Future Directions in HPC and Grids

Numerical Relativity as Driver

US Status Report: personal perspective

Ed Seidel

Center for Computation & Technology

Louisiana State University, USA

Max-Planck-Institut für Gravitationsphysik

Potsdam, Germany

The Challenges

Comp. Science as "Third Pillar"

- Applications
- Data
- Communities

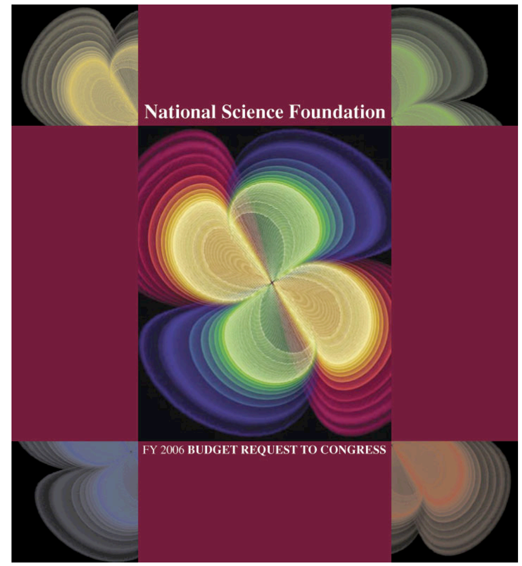

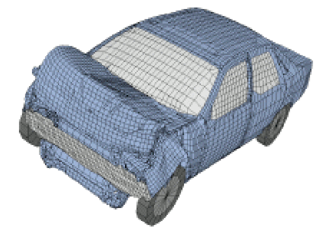

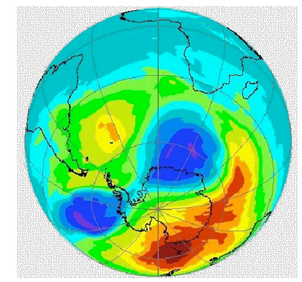

GWEN

Gravitational Wave European Network

From Astrophysical Theory to Detection and Understanding

Huge Black Hole Collision Simulation

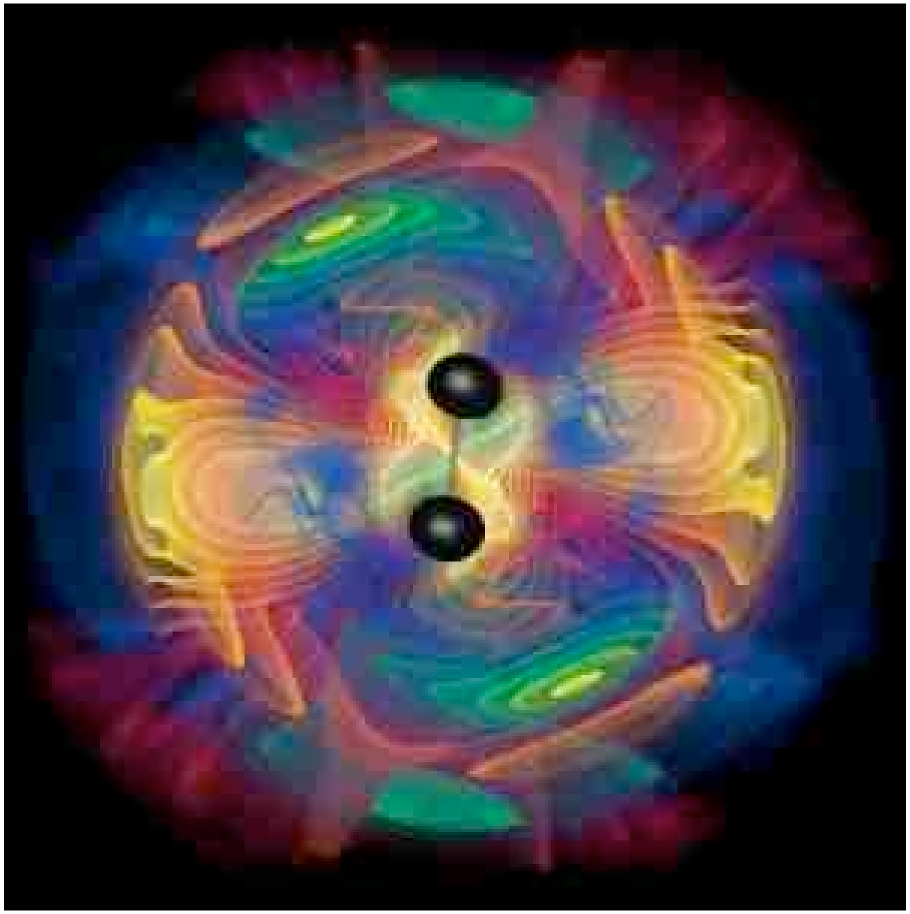

- Data Tsumani
  - Smarr: supercomputers as petabyte generators
    - Num. Rel.: 10^6 sims for GW param space, each 25TB
  - How to retrieve, track, manage?
- Many components developed by distributed collaborations
  - How to build codes for modern environments?
- Many computational resources available
  - How to find best ones to start?
  - How to distribute work? Interact w/code?
- Time-changing computations!
  - Multiscale, multiphysics: low efficiency
  - Adapt, monitor changes?
- Interact with experiments? Coming!

