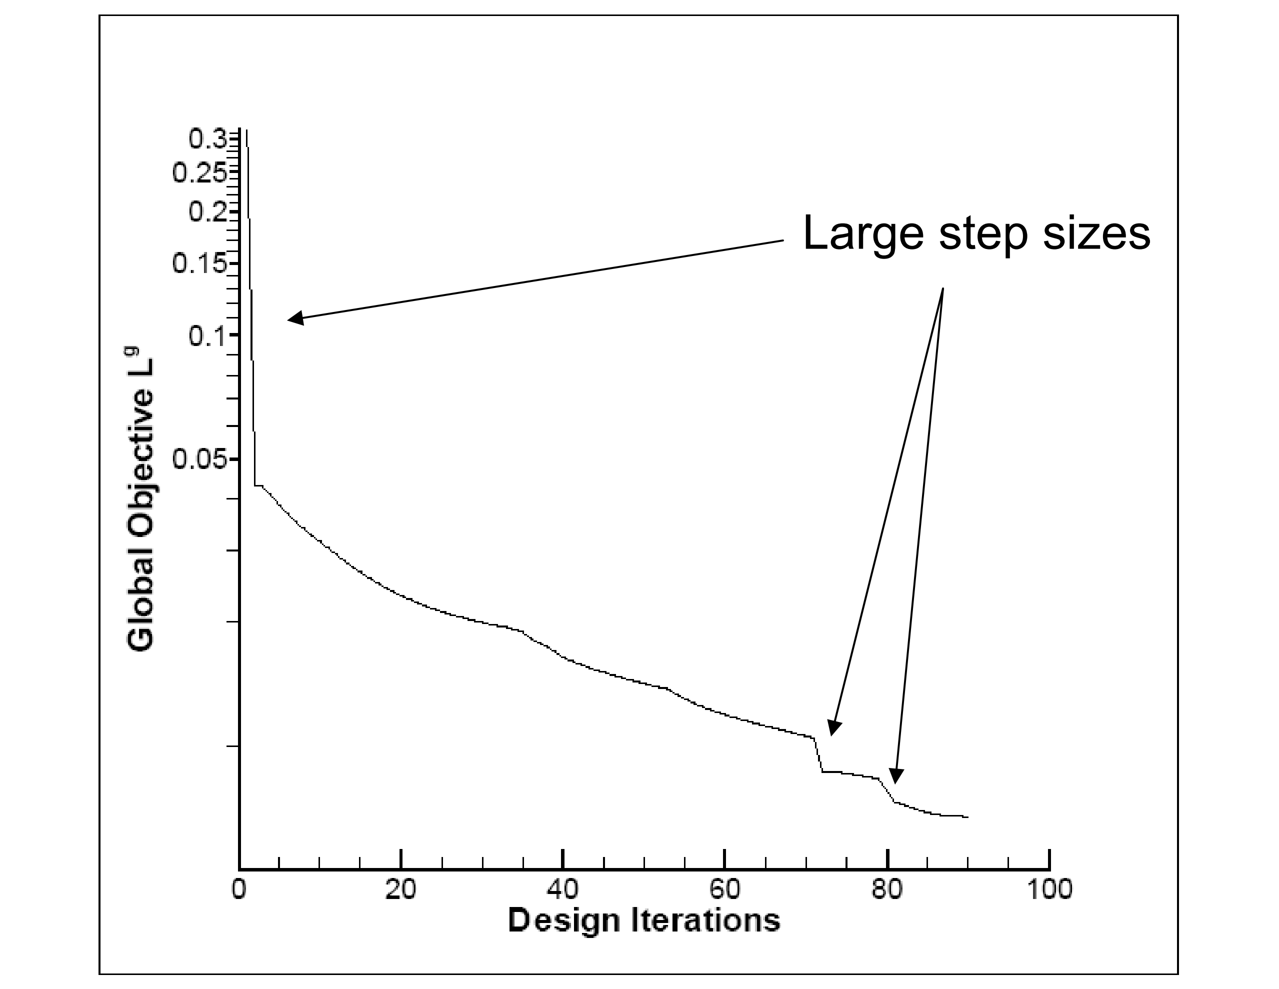

Computational Requirements

- One analysis cycle
  - 100 million grid points, one revolution
  - ~30 hours on 100 cpus
- One design cycle (twice cost of analysis)
  - Forward time dependent simulation
  - Backwards time dependent adjoint solution
- 50 to 100 design cycles
- 30 to 60 hours on 10,000 cpus

Design under Uncertainty

- Most analysis and design assumes problems are deterministic
- Aerospace vehicles are subjected to variability, e.g.
  - Manufacturing
  - Wear
  - Operating & environmental conditions
- NASA has been a leader in the development and application of probabilistic methods in structural applications (e.g. NASA-developed NESSUS for stochastic structural analysis)
- Common applications include turbine disks, high cycle fatigue of rotors, high-temperature material lifing, etc.

