U.S. Department of Energy

Office of Science

View from Washington and Germantown

Michael Strayer

Associate Director, Advanced Scientific Computing Research

August 5, 2008

Talk Outline

- ASCR Update
  - Staffing
  - ASCR Base Program Solicitation
  - Facility News
  - INCITE
- Top Ten Accomplishments
- SciDAC – Past, Present and Future

ASCR Staffing

- New Hires:
  - Sandy Landsberg, started Monday August 4th as an Applied Math Program Manager in the Research Division
  - Math vacancy announcement closed July 23, 2008
  - CS vacancy announcement should be posted soon
- IPAs:
  - Applied Math: Bill Spotz, SNL, started June 9th
  - CS: Steve Scott, ORNL and Osni Marques, LBNL expected to start in mid to late August
- Detailees
  - Applied Math: Steven Lee, LLNL, started July 1
  - Collaboratories: Susan Turnbull, GSA, started May 19
  - Research Division: Sue Morss, ANL will start later in August

ASCR Base Programs
FY 2008 Request for Proposals

- Multiscale Mathematics and Optimization for Complex Systems
  - Letters of Intent
    - 426 one-page Letters of Intent (LOIs) were submitted by March 3
      - Optimization (138 LOIs): 76 led by Labs, 62 led by universities
      - Multiscale Mathematics (288 LOIs): 186 led by Labs, 102 led by universities
    - From LOIs, ASCR encouraged 114 full proposals due by April 28
  - Optimization peer-review panel convened June 10-11
    - **38** proposals: multilevel, stochastic, dimension reduction, mixed integer, inverse problems
  - Multiscale Mathematics peer-review panels convened June 23-24 and June 25-26
    - **35** proposals: hybrid methods, geoscience, fluids and plasmas, materials, data-model fusion
    - **25** proposals: uncertainty and sensitivity analysis, stochastic methods, additional topics

- Petascale Tools
  - 97 proposals received representing 34 projects
  - Topics included Performance tools, Correctness tools, Development Environment and Scalable Infrastructure
  - Review scheduled for August 26-27

Facility News
ORNL

First 17 XT5 cabinets arrived July 30, 2008

Facility News
ANL

World’s Fastest Open Science Computer

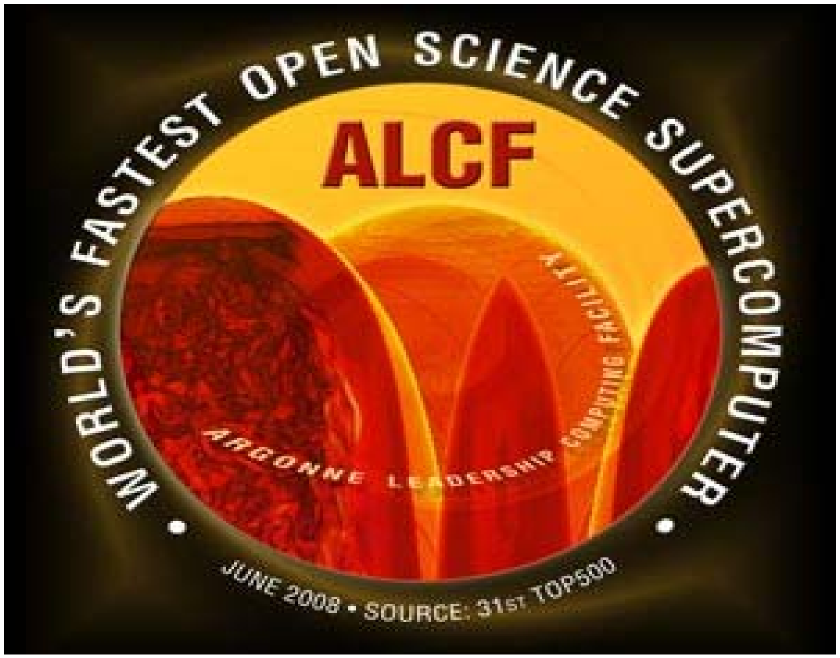

5 ORNL Cray XT4

15 NERSC Cray XT4

The U.S. Department of Energy's (DOE) Argonne National Laboratory's IBM Blue Gene/P high-performance computing system is now the fastest supercomputer in the world for open science, according to the June, 2008 Top500 List of the world's fastest computers.

The Blue Gene/P known as Intrepid and located at the Argonne Leadership Computing Facility (ALCF) also ranked third fastest overall. Both rankings represent the first time an Argonne-based supercomputing system has ranked in the top five of the industry's definitive list of supercomputers.

Facility News
NERSC

NERSC-6 Lehman Review July 9-10, 2008

(left to right) Mike McCoy, Yukiko Sekine, Jim Tomkins, Jackie Chen, Warren Yip, Stephane Ethier, Dan Hitchcock, Doug Kothe, Kurt Fisher, Dan Lehman, Stephen Meador, Kathlyn Boudwin, John Morrison. Pete Beckman, Frank Gines, (missing: Mark Seager, Ira Goldberg)

"NERSC presented a thorough and methodical science requirements analysis that made a compelling case for delivery of the NERSC-6 system to the science community in as timely manner as possible. The science output generated on existing NERSC systems is impressive, with a long-lasting legacy of high-impact accomplishments. NERSC convincingly articulated that science user demand continues to outstrip supply…"

Lehman Review Report from NERSC Review July 9-10, 2008 to evaluate the LBNL/NERSC plan for procuring NERSC-6.

Innovative and Novel Computational Impact on Theory and Experiment (INCITE)

- **Provides Office of Science computing resources to a small number of projects that are**
  - **Computationally intensive**
  - **Large scale**
  - **High impact**
- **Open to national and international researchers, including industry**
- **No requirement of DOE Office of Science funding**
- **Peer and computational reviews**

INCITE
2008

Received 88 new and 24 renewal proposals requesting over 600 Million processor hours