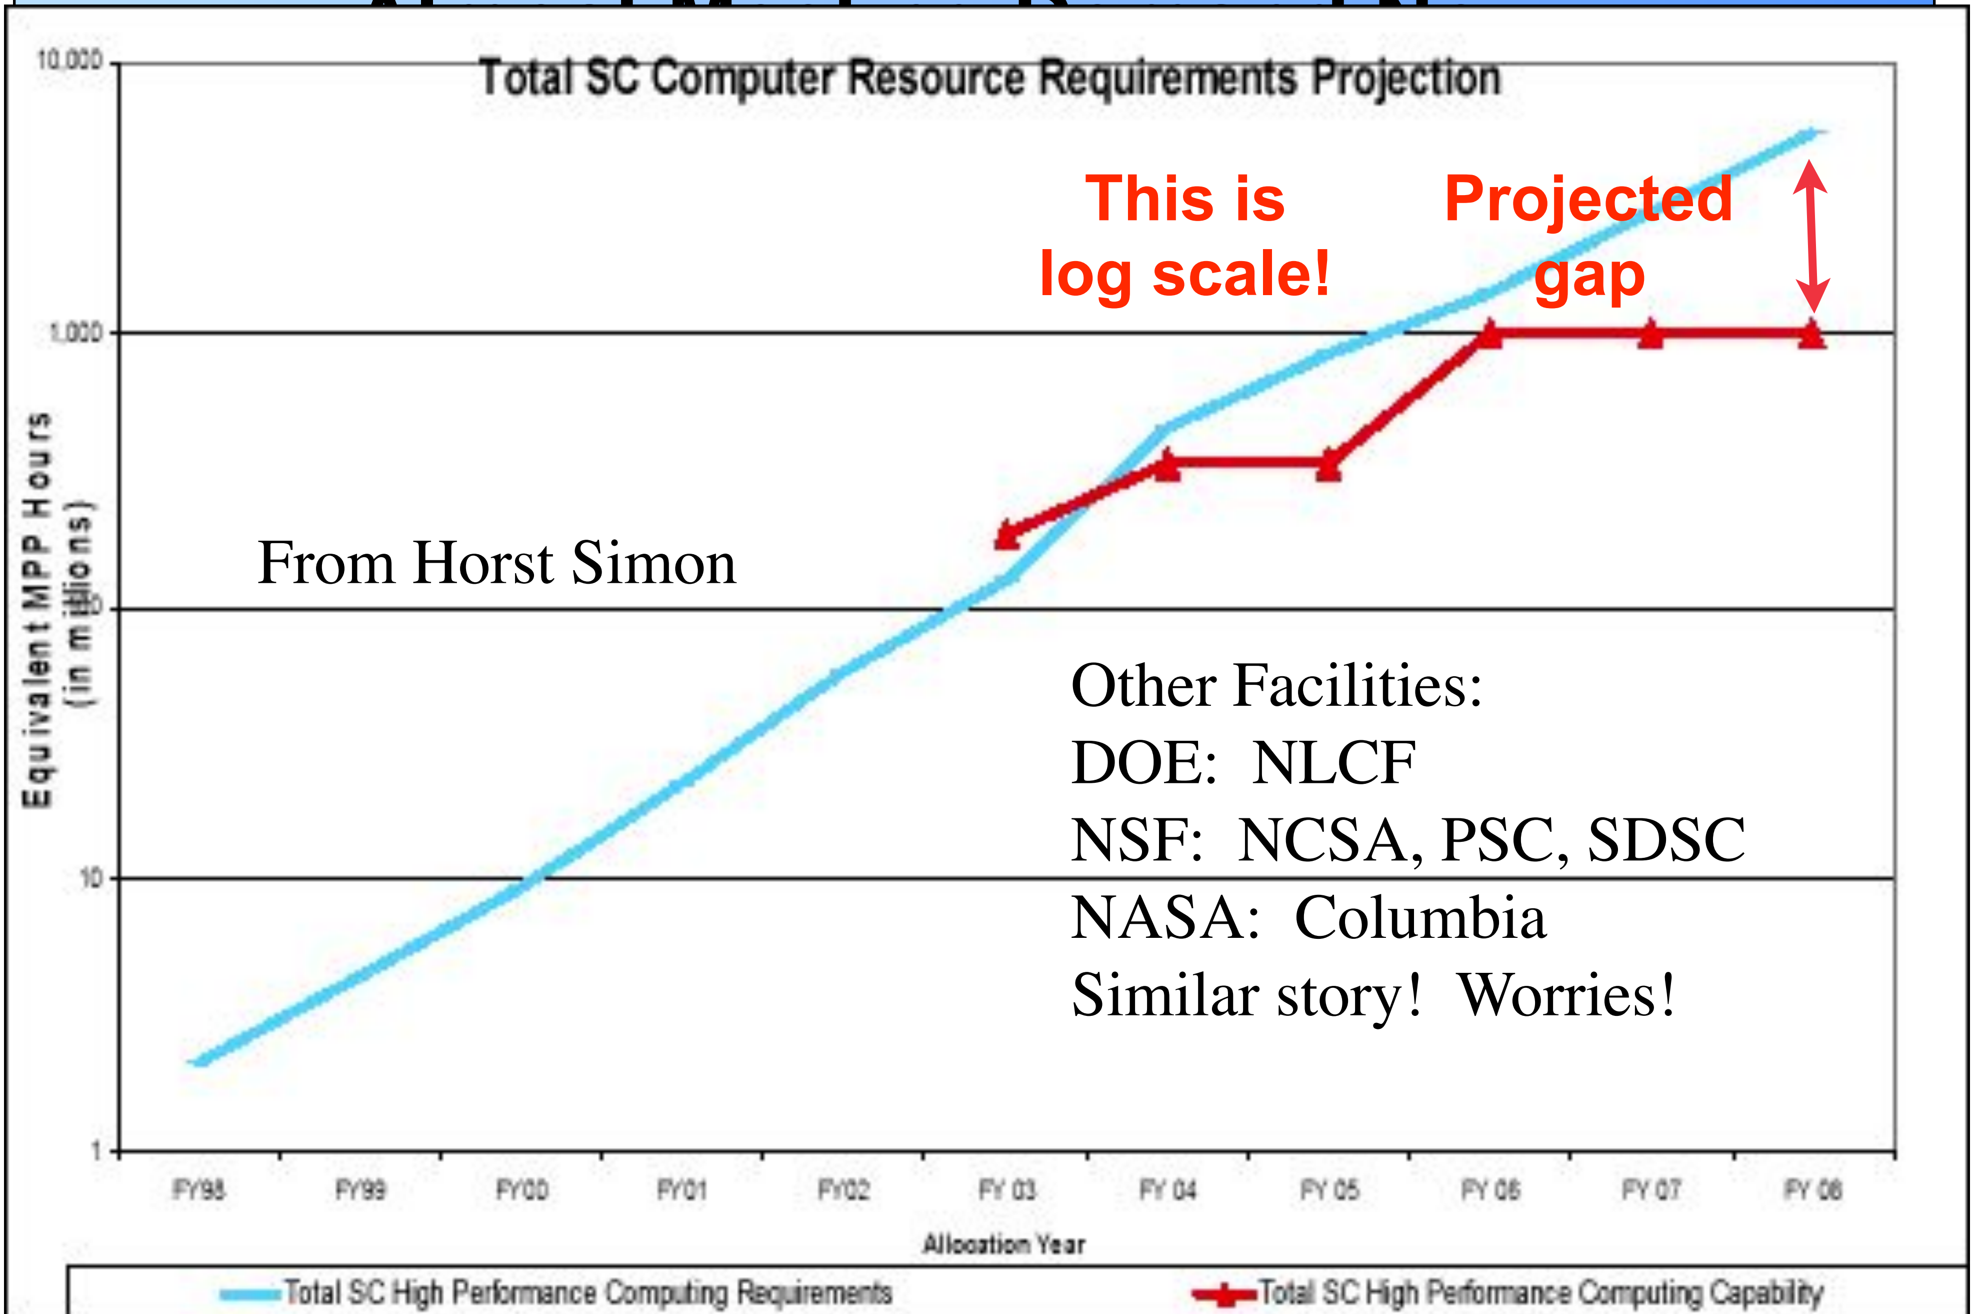

From Horst Simon

Other Facilities:
DOE: NLCF
NSF: NCSA, PSC, SDSC
NASA: Columbia
Similar story! Worries!