Non-deterministic Aerodynamics

- Few applications of non-deterministic methods involve high-fidelity aerodynamics
- Issues again center on:
  - Cost of high-fidelity models
  - Lack of robustness for high-fidelity models

Impact of Manufacturing Variability on Compressor Aerodynamic Performance

- Garzon & Darmofal studied the impact of geometry variability due to manufacturing on compressor aero performance (2002)

Compressor efficiency

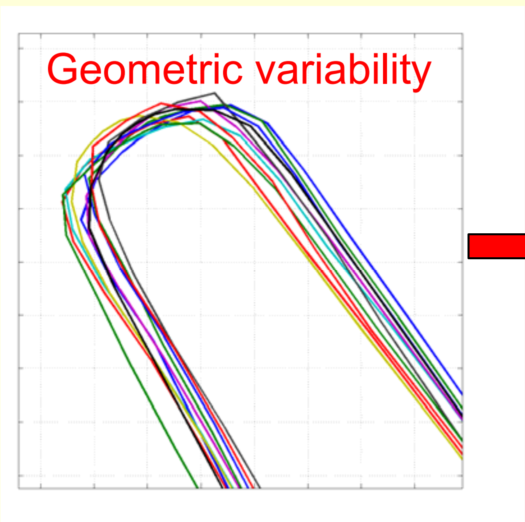

Probabilistic Aerodynamic Analysis

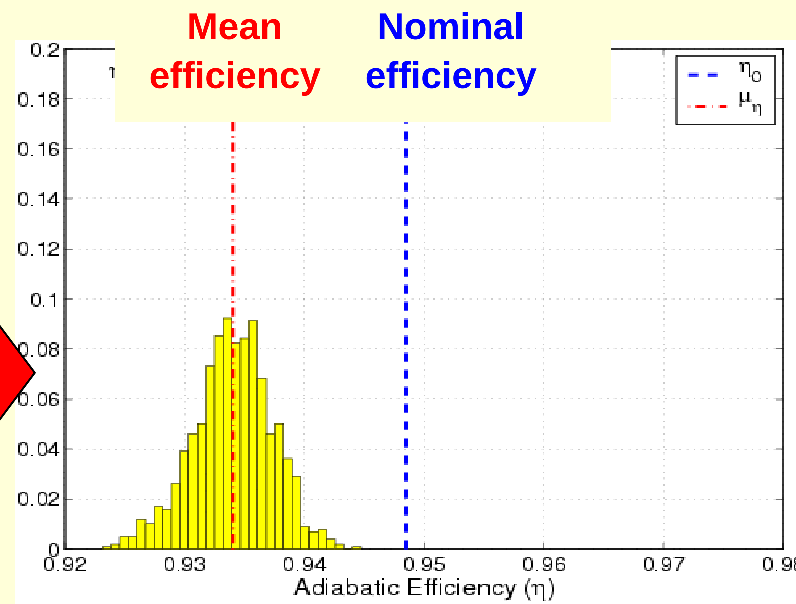

- 2-D coupled Euler-boundary layer model (about 5 seconds per run) used for a 2000 run Monte Carlo (about 3 CPU hours)

- Mean efficiency ~1.5 points lower than nominal
- 0% probability of compressor achieving design efficiency

Grand Challenge: Probabilistic Design of Cooled Turbine Blades

- Temperature-related damage of turbine blades is a leading cause of unscheduled engine repairs
- An increase of 20C decreases life by 50%
- Life estimation for cooled turbine blade requires multiple disciplines:
  - Main gaspath flow
  - Cooling passage flow
  - Structural dynamics
  - Heat transfer
- High fidelity modeling important because failure is often due to localized features (e.g. hot spots)

Probabilistic Design of Cooled Turbine Blades: Cost Estimate

- Burdet & Abhari (2007) estimate that fluid mesh for combined external and internal flows would be between 50-100 million points for existing methods (single stator-rotor stage)
- Unsteady flow effects due to upstream wakes are important (Hildebrandt et al 2006)
- Assume:
  - 1,000 run Monte Carlo
  - 25,000 CPU hours per run
  - 10,000 CPU's

→ 2,500 Wall Clock Hours

NASA Computational Environment

- Columbia processes mostly O(100) cpu jobs
- 2048 sub-system occupied with 512 jobs
- Few benchmarks above 512 cpus
- Some 2048 benchmarks (production ?)

