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

- Two principal intertwined themes
  - 1: NASA simulation capability risks becoming commoditized
    - Rapid advance of parallelism (> 1M cores)
    - Fundamental improvements in algorithms and development tools not keeping pace
    - Hardware and software complexity outstripping our ability to simulate (J. Alonso)
    - Clear vision of enabling possibilities is required
      - What would you do with 1000 times more computational power ?
  - 2: HPC Resurgent at National Level : Competitiveness
    - Aerospace industry is at the heart of national competitiveness
    - NASA is at the heart of aerospace industry
    - Can NASA leverage/benefit from renewed national HPC emphasis ?

ARMD’s Historic HPC Leadership
(Code R)

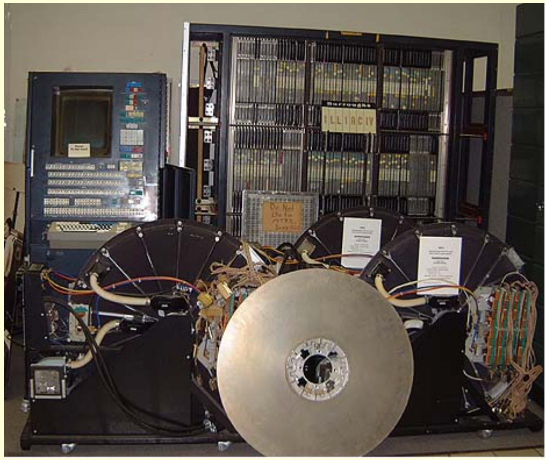

- ILLIAC IV (1976)
- National Aerodynamic Simulator (1980’s)
- 1992 HPCCP Budget:
  - $596M (Total)
    - $93M Department of Energy (DOE)
    - $71M NASA
      - Earth and Space Sciences (ESS)
      - Computational Aerosciences (CAS)

- Computational Aerosciences (CAS) Objectives (1992):
  - “…*integrated, multi-disciplinary simulations and design optimization of aerospace vehicles throughout their mission profiles*”
  - “… *develop algorithm and architectural testbeds … scalable to sustained teraflops performance*”

HPC Today at NASA

- NASA Columbia Supercluster:
  - 10,240 cpus
- Mostly used as capacity (not capability) facility
  - Many “small” jobs of order O(100) cpus
  - Incremental progress since 1990’s
    - Published NASA code benchmarks stop at 512 cpus
    - 512 cpu runs on Intel Touchstone Delta Machine (ICASE/NASA Langley at Supercomputing ’92)
    - Supercomputing’05: Only 1 NASA Paper
  - Not demonstrating new aerospace engineering frontiers to be opened by rapid increases in computational power

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

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
    - SGI ICE system
  - Easy access to massively parallel architectures required to stimulate need
    - Restrict capacity use (small jobs)

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
  - Geometry management
- Current simulation capabilities (NASA/DOE/others) rests on algorithmic developments, many funded by NASA

Science Runs on RedStorm

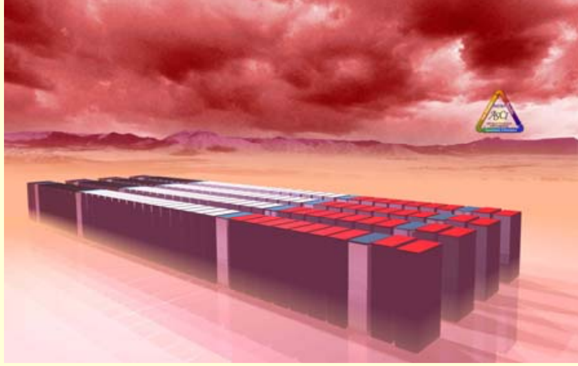

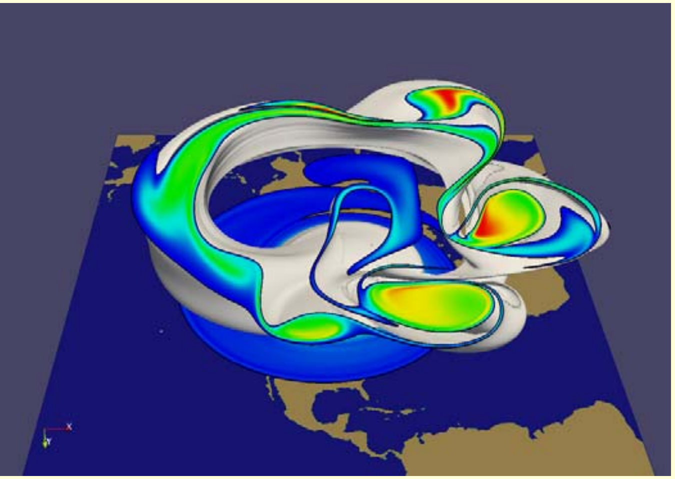

- SEAM (Spectral Element global Atmospheric Model) Simulation of the breakdown of the polar vortex used to study the enhanced transport of polar trapped air to mid latitudes.
- Record setting 20 day simulation, 7200 cpus for 36 hours. 1B grid points (3000x1500x300), 300K timesteps, 1TB of output.
- Benchmarks up to 64K cpus
- 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

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 stated Grand Challenges are of this scale and could be done today

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)….