Sensors: NSF EarthScope Project

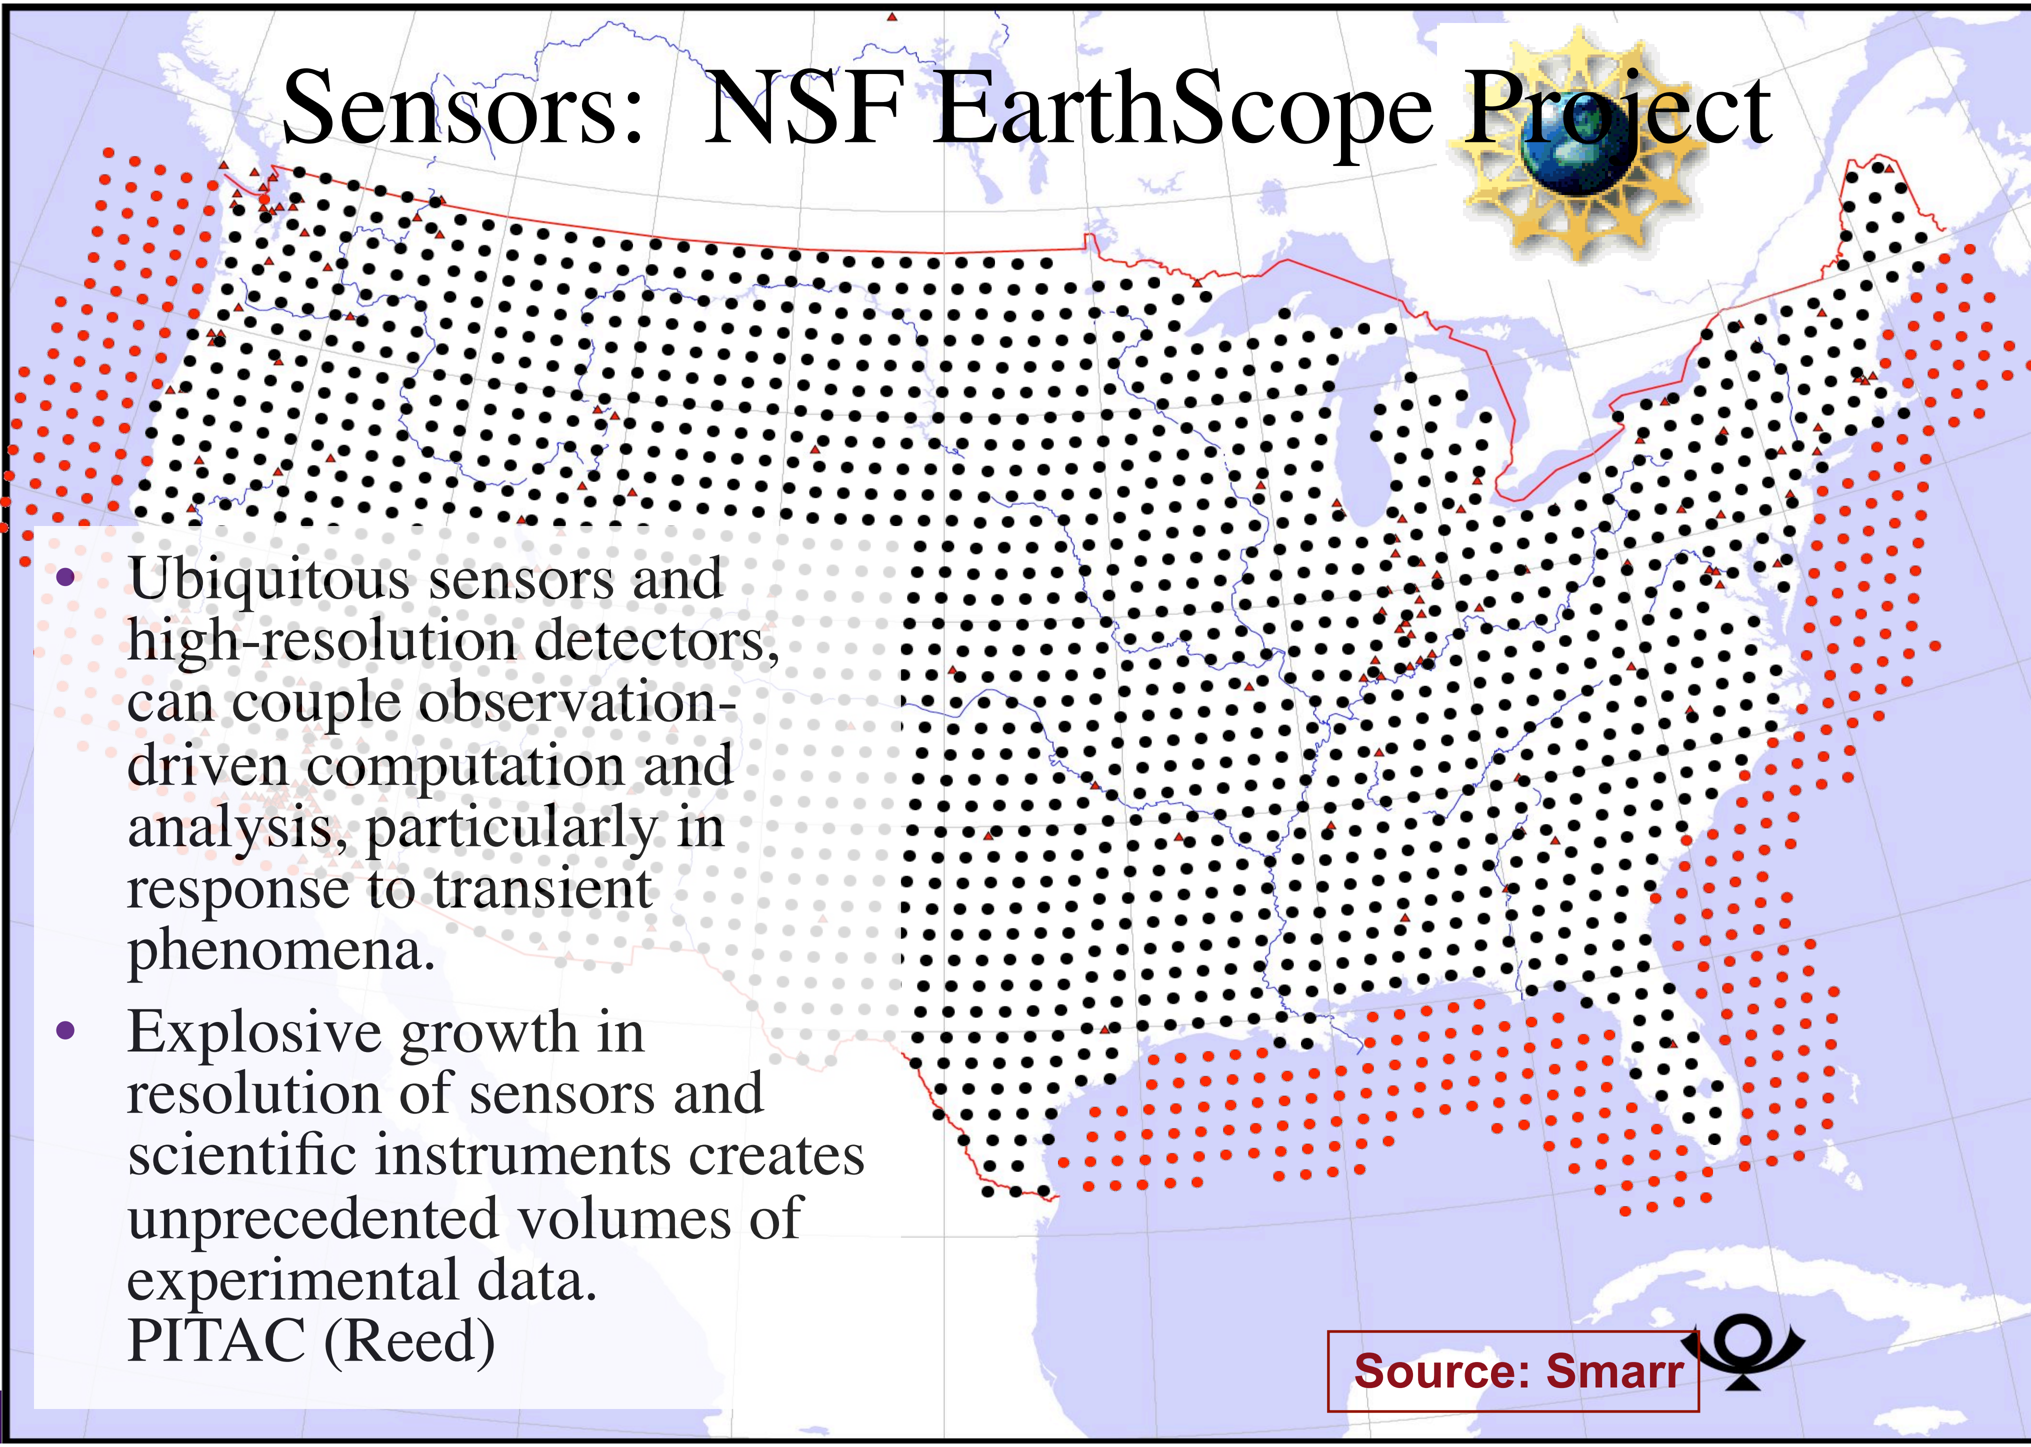

- Ubiquitous sensors and high-resolution detectors, can couple observation-driven computation and analysis, particularly in response to transient phenomena.
- Explosive growth in resolution of sensors and scientific instruments creates unprecedented volumes of experimental data.
PITAC (Reed)