- 44% from Universities
- 46% funded from non-DOE sources

Over 265 Million hours awarded to 55 new and renewal projects

INCITE
2008

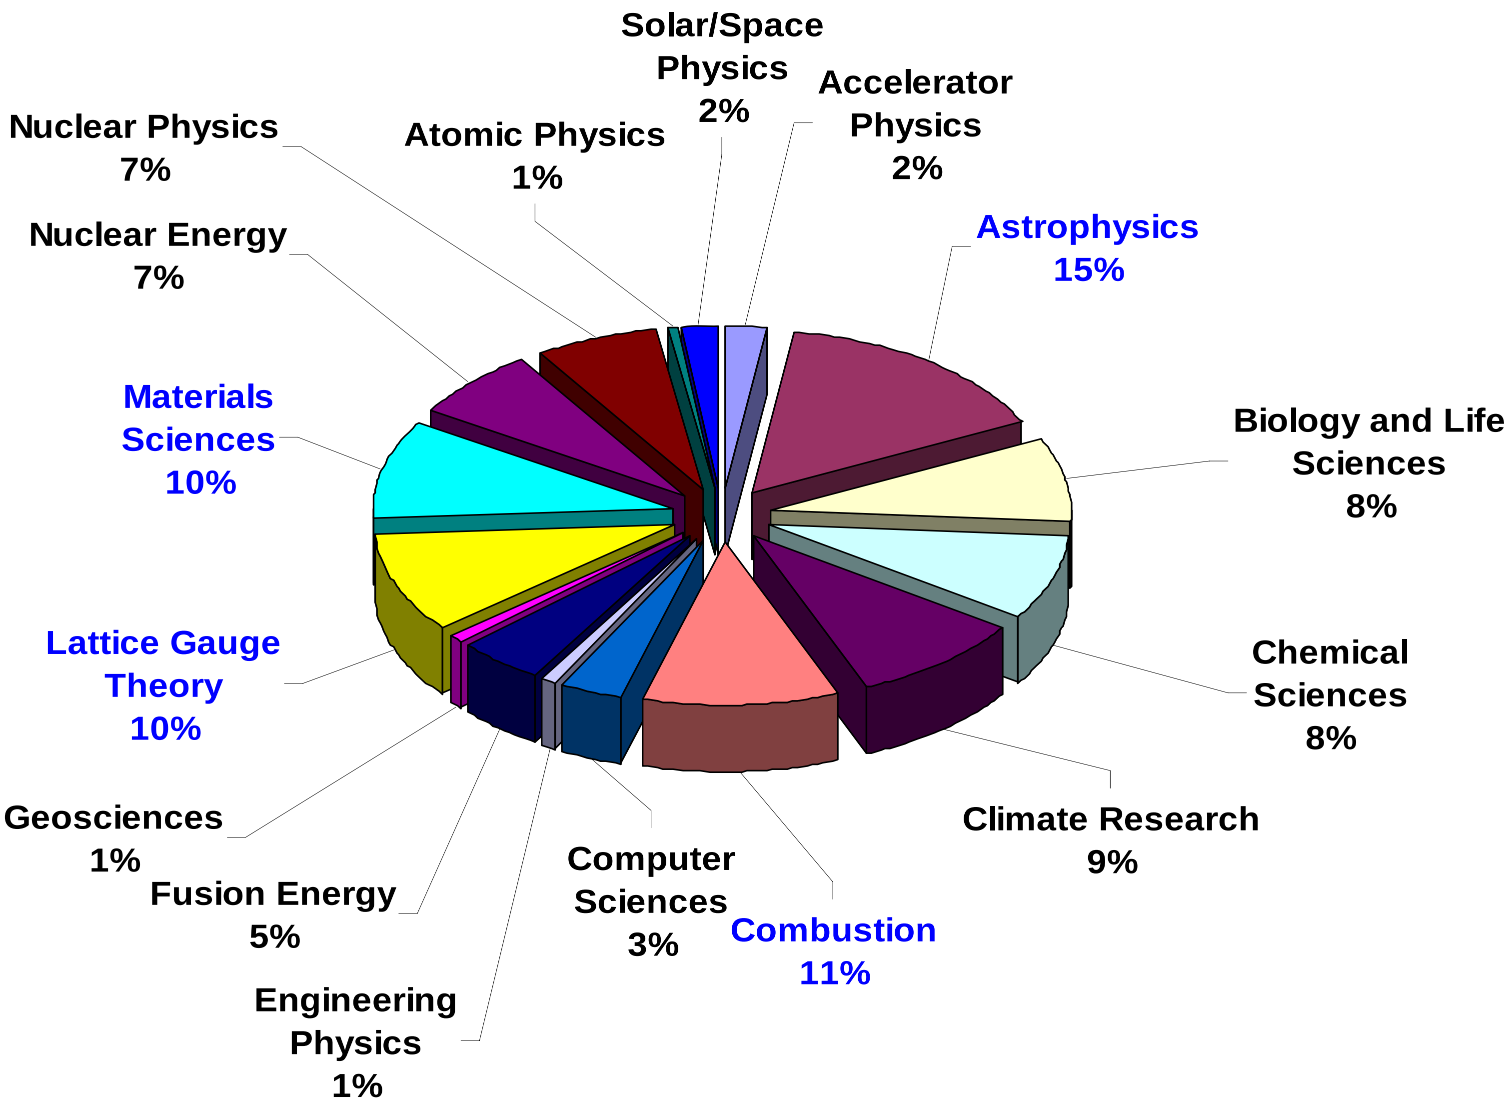
Solar/Space Physics 2%
Accelerator Physics 2%
Astrophysics 15%
Biology and Life Sciences 8%
Chemical Sciences 8%
Climate Research 9%
Combustion 11%
Computer Sciences 3%
Engineering Physics 1%
Fusion Energy 5%
Geosciences 1%
Lattice Gauge Theory 10%
Materials Sciences 10%
Nuclear Energy 7%
Nuclear Physics 7%
Atomic Physics 1%

INCITE 2009

hpc.science.doe.gov

- August 11, 2008: Deadline for New Proposal Submission
- September 12, 2008: Deadline for Renewals

Over a half a billion INCITE hours available in 2009 on systems at ANL, LBNL, ORNL and PNNL

Top 10 Accomplishments in Computational Science

Computational Science Accomplishments Panel

The Panel

Panelist	Area
Beckman (ANL)	Computer Science
Chen (Sandia)	Accelerator Physics, Combustion, Fluid Turbulence, Engineering Physics
Galli (UCD)	Chemical Sciences, Nano and Materials Sciences, Physical Chemistry
Hack (ORNL)	Climate Research, Environmental Sciences, Geosciences
Keyes (Columbia)	Applied Mathematics
Kothe (ORNL)	
Messina (ANL)	
Mezzacappa (ORNL)	Astrophysics, Solar/Space Physics
Rebbi (BU)	Lattice Gauge Theory, Nuclear Physics
Samatova (NCSU)	Biology, Life Sciences
Tang (PPPL)	Fusion, Fusion, Energy, Plasma Physics
Yelick (LBL)	

Center directors played an instrumental role in:

- *Filling gaps.*
- *Facilitating comparisons across areas.*
- *Facilitating selections involving panel members.*

Process Outline

Final Set
(10 Projects)

Entire panel considered and ranked final candidates.

One Per Application Area +
(1-2 Dozen Projects)

Panelists solicited 1-3 "one-pagers" from teams in their area of expertise, following a template.