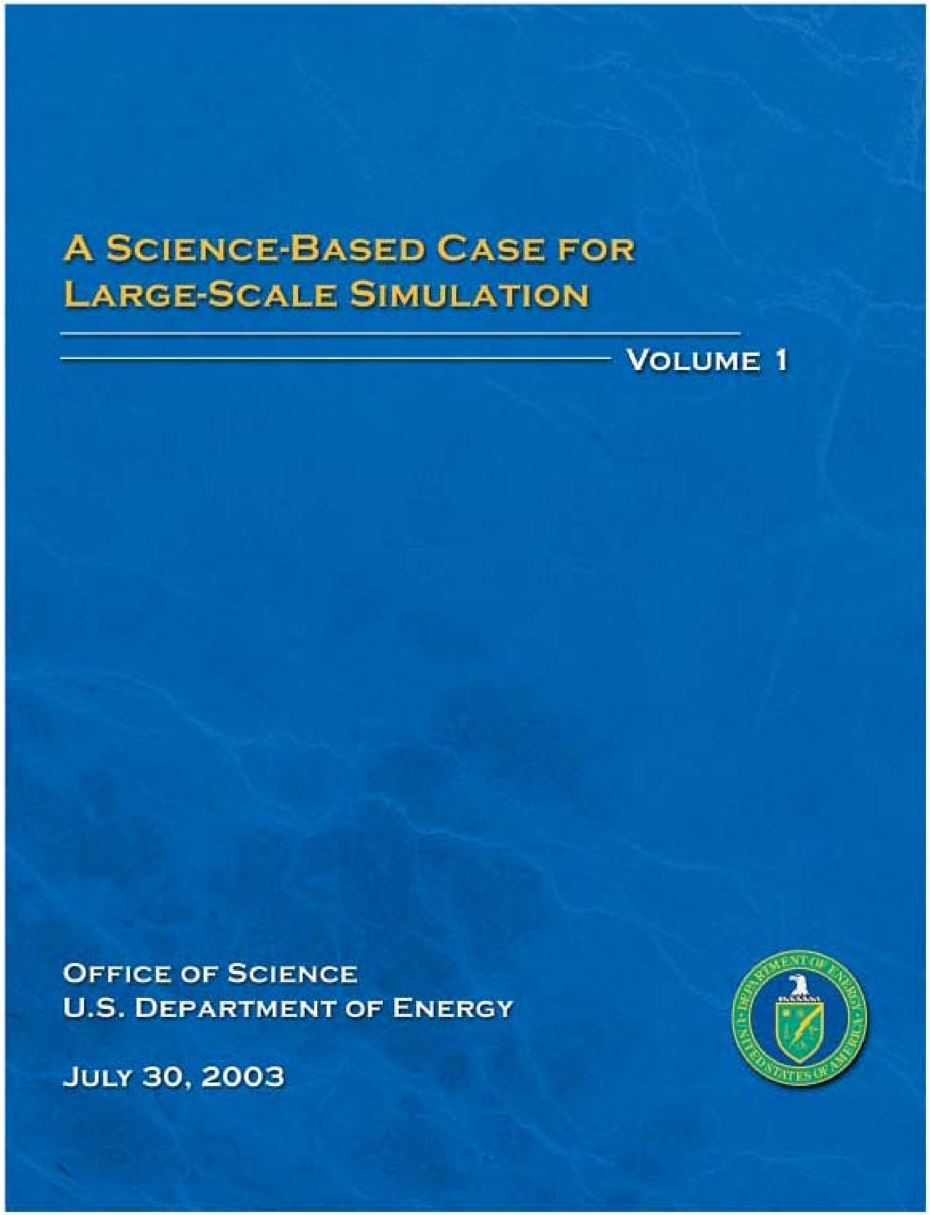
A SCIENCE-BASED CASE FOR
LARGE-SCALE SIMULATION
VOLUME 1
OFFICE OF SCIENCE
U.S. DEPARTMENT OF ENERGY
JULY 30, 2003