Source: Smarr

cct Collaborations for Complex Problems

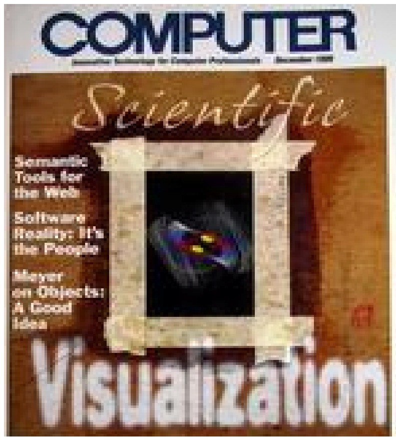

NASA Neutron Star Grand Challenge

- 5 US Institutions
- Attack colliding neutron star problem

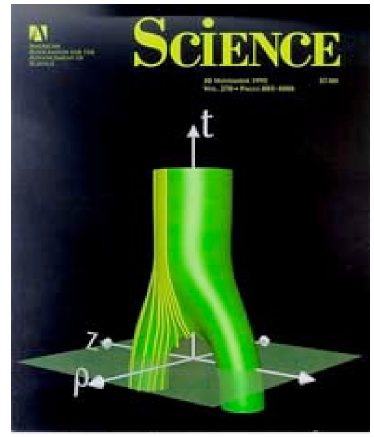

NSF Black Hole Grand Challenge

- 8 US Institutions
- 5 years
- Attack colliding black hole problem

EU Astrophysics Network, GWEN

- 10 EU Institutions
- 3 years
- Continue these problems

Examples of Future of Science & Engineering

- Require large scale data, simulations, beyond reach of existing supercomputers
- Require large geo-distributed cross-disciplinary collaborations
- Require Grid technologies

Grids are not a substitute for SC's!

Issues Meeting the Challenges

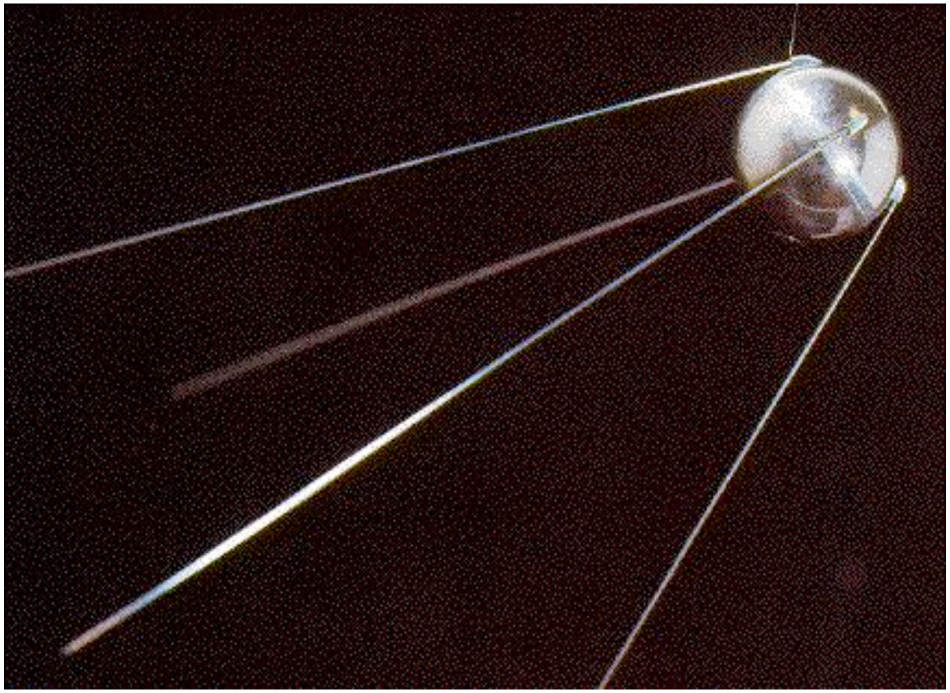

Sputnik: October 4, 1957

System trends towards Petascale

- Processors
  - Moore's Law slowing, typical efficiencies lower than ever.
  - Heat!! Memory Bandwidth!!
- System Taxonomy
  - Commodity
  - Hybrid
  - Custom

"Challenge: Prove efficiency of petascale system is bounded away from zero" Mike Heath

- Power consumption impact
  - Large systems using order MWatts of power.
  - Prevents dense packaging of HPC systems.
  - Adds facility costs
- Worries: we will need something new! Who is funding this? See "High End Crusader", HPCWire