SciDAC, INCITE, Base Program
(Large Number of Projects)

Top 10 Computational Science Accomplishments

Rank	Title
1	**Modeling the Molecular Basis of Parkinson's Disease (Tsigelny)**
2	**Discovery of the Standing Accretion Shock Instability and Pulsar Birth Mechanism in a Core-Collapse Supernova Evolution and Explosion (Blondin)**
3	**Prediction and Design of Macromolecular Structures and Functions (Baker)**
4	**Understanding How Lifted Flame Stabilized in a Hot Coflow (Yoo)**
5	**New Insights from LCF-enable advanced kinetic simulations of global turbulence in fusion systems (Tang)**
6	**High Transition Temperature Superconductivity: A High-Temperature Superondructive State and a Pairing Mechanism in 2-D Hubbard Model (Scalapino)**
7	**PETsc: Providing the Solvers for DOE High-Performance Simulations (Smith)**
8	**Via Lactea II, A Billion Particle Simulation of the Dark Matter Halo of the Milky Way (Madau)**
9	**Probing the properties of water through advanced computing (Galli)**
10	**First Provably Scalable Maxwell Solver Enables Scalable Electromagnetic Simulations (Kovel)**

Modeling the Molecular Basis of Parkinson's Disease

⇨ Igor Tsigelny (SDSC)
⇨ Eliezer Masliah (UCSD)
⇨ Stanley Opella (UCSD)

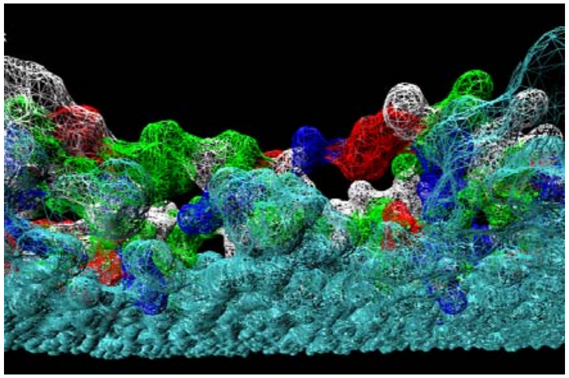

Accomplishment

• Elucidated the molecular mechanism of the progression of Parkinson's disease.
• Insights will help focus the search for treatment.

Broader Implications

• Provided a test bed for identifying possible therapeutic interventions through computational modeling.
• Overall approach has broad applicability to other diseases.

Support: INCITE

I.F. Tsigelny et al. 2007. Dynamics of α-synuclein aggregation and inhibition of pore-like oligomer development by β-synuclein, *FEBS J.*, **274**, 1862-1877.

Prediction and Design of Macromolecular Structures and Functions

Accomplishment

First predictions of protein structures with atomic-level accuracy.

Atomic level protein structure prediction has been a Holy Grail of molecular biology for over thirty years.

Broader Implications

Accurate protein structure prediction will greatly speed the interpretation of the genome sequence information generated in ongoing large-scale sequencing projects such as the Human Genome Project.

Support: INCITE

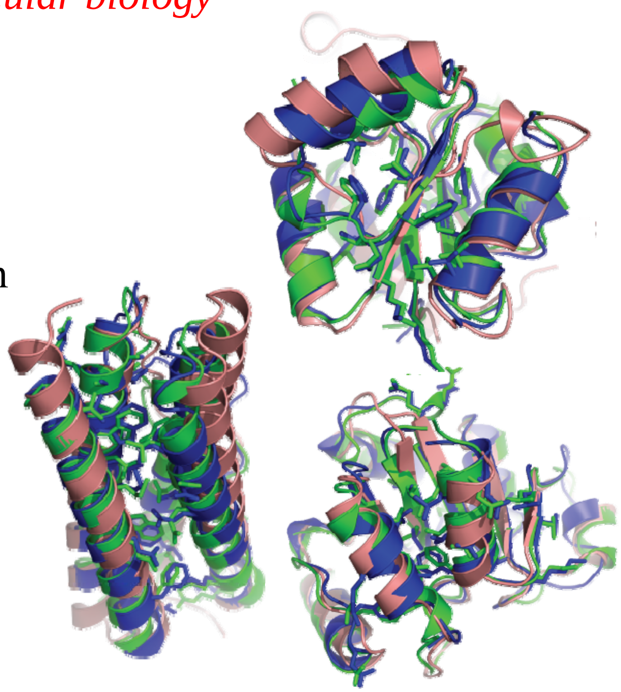

1. R�hlisberger D, Khersonsky O, Wollacott AM, Jiang L, Dechancie J, Betker J, Gallaher JL, Althoff EA, Zanghellini A, Dym O, Albeck S, Houk KN, Tawfik DS, Baker D., Kemp elimination catalysts by computational enzyme design, ***Nature***. 2008 Mar 19; [Epub ahead of print]
2. Jiang L, Althoff EA, Clemente FR, Doyle L, R�thlisberger D, Zanghellini A, Gallaher JL, Betker JL, Tanaka F, Barbas CF 3rd, Hilvert D, Houk KN, Stoddard BL, Baker D., De novo computational design of retro-aldol enzymes, ***Science***. 2008 Mar 7;319(5868):1387-91.
3. Qian B, Raman S, Das R, Bradley P, McCoy AJ, Read RJ, Baker D. High-resolution structure prediction and the crystallographic phase problem, **Nature**. 2007 Nov 8;450(7167):259-64. Epub 2007 Oct 14.

LCF-Enabled Simulations of Global Turbulence in Fusion Systems

PPPL, UCI, ORNL, Columbia, UCSD

Accomplishment

3D simulations of unprecedented resolution led to a fundamentally new understanding of thermal energy loss in magnetically confined toroidal plasmas, such as in fusion tokomaks.

Broader Implications

Critical to the design of ITER.

⇨ Size and cost depend on balance between losses and self heating.

Support: FES Base Program, SciDAC, and INCITE

ITER

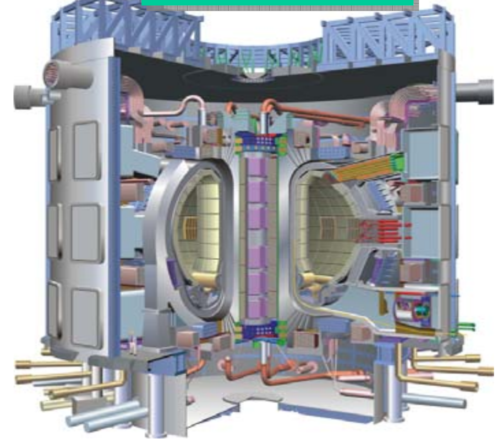

FUSION TURBULENCE SIMULATION

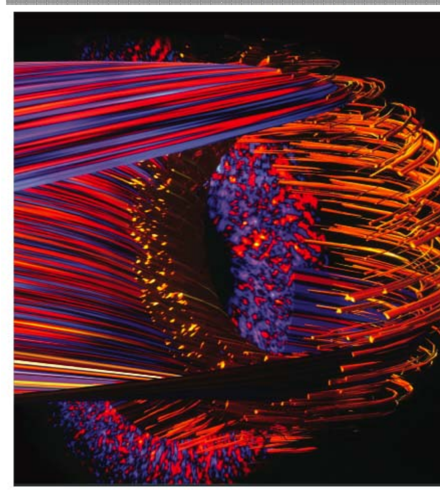

1. Z. Lin, I. Holod, L. Chen, P. H. Diamond, T. S. Hahm, and S. Ethier, *Phys. Rev. Lett.* **99**, 265003 (2007)
2. Wang W.X., Hahm T.S., Lee W.W., et al., *Phys. Plasmas* **14**, 072306 (2007)