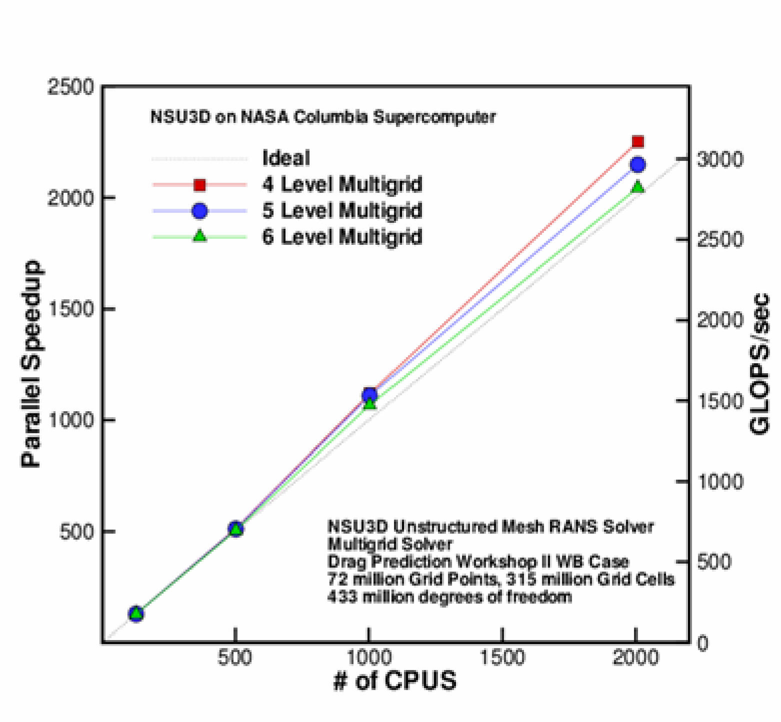

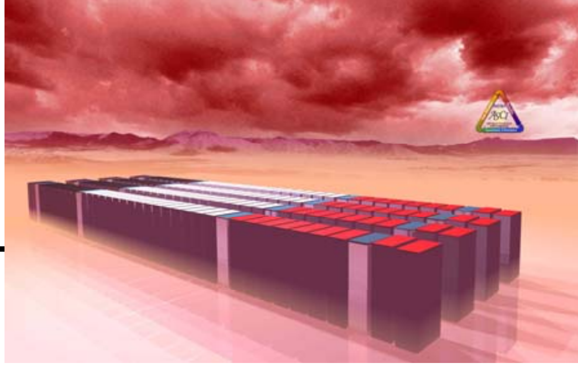

Science Runs on Red Storm

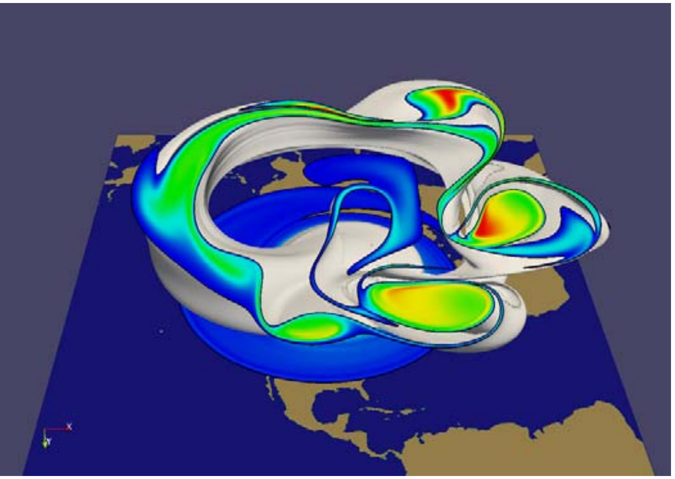

- SEAM (Spectral Element global Atmospheric Model) Simulation of the breakdown of the polar vortex used to study the enhanced transport of polar trapped air to mid latitudes.
- Record setting 20 day simulation, 7200 cpus for 36 hours. 1B grid points (3000x1500x300), 300K timesteps, 1TB of output.
- Spectral elements replace spherical harmonics in horizontal directions
- High order (p=8) finite element method with efficient Gauss-Lobatto quadrature used to invert the mass matrix.
- Two dimensional domain decomposition leads to excellent parallel performance.