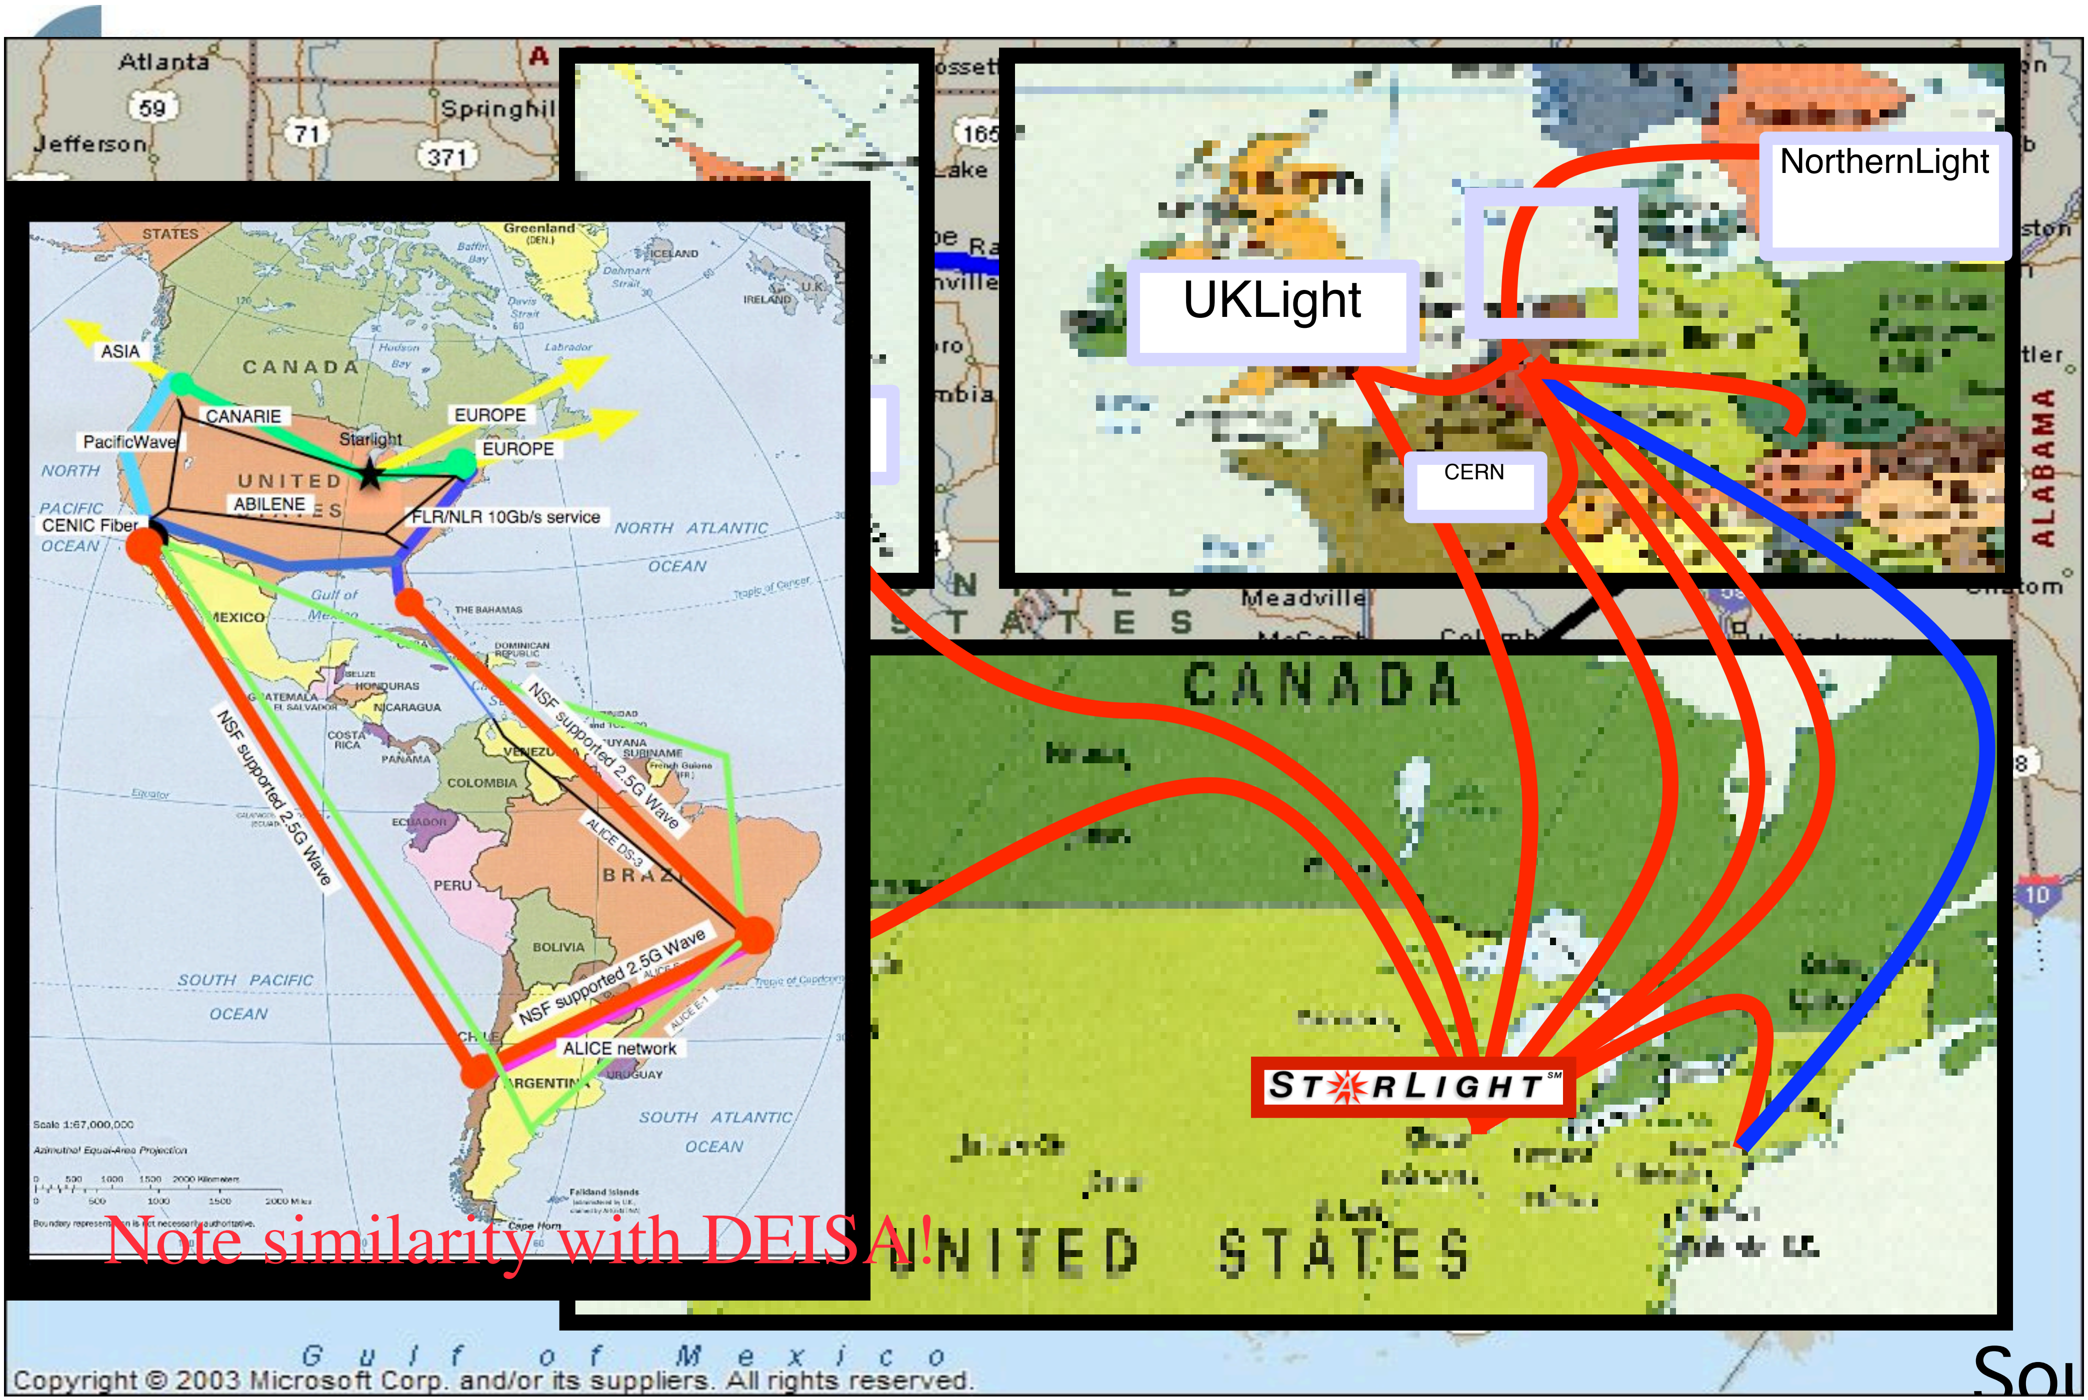
NorthernLight
UKLight
CERN
STARLIGHT
CANADA
UNITED STATES
ASIA
CANARIE
PacificWave
Starlight
EUROPE
EUROPE
ABILENE
CENIC Fiber
FLR/NLR 10Gb/s service
NORTH ATLANTIC OCEAN
NORTH PACIFIC OCEAN
NSF supported 2.5G Wave
NSF supported 2.5G Wave
NSF supported 2.5G Wave
ALICE network
SOUTH PACIFIC OCEAN
SOUTH ATLANTIC OCEAN
Note similarity with DEISA!
Gulf of Mexico
Copyright © 2003 Microsoft Corp. and/or its suppliers. All rights reserved.

Grids: Bringing all Together

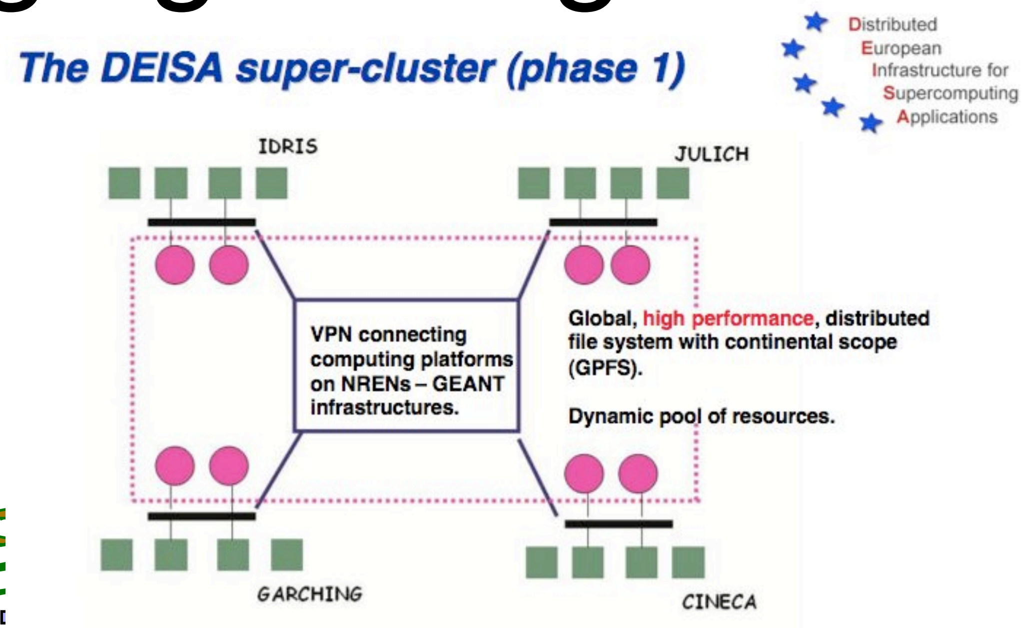

- Computational Devices Scattered Across the World
  - Compute servers (double 18 months)
  - Networks (double each 9 months)
  - Sensors (exploding field)