U.S. Department of Energy

Office of Science

High-Transition Temperature Superconductivity

D. Scalapino (UCSB)
T. Maier, P. Kent, and T. Schulthess (ORNL)
M. Jarrell and A. Macridin (University of Cincinnati)
D. Poilblanc (Laboratoire de Physique Thˇorique, CNRS and Universitˇ de Toulouse)

Accomplishment

- Given new numerical techniques and leadership-class resources, 2D Hubbard Model solved computationally.
- Proved model does in fact describe high-temperature superconductivity.

⇨ *Settled a debate that raged for two decades.*

Broader Implications

- Provides a deeper understanding of high-T superconductivity.
- Key step toward the identification of high-T superconducting materials, *the Holy Grail of superconductivity research*.

Support: BES Base Program, SciDAC

Maier, T.A., A. Macridin, M. Jarrell, and D. J. Scalapino, ÒSystematic analysis of a spin-susceptibility representation of the pairing interaction in the two-dimensional Hubbard modelÓ *Phys. Rev. B*. **76**, 144516 (2007).
Maier, T.A., M. Jarrell, and D. J. Scalapino, ÒSpin susceptibility representation of the pairing interaction for the two-dimensional Hubbard modelÓ*Phys. Rev. B*. **75**, 134519 (2007).

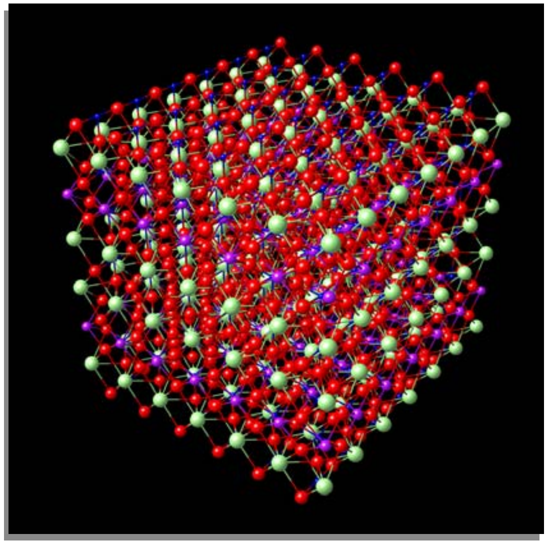

Model of a $YBa_2Cu_3O_7$ high-temperature superconductor crystal

Via Lactea II: A Billion Particle Simulation of the Dark Matter Halo of the Milky Way Galaxy

- ⇨ Piero Madau (UCSC)
- ⇨ Juerg Diemand (UCSC)
- ⇨ Michael Kuhlen (IAS)
- ⇨ Marcel Zemp (UCSC)

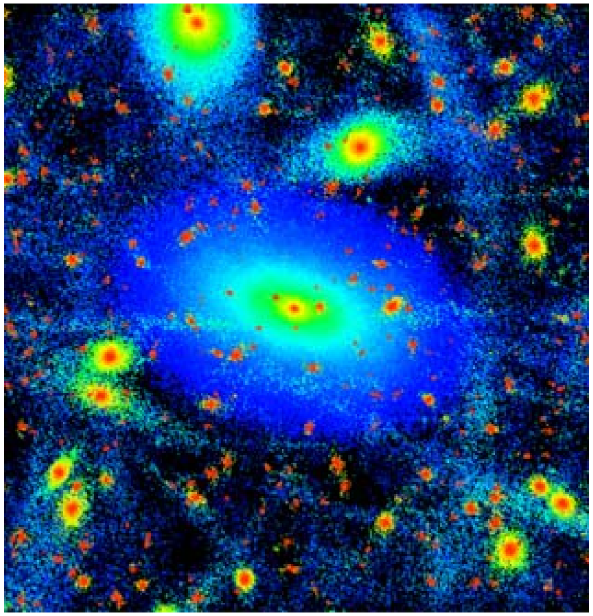

Accomplishment

Largest simulations ever performed of the formation of the dark matter halo of the Milky Way galaxy.

⇨ *Resolved and predicted structure at small scales.*

Broader Implications

Structures predicted for small scales have observable gamma-ray signatures for certain classes of dark matter candidates.

✶ *Simulations may play an important role in identifying the dark matter in the Universe.*

Support: INCITE

Diemand, Kuhlen, Madau, Zemp, Moore, Potter, and Stadel 2008, Clumps and Streams in the Local Dark Matter Distribution, *Nature*, in press.

PETSc: Providing Solvers for DOE Simulations

- Developed portable, scalable solvers for solving large-scale algebraic systems.
  - Have solved problems with over 3 billion unknowns and scaled to over 16,000 processes on DOE leadership-class computers.
- Enabled DOE applications using PETSC on leadership-class systems.
  - UNIC (nuclear reactor simulation), PFLOTRAN (reservoir modeling), STOMP (subsurface transport), M3D (plasma simulation), and GTC (gyrokinetics).

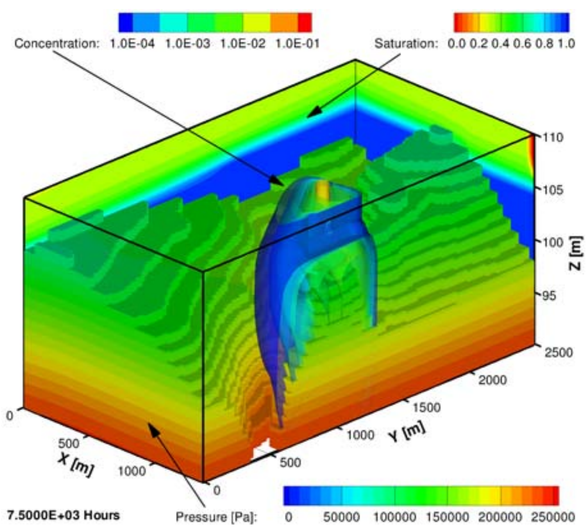

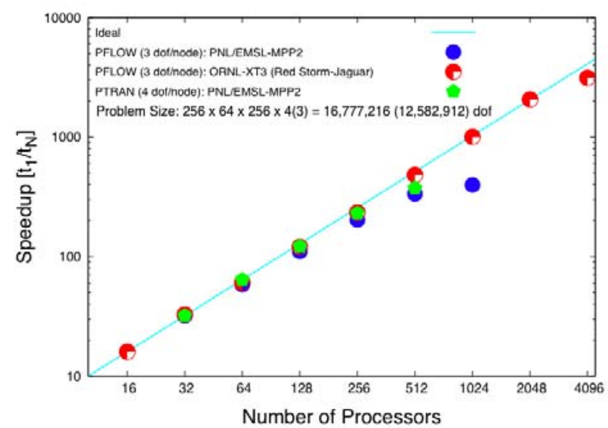

Sample computation using the PETSc powered PFLOTRAN subsurface flow simulator and its parallel performance on the DOE leadership-class Cray XT3 system

SciDAC – Past Present and Future

Scientific Discovery through Advanced Computing (SciDAC)

"SciDAC is unique in the world. There isn't any other program like it anywhere else, and it has the remarkable ability to do science by bringing together physical scientists, mathematicians, applied mathematicians, and computer scientists who recognize that computation is not something you do at the end, but rather it needs to be built into the solution of the very problem that one is addressing."