c/o Mark Taylor, Sandia National Laboratories

Sandia National Laboratories

SEAM on Red Storm and BG/L

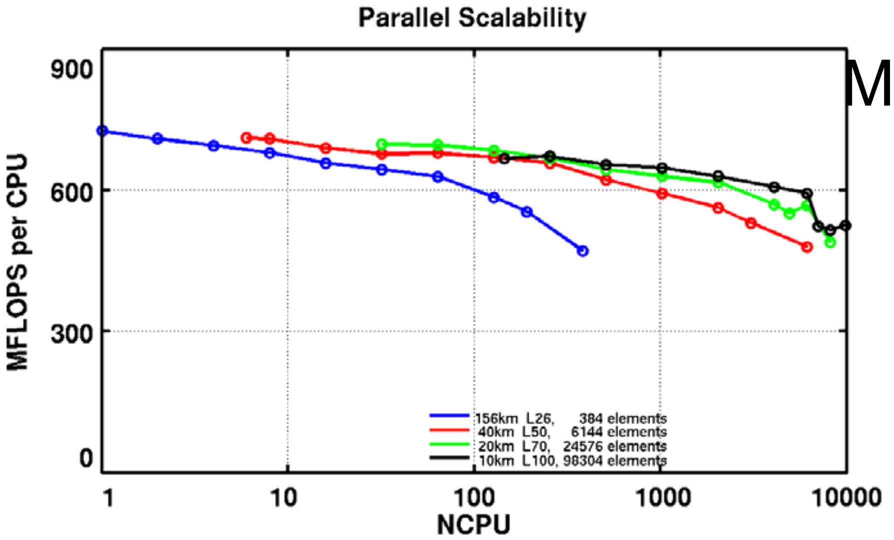

Max: 5TF

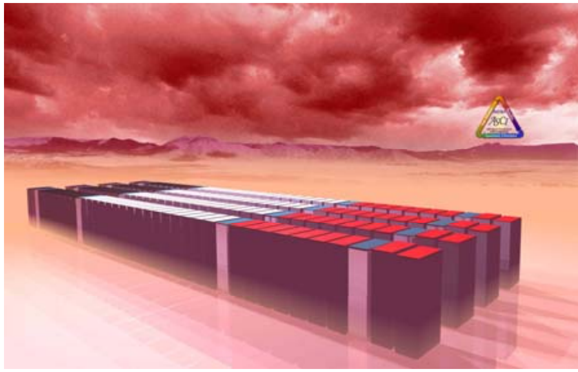

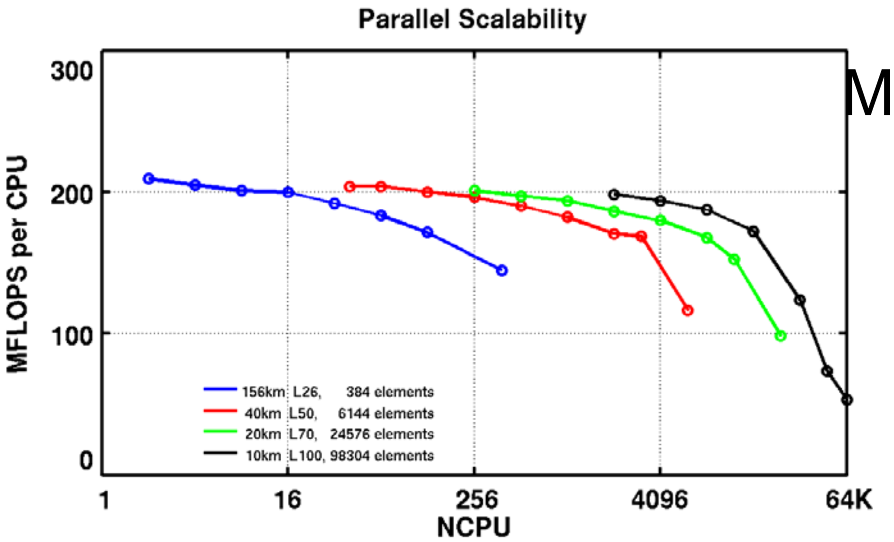

Max: 4TF

Performance of 4 fixed problem sizes, on up to 6K CPUs. The annotation gives the mean grid spacing at the equator (in km) and the number of vertical levels used for each problem.