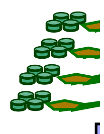

- How to take advantage of this for science, engineering, business, art?
  - Harness multiple sites and devices
  - Deploy cyberinfrastructure: Globus, Unicore, many others

- Worries in US!
  - DOE, NSF cyberinfrastructure funding

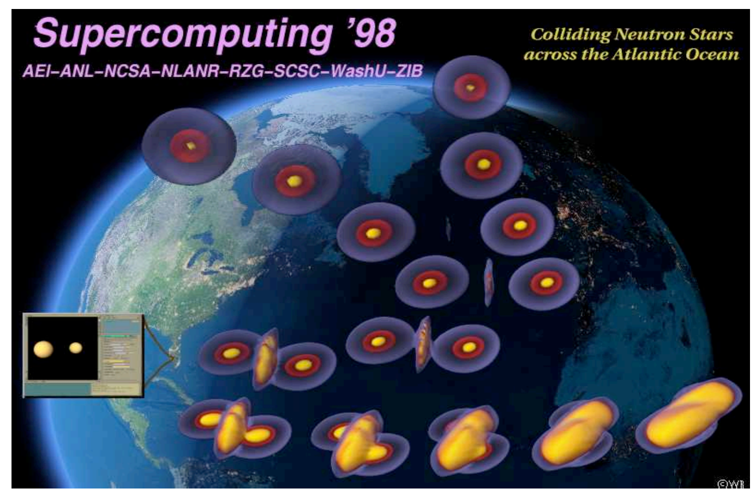

Gordon Bell Prize, 2001

Computational Science Software

- There is a *crisis* in computational science software
  - many years of inadequate investments
  - lack of useful tools
    - Cactus, Samrai, Overture, etc desperately needed
  - dearth of widely accepted standards and best practices
  - paucity of third party software vendors
    - DARPA survey: only handful could possibly scale to petascale
- Improvement in computational science software is needed urgently along multiple dimensions
- Half-life
  - few years for machines
  - decades for software!!

Source: Dan Reed

IT Funding Crisis in US

- Computing growth path in doubt
- Academic IT funding at Darpa down 50% in 3 years
- NASA downsizing IT
- DOE, NIH facing cuts
- NSF in disarray
  - funds 86% of IT research
  - success rate drops below 15%

Science EDITORIAL

An Endless Frontier Postponed

Next month, U.S. scientists Vinton G. Cerf and Robert E. Kahn will receive computing's highest prize, the A. M. Turing Award, from the Association for Computing Machinery. Their Transmission Control Protocol (TCP), created in 1973, became the language of the Internet. Twenty years later, the Mosaic Web browser gave the Internet its public face. TCP and Mosaic illustrate the nature of computer science research, combining a quest for fundamental understanding with considerations of use. They also illustrate the essential role of government-sponsored university-based research in producing the ideas and people that drive innovation in information technology (IT).

Recent changes in the U.S. funding landscape have put this innovation pipeline at risk. The Defense Advanced Research Projects Agency (DARPA) funded TCP. The shock of the Soviet satellite Sputnik in 1957 led to the creation of the agency, which was charged with preventing future technological surprises. From its inception, DARPA funded long-term nonclassified IT research in academia, even during several wars, to leverage all the best minds. Much of this research was dual-use, with the results ultimately advancing military systems and spurring the IT industry.

U.S. IT research grew largely under DARPA and the National Science Foundation (NSF). NSF relied on peer review, whereas DARPA bet on vision and reputation, complementary approaches that served the nation well. Over the past 4 decades, the resulting research has laid the foundation for the modern microprocessor, the Internet, the graphical user interface, and single-user workstations. It has also launched new fields such as computational science. Virtually every aspect of IT that we rely on today bears the stamp of federally sponsored research. A 2003 National Academies study provided 19 examples where such work ultimately led to billion-dollar industries, an economic benefit that reaffirms science advisor Vannevar Bush's 1945 vision in *Science: The Endless Frontier*.

However, in the past 3 years, DARPA funding for IT research at universities has dropped by nearly half. Policy changes at the agency, including increased classification of research programs, increased restrictions on the participation of noncitizens, and "go/no-go" reviews applied to research at 12- to 18-month intervals, discourage participation by university researchers and signal a shift from pushing the leading edge to "bridging the gap" between fundamental research and deployable technologies. In essence, NSF is now relied on to support the long-term research needed to advance the IT field.

Other agencies have not stepped in. The Defense Science Board noted in a recent look at microchip research at the Department of Defense (DOD): "[DARPA's] withdrawal has created a vacuum . . . The problem, for DOD, the IT industry, and the nation as a whole, is that no effective leadership structure has been substituted." The Department of Homeland Security, according to a recent report from the President's Information Technology Advisory Committee, spends less than 2% of its Science and Technology budget on cybersecurity, and only a small fraction of that on research. NASA is downsizing computational science, and IT research budgets at the Department of Energy and the National Institutes of Health are slated for cuts in the president's fiscal year 2006 budget.

These changes, combined with the growth of the discipline, have placed a significant burden on NSF, which is now showing the strain. Last year, NSF supported 86% of federal obligations for fundamental research in IT at academic institutions. The funding rate for competitive awards in the IT sector fell to 16%, the lowest of any directorate. Such low success rates are harmful to the discipline and, ultimately, to the nation.*

At a time when global competitors are gaining the capacity and commitment to challenge U.S. high-tech leadership, this changed landscape threatens to derail the extraordinarily productive interplay of academia, government, and industry in IT. Given the importance of IT in enabling the new economy and in opening new areas of scientific discovery, we simply cannot afford to cede leadership. Wher

Where will the Turing Awardees 30 years hence r

"not in the United States."

Lazowska, Patterson

Recommendations

- Good advice for all from PITAC, others

Recommendation

The Federal government must rebalance research and development investments to:

a) *create a new generation of well-engineered, scalable, easy-to-use software* suitable for computational science that can reduce the complexity and time to solution for today's challenging scientific applications and can create accurate simulations that answer new questions;

b) *design, prototype, and evaluate new hardware architectures* that can deliver larger fractions of peak hardware performance on scientific applications; and

c) *focus on sensor- and data-intensive computational science applications* in light of the explosive growth of data.

Lessons Not Learned

- *Panel on Large Scale Computing in Science and Engineering,* interagency, 1982
- *From Desktop to Teraflop: Exploiting the U.S. Lead in High Performance Computing,* NSF, 1993
- *Information Technology Research: Investing in Our Future*, PITAC 1999
- *The Biomedical Information Science and Technology Initiative*, NIH, 1999
- *Making IT Better,* National Academies, 2000
- *Embedded Everywhere,* National Academies, 2001
- *High-Performance Computing for the National Security Community*, DOD, 2002
- *Knowledge Lost in Information,* NSF, 2003
- *Revolutionizing Science and Engineering Through Cyberinfrastructure*, 2003
- *A Science-Based Case for Large-Scale Simulation (ScALES),* DOE, 2003
- *Roadmap for the Revitalization of High End Computing*, Interagency, June 2003
- *Supercharging U. S. Innovation & Competitiveness*, Council on Competitiveness, 2004
- *Getting up to Speed: the Future of Supercomputing*, National Academies, 2004

Futures

DR.DUDD
I can assure you that despite the recent set backs, the government is still fully committed to a digital future
www.drdudd.com
(c) Copyright Tony Frampton 2002

NSF, NASA

- DARPA High Productivity Computing Systems
  - 3 phase program [illegible] 8 years 2002-2010
    - aiming at co[illegible]ensive, balanced environments with > 4 P[illegible]
  - Aims t[illegible]elop disruptive technologies for fut[illegible]enerations of HPC applications
    - [illegible]rdware
    - software
  - Currently Phase II: Sun, IBM, Cray
  - Phase III will have two vendors

But this is nowhere near enough!