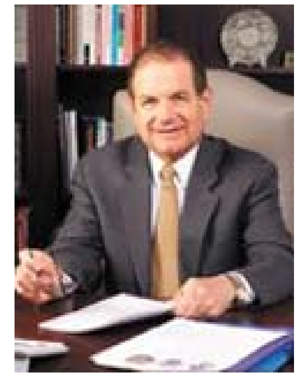

Dr Raymond Orbach, Under Secretary for Science and Director, Office of Science, US Department of Energy in SciDAC Review

SciDAC 1

- **Create**
  - **a new generation of Scientific Simulation Codes**
  - **Mathematical and Computing Systems Software**
  - **Collaboratory Software Environment**

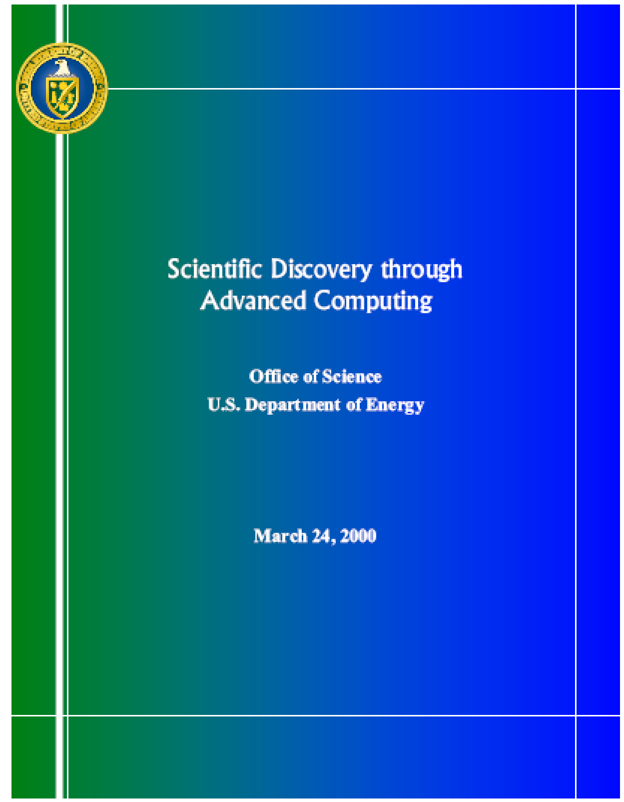
Scientific Discovery through Advanced Computing

Office of Science
U.S. Department of Energy

March 24, 2000

Supported by upgrades to ASCR's Scientific Computing Infrastructure

SciDAC 1
The New Paradigm

Advancing Science through Modeling and Simulation

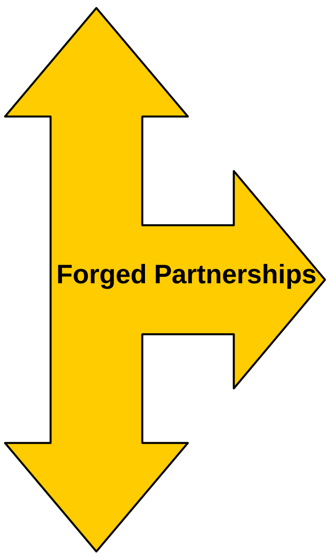

Compute & Network Resources

First Federal program to implement Computational Science and Engineering

SciDAC 1

- Accelerators
- Climate
- Fusion
- QCD
- Astrophysics
- Chemistry
- Materials

- National Collaboratories
- Integrated Software Infrastructure Centers (ISICs)

NERSC Allocation

NERSC	3 Teraflops (TF)
ESnet	622 Mbps bandwidth

U.S. Department of Energy

Office of Science

SciDAC 1

Applied Mathematics Integrated Software Infrastructure Centers (ISICS)

- **Applied Partial Differential Equation Center (APDEC)**
  - Develop a high performance algorithmic and software framework for multiscale problems based on the use of block-structured adaptive mesh refinement (AMR) for representing multiple scales
- **Terascale Simulation Tools and Technologies (TSTT)**
  - Develop interoperable services that enable mesh generation for representing complex and possibly evolving domains, high-order discretization techniques for improved numerical solutions and adaptive strategies for automatically optimizing the mesh to follow moving fronts.
- **Terascale Optimal PDE Simulations (TOPS)**
  - Develop and implement algorithms and support DOE-sponsored scientific investigations; develop and deploy an integrated toolkit of open-source, optimal complexity solvers for nonlinear partial differential equations.

SciDAC 1
Computer Science ISICs

- **Performance Evaluation Research Center (PERC)**
  – Develop a *science* for understanding performance of scientific applications on high-end computer systems and *engineering* strategies for improving performance on these systems
- **Scientific Data Management (SDM)**
  – Optimize and simplify access to very large datasets, distributed and heterogeneous data and data mining of very large data sets
- **Scalable Systems Software for TeraScale Computers (SSS)**
  – Define, with industry, standard interfaces between systems components for interoperability and create scalable, standardized management tools for efficiently running large computing centers
- **Component Technology for Terascale Simulation Software (CCA)**
  – Integrate legacy software investment by extending commodity component technology to HPC

SciDAC 1
Collaboratories

- **DOE Science Grid**
  – Provide persistent Grid services to advanced scientific applications and problem-solving frameworks
- **National Fusion Collaboratory**
  – Develop a system for secure sharing of fusion computational, visualization and data resources over the Interne
- **Particle Physics Data Grid**
  – Collaboration between physicists and computer scientists to adapt and extend Grid technologies and the physicist's applications to build distributed systems to handle the storage, retrieval and analysis of the particle physics data at DOE's critical research facilities
- **Earth Science Grid**
  – Collaboratory interdisciplinary project aimed at addressing the challenge of enabling management, discovery, access and analysis of climate simulation datasets