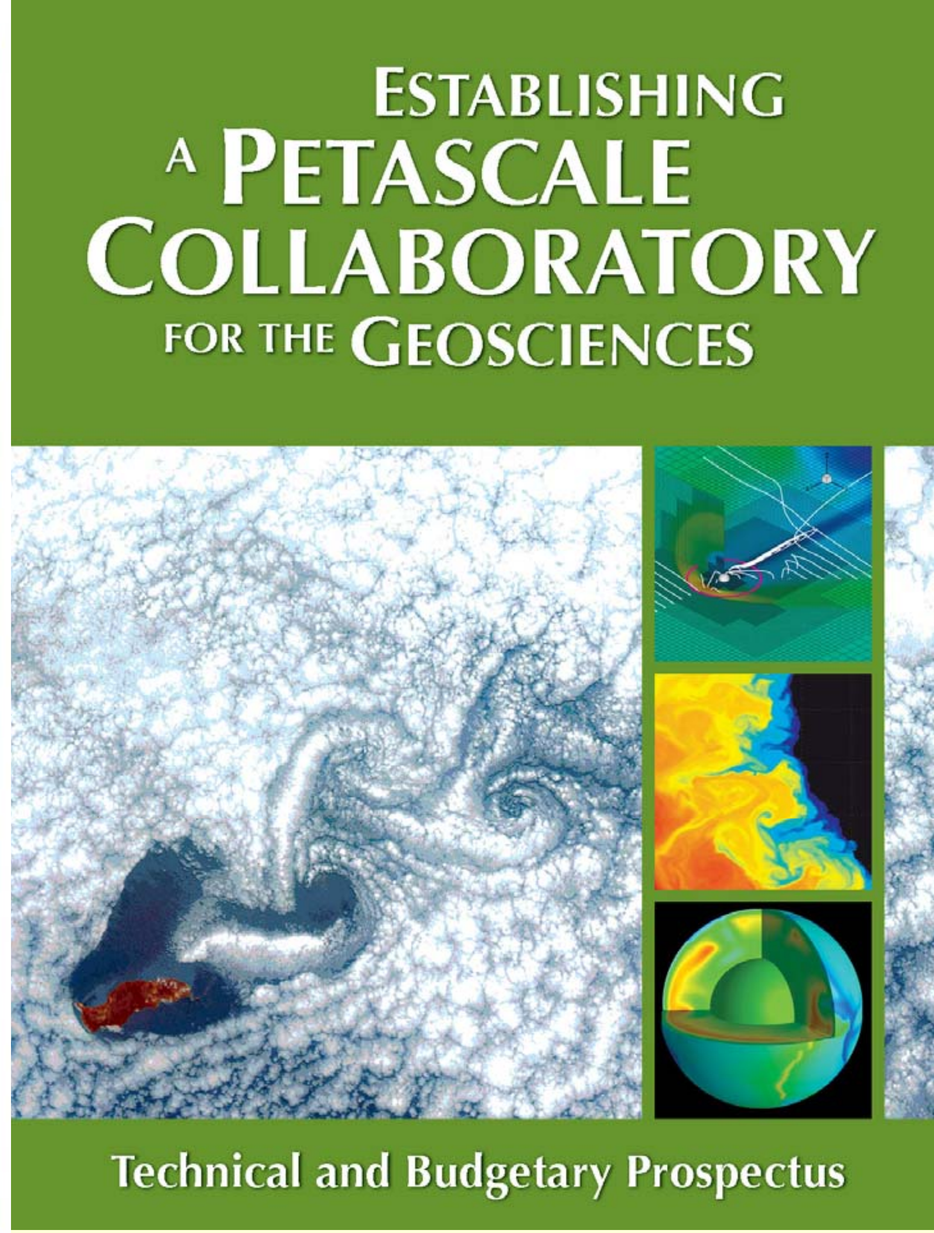
ESTABLISHING
A PETASCALE
COLLABORATORY
FOR THE GEOSCIENCES
Technical and Budgetary Prospectus

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)