EU GridLab Project

Tools for New Paradigms on Grids

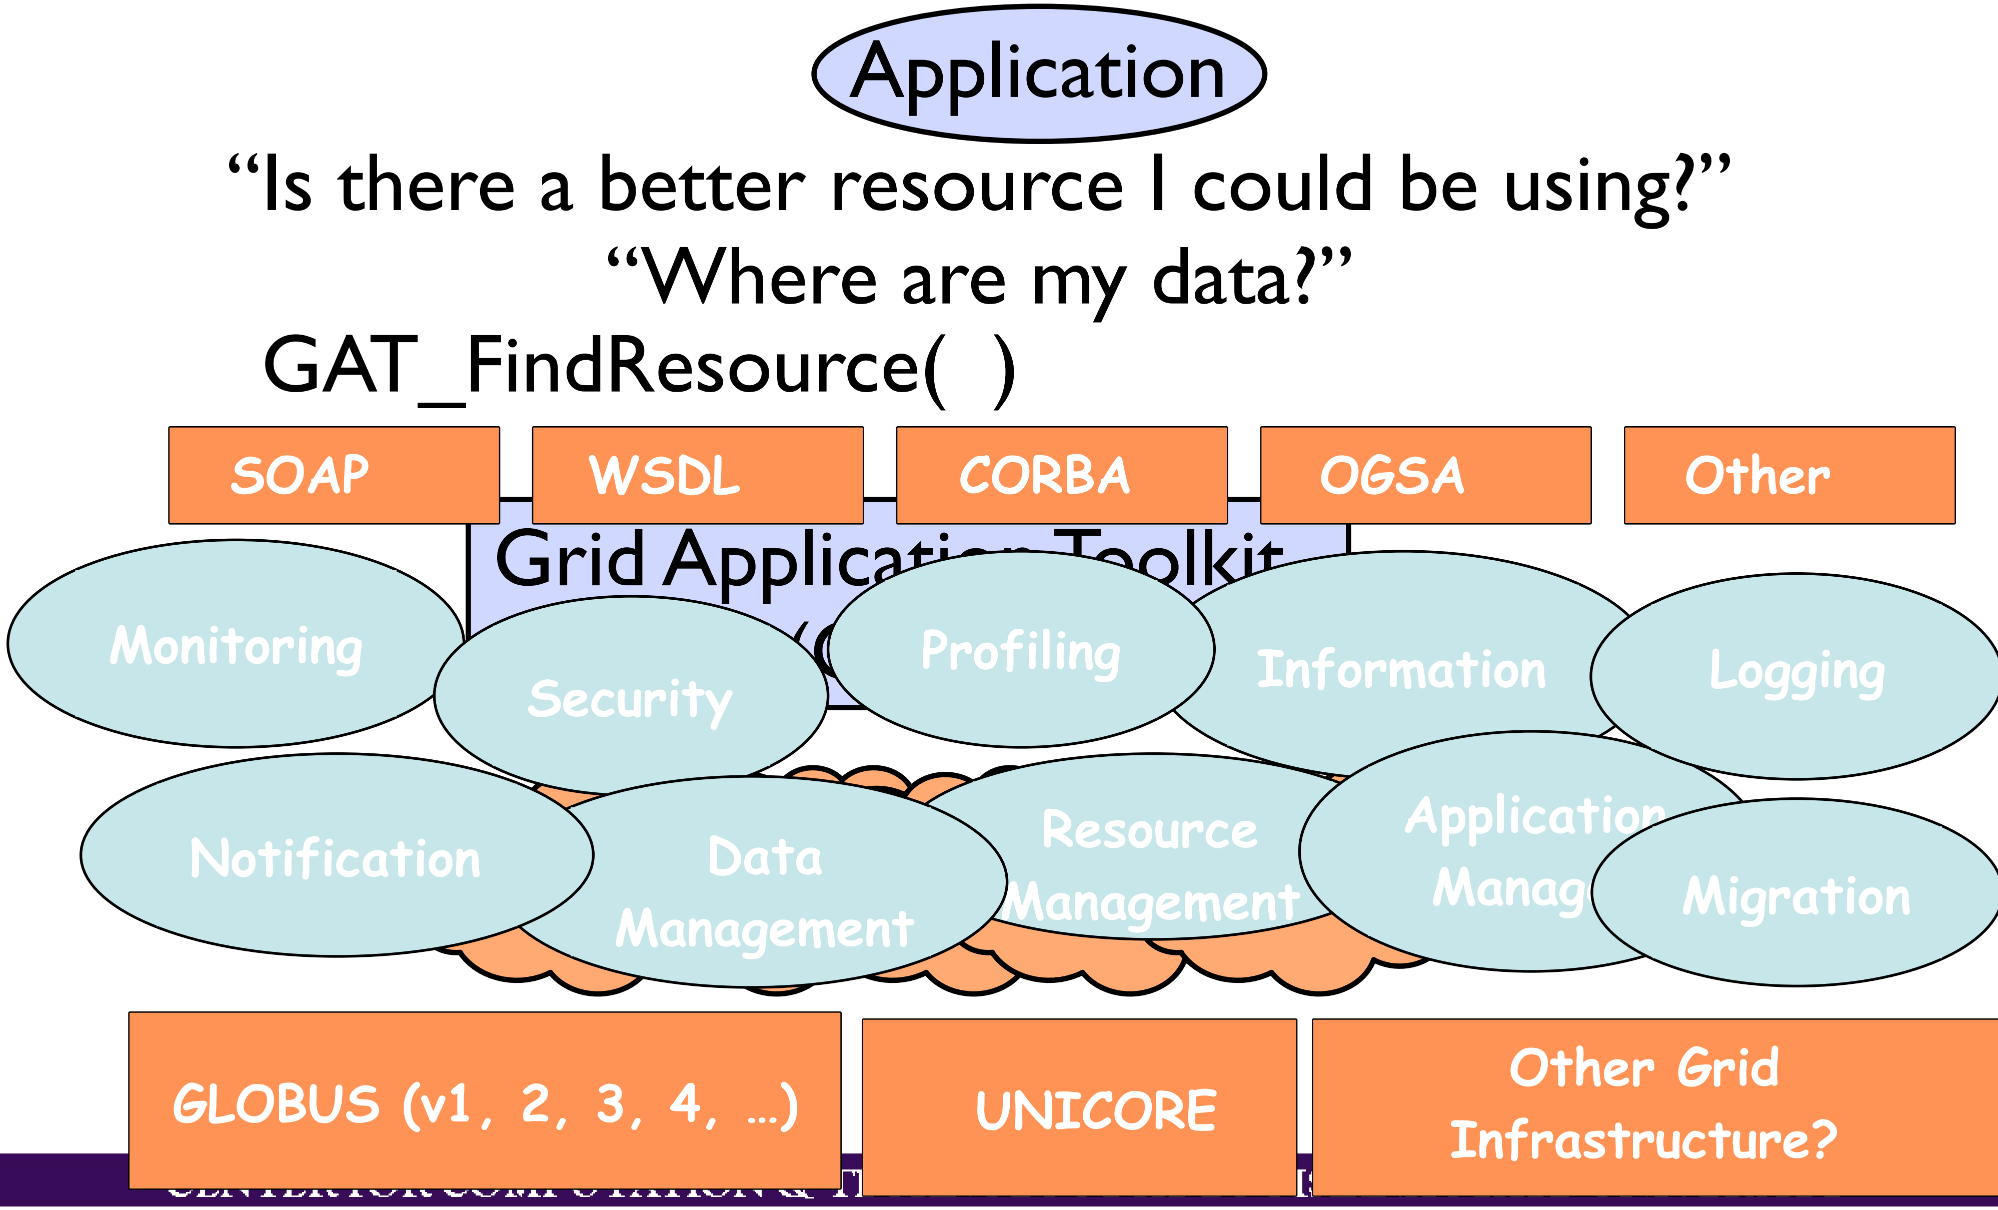

GGF SAGA-RG

- GAT evolves into GGF standard
  - Numerous attempts to address: GAT most ambitious, but also CoG, DRMAA, GridRPC, many others *www.gridlab.org/GAT*
- SAGA: Simple API for Grid Applications
  - Bringing all these efforts together
  - Chicago, Berlin, Brussels, LSU, Berkeley, Seoul, and Chicago in June, 2005
- GGF focussing now on standardization
  - SAGA API spec done
  - Much momentum in SAGA now