U.S. Department of Energy

Office of Science

SciDAC 1
Highlights

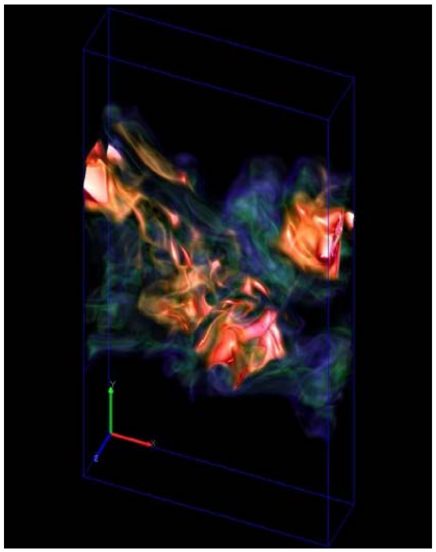

Hydroxyl radical in a turbulent jet flame

- **SciDAC teams created first laboratory-scale flame simulation in three dimensions to better understand combustion**
- **Partnerships improved effectiveness of scientific applications codes between 275% to over 10,000%**
- **The SciDAC data mining tool Sapphire awarded a 2006 R&D100 award**

Publications	
2001	94
2002	186
2003	277
2004	247
2005-06	377

SciDAC 1 → SciDAC 2

Program re-competed in FY2006
Partners- All Office of Science programs, NSF and NNSA

Goals

- **Create comprehensive, scientific computing software infrastructure to enable scientific discovery in the physical, biological, and environmental sciences at the petascale**
- **Develop new generation of data management and knowledge discovery tools for large data sets (obtained from scientific user facilities and simulations)**

Over 230 proposals received requesting approximately $1 Billion

U.S. Department of Energy

Office of Science

SciDAC 2

Path to Petascale

Scientific Discovery

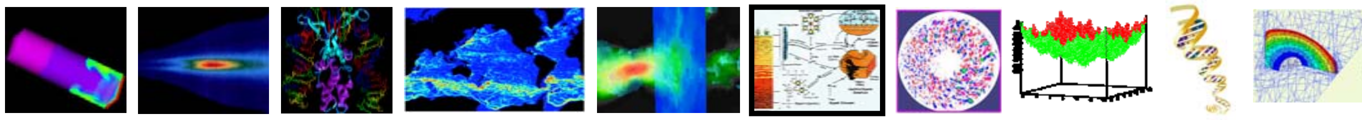

Applications

Computing/ Networking

SciDAC 2

- Accelerator science and simulation
- Climate modeling and simulation
- Fusion science
- Petabyte high-energy/nuclear physics
- Nuclear physics
- Radiation transport
- Astrophysics
- Computational Biology
- High-energy physics
- Materials science and chemistry
- QCD
- Turbulence
- Groundwater reactive transport modeling and simulation

- Centers for Enabling Technology
- Scientific Applications Partnerships
- Institutes (University-lead)

INCITE
NERSC Allocation

Leadership Computing-ANL	556 TF IBM BG/P
Leadership Computing-ORNL	263 TF Cray XT4 → 1 PF Cray XT5
Production Computing-NERSC	104 TF Cray XT4 → ~360 TF Cray XT4
ESnet	On path toward Dual rings 40Gbps/ 10 Gbps fault tolerant

SciDAC 2
Applied Math

Centers:

- **APDEC – Applied Partial Differential Equations.**
  - **Development of simulation tools for solving multi-scale and multi-physics problems based on finite difference and finite volume methods on logically-rectangular structured grids combined with block structured adaptive mesh refinement.**
- **ITAPS – Interoperable technologies for advanced petascale simulations**
  - **Delivery of interoperable and interchangeable mesh, and geometry and field manipulation services**
- **TOPS– Towards Optimal Petascale Simulations**
  - **Development, testing and dissemination of solver software for systems governed by partial differential equations.**

Institutes:

- **CSCAPES – Combinatorial Scientific Computing and Petascale Simulations.**
  - **Development of**
    - **new capabilities in load balancing and parallelization toolkits for petascale computers;**
    - **new automatic differentiation capabilities; and**
    - **sparse matrix software tools.**

APDEC

Simulation of a Low-Swirl Hydrogen Burner

Motivation: Office of Fossil Energy program to design turbines for power generation based on low NO_x lean premixed hydrogen combustion.

Goal: use simulation to map out the engineering design space for such systems. The large differences between hydrogen and traditional hydrocarbon fuels makes this a necessity.

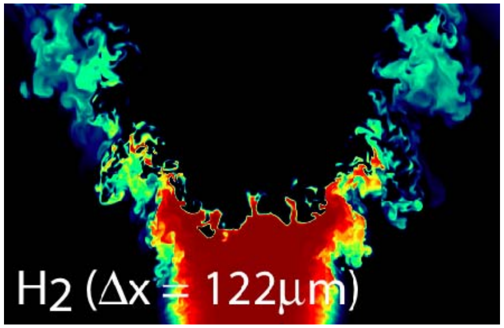

Adaptive mesh refinement calculation of Low Mach number combustion. Five levels of refinement, with the finest grid occupying ½% of the domain. Computed using 2400 processors on Cray XT4 at NERSC.

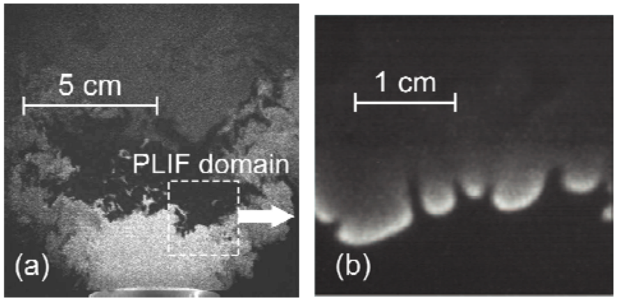

Collaboration with laboratory experimental program (R. Cheng, LBNL).

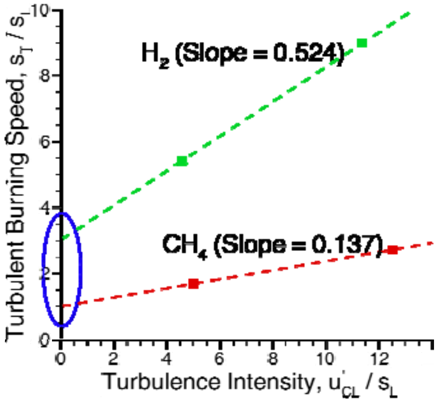

Quantifying dependence of burning rate on turbulence intensity for different fuels, with substantial impact on engineering design.

APDEC

APDEC

Embedded Boundary Methods for PDE in Complex Geometries

New methods for grid generation based on implicit function representations of a domain $\Omega = \{x : \phi(x) = 0\}$

- Applicable to broad range of image, geographical, engineering representations.
- Fast, automatic; generalizes to any order of accuracy, any number of dimensions.