Other Sample Science Simulations

- Magnetically Confined Fusion:
  - Tokomak core turbulence: 3.3 Tflops on 6,400 cpus Cray XT3 at ORNL: 70 hour runs
- Molecular Dynamics:
  - Solidification of metals using 0.5 trillion atoms
  - 100 TFlops on 131,072 cpus of IBM Blue Gene at LLNL: 7 hour runs
- These types of simulations are considered intractable within NASA aeronautics and most engineering communities
  - Some of the previous Grand Challenges are of this scale and could be done today

Barriers: Massive Parallelism

- All Hardware trends point to large increases in parallelism ~ 100,000 cpus or cores, maybe 1 Million
- 1992: HPCCP took us from O(8) vector cpus to O(100) cache-based cpus
- Today: Need to go from O(100) cpus to $O(10^5)$ cpus. WHY?
  1. Advance state-of-art through leading-edge simulations
  2. Future mid-size machines will have $O(10^4)$ cores or more
     - 2048 cpu IBM Blue Gene < $1M today
     - Cheap 16 core nodes available today

Massive Parallelism

- Explosive growth in parallelism is coming fast and needs to be met head on
  - Will require investment in scalable solvers research and deployment
  - Will require availability of massively parallel architectures for developing/testing solvers
  - Easy access to massively parallel architectures required to stimulate need
    - Restrict capacity use (small jobs)

Algorithm Development

- NASA ARMD research portfolio geared to programmatic interests in subsonic, rotorcraft, supersonic, and hypersonic applications
- Foundational algorithm developments needed in all areas
- Programmatic focus on realized deliverables (e.g. CFD software) not a good fit to foundational research over a 2-3 year period
- Many now-standard computational techniques were developed through foundational research funded by NASA during the 70's and 80's
- These foundational results were then often transitioned to applications by NASA personnel

Algorithm Development Opportunities

- Modest investment in cross-cutting algorithmic work would complement mission driven work and ensure continual long-term progress (including NASA expertise for determining successful future technologies)
  - Scalable non-linear solvers
  - Higher-order and adaptive methods for unstructured meshes
  - Optimization (especially for unsteady problems)
  - Reduced-order modeling
  - Uncertainty quantification

Other Important Areas

- Physical modeling
  - High fidelity, well resolved simulations can fail to provide useful engineering results due to failure of critical physical model
    - e.g. transition in hypersonics for heating rates
- Supporting experiments
  - NASA has rich history of experimental support
  - New (third) type of experiment: Numerical validation
- Software issues
  - Software complexity becoming important bottleneck
  - DOE has invested considerable effort in this area
    - http://www.the-diminishing-return-of-adding-personnel-to-a-task
- Educational issues
  - Fellowships, summer collaborations
  - ISCR at Livermore: 64 summer students, 8 academic collaborations

Conclusions

- NASA aeronautics has seen more than its share of cuts and reprogramming over the last decade
- Numerous studies have been published in support of increased aeronautics funding
  - These have always concentrated on the impact of NASA aero on
    - Aerospace industry
    - National air transportation system

Conclusions

- Aeronautics HPC impact and role much broader
  - Traditionally a driver for engineering simulation
  - Similar to DOE Office of Science:
    - Broad support for national science research
  - DOE/NASA complement other more academic agencies (NSF) in Science/Engineering
- This viewpoint requires NASA Aero to participate in national HPC initiatives
  - Engineering HPC requirements need to be voiced
  - Reformulated NASA Fundamental aero well positioned to be this voice

Conclusions

- Other communities have spent great effort to formulate the case for increased HPC investment
  - DOE SciDAC:
    - Scientific Simulation Initiative
    - Advanced Scientific Computing (1998)
    - SciDAC Report (2000)
    - Science Based Case for Large Scale Simulation (ScaLES: 2003, 2004)
  - Petascale Collaboratory for the GeoSciences (2006)
  - NSF Office of Cyber Infrastructure
    - 62 testimonials, 700 survey responses, Panel of 9

Conclusion

- We have provided an example of how this may be done for NASA aeronautics
- Similar efforts should be considered for formulating the case for increased investment in HPC for NASA Aeronautics
  - Impact on NASA's own missions (exploration)
  - Impact on aerospace industry
  - Impact on national engineering community
- Need to be an active participant in the discussion of HPC nationally