forge.gridforum.org/projects/saga-rg

Task Farming, Spawning & Migration

Main Cactus BH Simulation starts in Berkeley

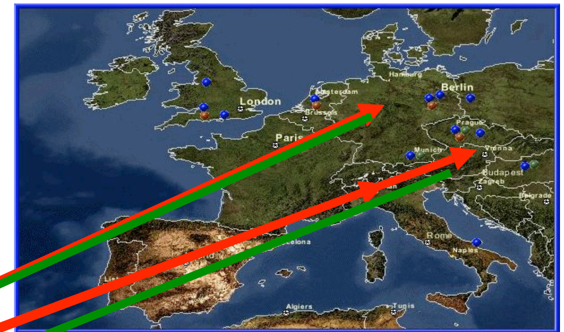

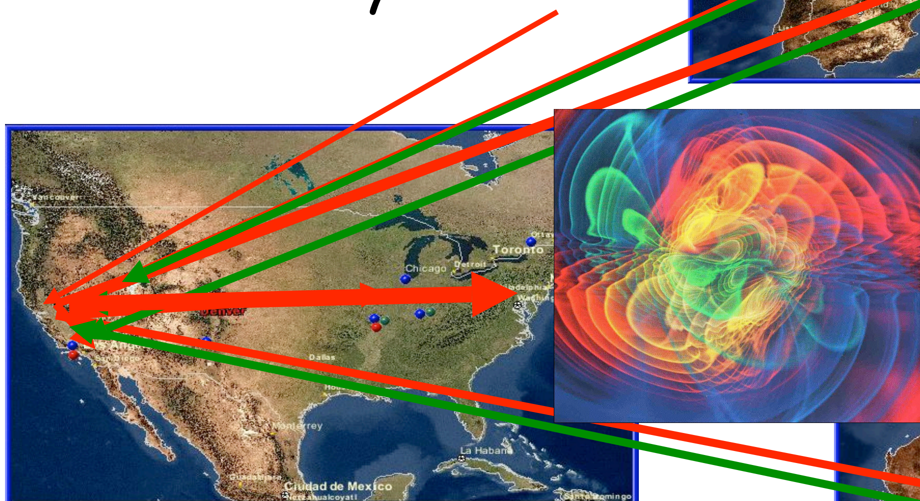

ens of small jobs sent to test parameters

Data returned for main job

Huge job generates remote data to be visualized in Baltimore: lambda on demand!

DDDAS Scenarios

- UCoMS (Petroleum Engineering)
  - Deploy sensor networks across Gulf
  - Data collected to provide input to simulations, tasks farmed out
  - Results collected (transmitted back)
  - http://www.ucoms.org
- SCOOP (Ocean Observing)
  - Data coming in from sensors all over Gulf Realtime Operational Grid
  - Feeds in to models on Grid sites

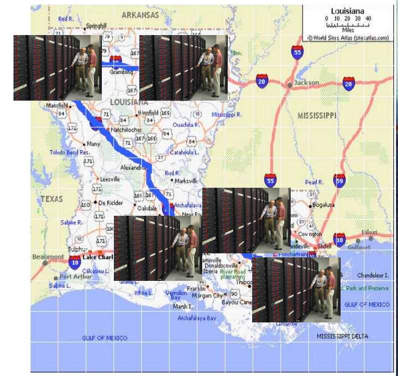

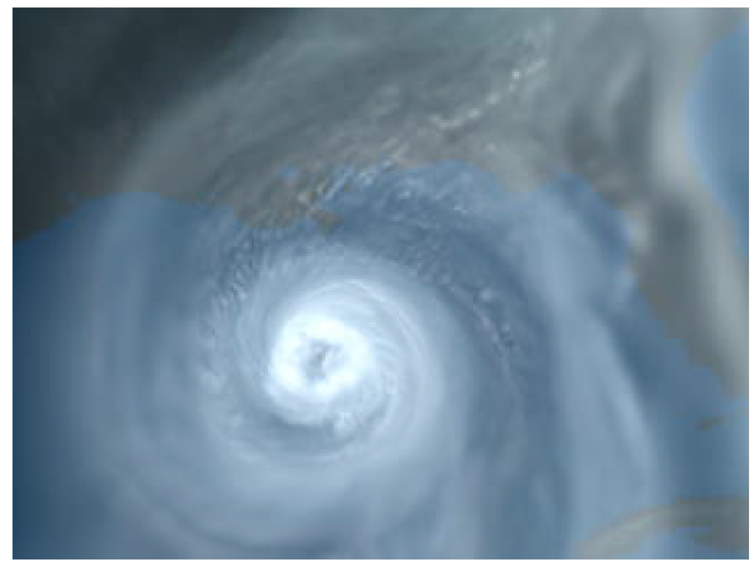

Summary

- Apps challenge HPC, Grids for years to come
- Some challenges partially met, many not
  - Hardware: C
  - Networks: B
  - Software: D
  - Funding (US): F

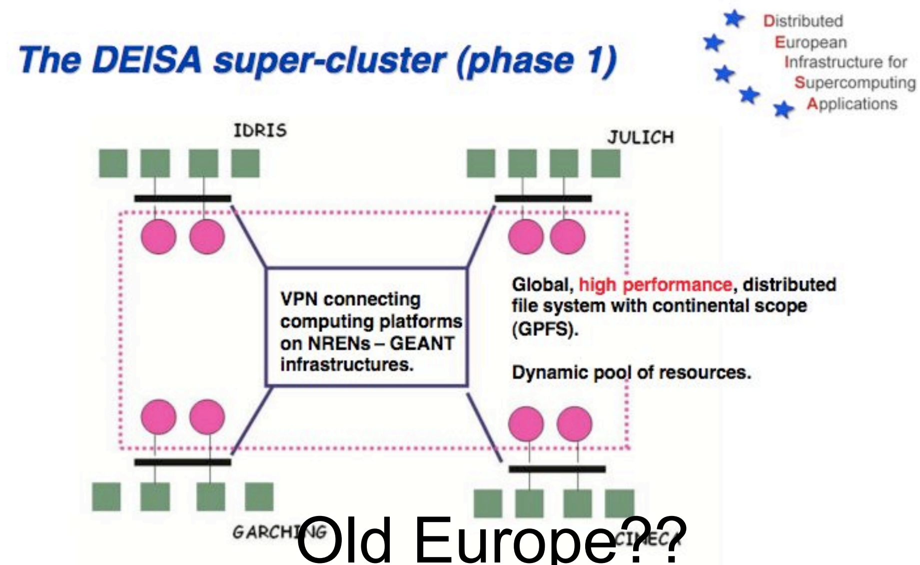

Old Europe??

- Many new application types coming
- DEISA, others, EU 7th Framework poised to fill some of the void, take some leadership
  - But: Grids do not replace supercomputers!! EU still has no SC Center