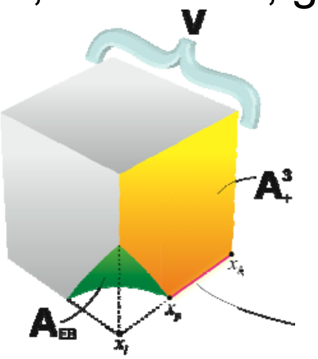

Grid generation for embedded boundary methods consists of computing areas, volumes, and moments associated with the intersection of each rectangular grid cell with the irregular domain.

Low-swirl injector for gas turbine. **Left:** surface description from embedded boundary grid generator. **Right**: blowup showing intersection of the grid with the vanes.

Plan: Details of cold flow through the low-swirl injector will be simulated and sampled to provide inflow conditions for hydrogen combustion calculation.

ITAPS

Curved mesh correction tool speeds SLAC electromagnetic analysis

- High-order meshes on curved geometries can be inverted
  - Solution blows-up unless negative contribution caused by inverted elements removed
- ITAPS tools automatically detect and correct inverted elements
- Applied to 8-cavity cryomodule simulations
  - 2.97 million curved regions
  - 1,583 invalid elements corrected - leads to stable simulation and that execute 30% faster

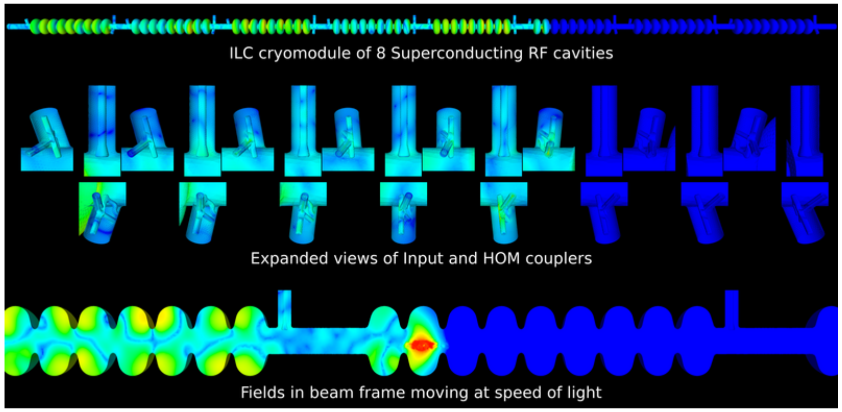

Valid curved mesh after operations Electromagnetic analysis for SLAC

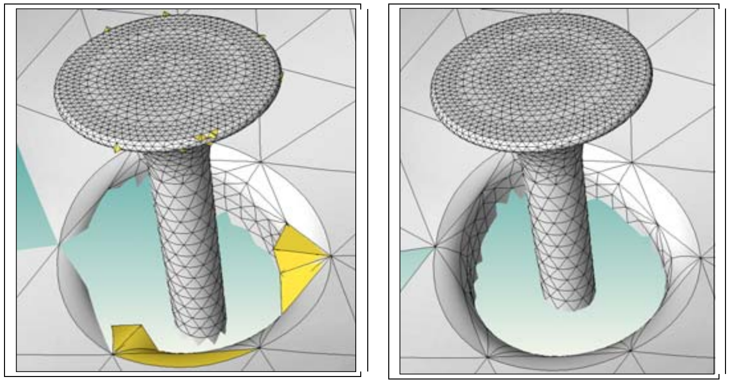

***POC**: Mark Shephard, RPI*

ITAPS
INTEROPERABLE TOOLS FOR ADVANCED PETASCALE SIMULATIONS

ITAPS

Mesh refinement tools increase efficiency of accelerator and fusion codes

- Mesh adaptation used to refine the small domain around the particle beams in accelerator simulations
  - Adaptively refined meshes gave same accuracy as uniform grids with 15-25% of the grid points
- Mesh adaptation developed for CEMM M3D-C1 magneto-hydrodynamics code
  - Provided API level for interacting with mesh and solvers (three-way CEMM, ITAPS,TOPS interaction)
  - Developed and successfully demonstrated initial anisotropic mesh adaptation procedure
  - Extending code to complex boundary conditions on curved domains

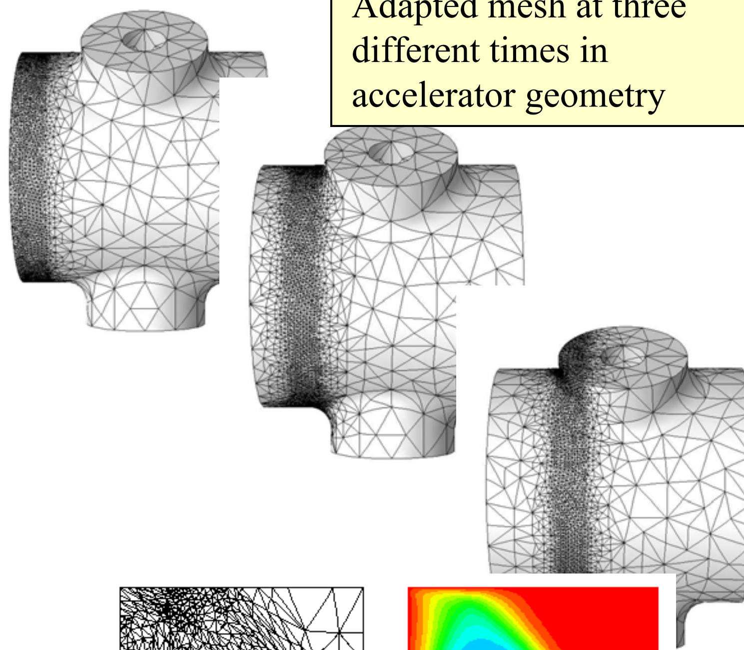

Adapted mesh at three different times in accelerator geometry

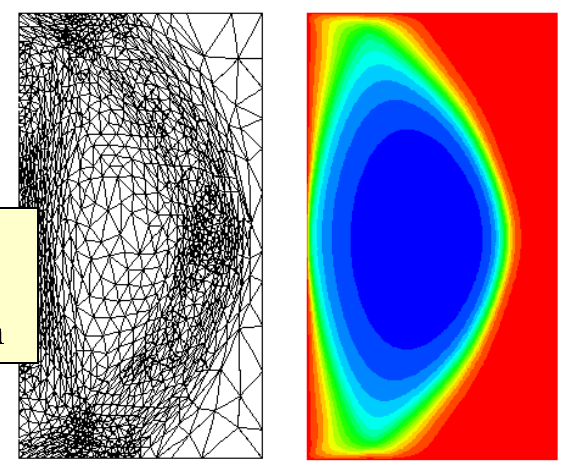

Anisotropic mesh adaptation in CEMM simulation

***POC**: Mark Shephard, RPI*

U.S. Department of Energy

Office of Science

ITAPS

Front tracking simulations leads to unexpected pellet ablation results in fusion

- ITAPS adaptive mesh refinement combined with front tracking used to perform first systematic study of microscale physics of pellet ablation in tokomak fueling
  - Revealed new properties of the ablation flow (strong sensitivity to warm up time, supersonic rotation of the ablation channel responsible for striation instabilities)
  - Corrected a published misconception regarding the reduction of the ablation rate in hydrodynamic models with directional heating
- Study suggests novel pellet acceleration techniques needed for ITER
- Collaborative effort with General Atomics