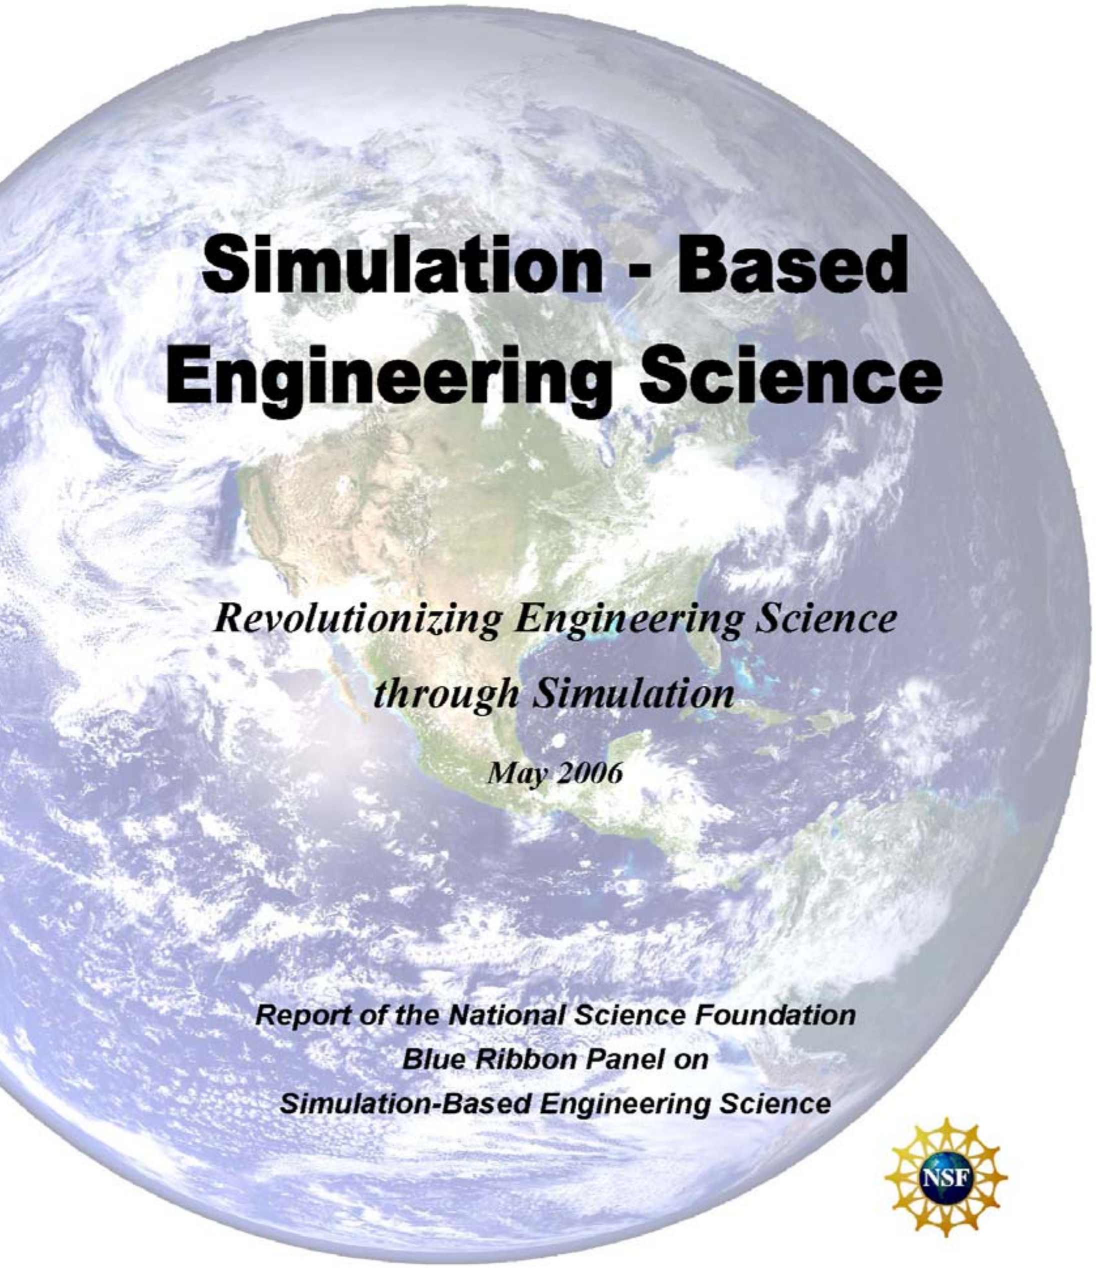

- Recent NSF Report
  - Engineering based simulation needs more attention
    - Science has been successful recently as advocate
  - Mainly structures, crash dynamics, materials
  - No mention of aeronautics activities

National HPC and Aeronautics

- 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/DOD not players

- Isn't Engineering as important (or more) than Science with respect to National Competitiveness ?

Aeronautics and HPC

- NASA Aeronautics has traditionally been at interface of HPC research and leading-edge engineering applications (more so than NSF report examples)
  - Fundamental Algorithmic research
  - Real world applications and close industry collaboration
  - Tools developed by ARMD serve other NASA Missions
- Aeronautics HPC impact and role much broader than considerations in OSTP National Aeronautics Plan
  - Traditionally a driver for engineering simulation
  - Similar to DOE Office of Science:
    - Broad support for national science research
- 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
- We have provided an example of how this may be done for NASA Aeronautics