***POC**: Roman Samulyak, BNL*

Distributions of the Mach number of the ablation flow near the pellet

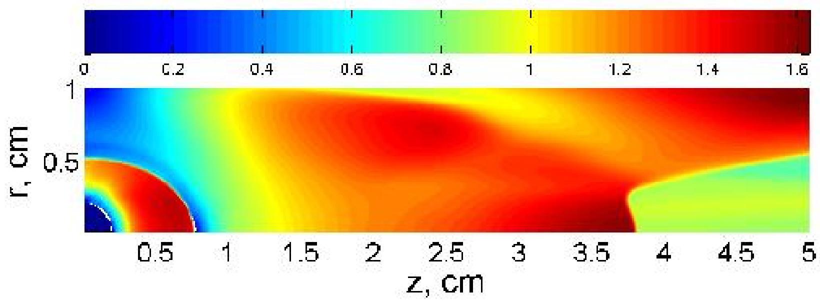

TOPS
Recent Accomplishments

- **Subsurface transport**
  - Strong scaling with 60% efficiency beyond 30,000 cores of XT4 for flow in porous media applied to nuclear waste tracking (PFLOTRAN code, LANL)
- **Neutronics**
  - Weak scaling with 85% efficiency beyond 70,000 cores of BG/L for discrete ordinates method applied to nuclear reactor core modeling (UNIC code, ANL)
- **Quantum chromodynamics**
  - Robust solver based on adaptive algebraic multigrid to bypass "critical slowing down" experienced in the standard conjugate gradient approach to inverting the Dirac operator in zero quark mass limit
- **Cray announces support of TOPS software in its libSci on the XT4 platforms**
  - Uses TOPS' (Berkeley) OSKI software (a sparse analog of ATLAS, but significantly more complex) to tune TOPS data structures for efficient exploitation of XT4 architecture

TOPS
Terascale Optimal PDE Simulations

TOPS PETSc
Example Applications

LANL's PFLOTRAN subsurface flow code
Linear solver application
270 million unknowns
Field data from the Hanford site:
ORNL's Cray XT 4

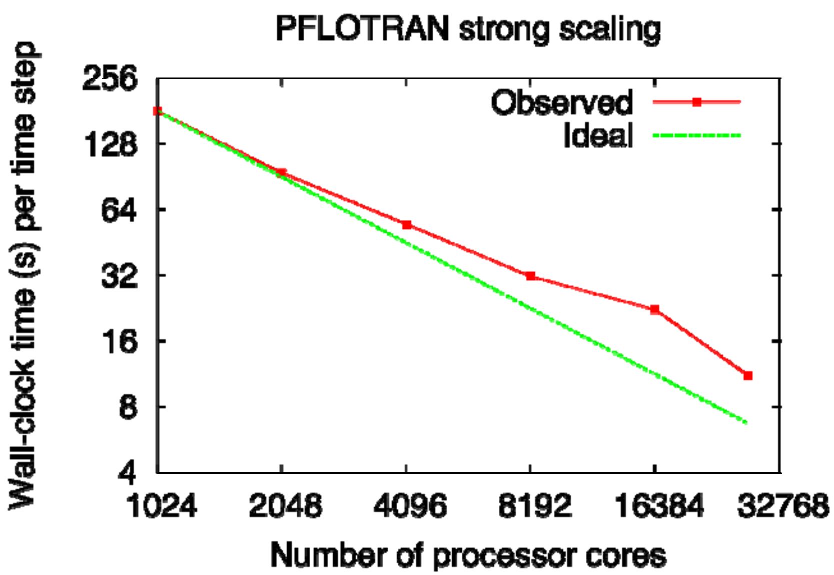

c/o Barry Smith, ANL, *et al.*

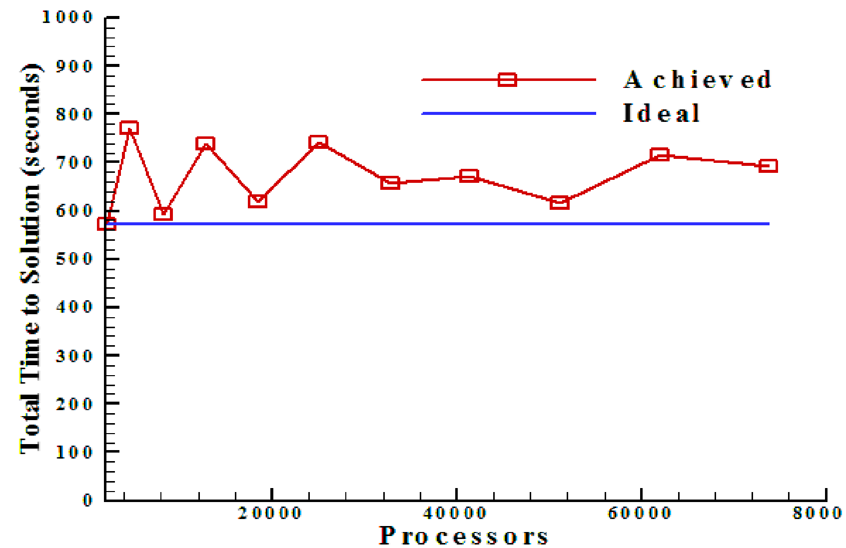

Neutron transport for full reactor core simulations
Nonlinear solver application
Weak scaling
2^{nd}-order Sn method
ANL's IBM BG/P

c/o Dinesh Kaushik, ANL, *et al.*

CSCAPES
Progress and Accomplishments

- CSCAPES is developing novel algorithms and software for a range of combinatorial problems. Capabilities are being deployed via:
  - The load-balancing and parallelization toolkit Zoltan
  - Automatic Differentiation tools such as OpenAD and ADOL-C
- Partitioning and Load-balancing:
  - Collaboration with COMPASS/SLAC on particle-in-cell simulation using PIC3D
  - ITAPS has integrated Zoltan into their software framework
  - Collaboration with TOPS to improve scalability of SuperLU parallel solver
  - Developed new sparse matrix partitioning algorithms for matrix-vector product; an important kernel in many SciDAC applications
- Automatic Differentiation:
  - Collaboration with NREL and TOPS researchers to provide first and second order sensitivities of ODEs for metabolic networks in the **Green Energy** SciDAC project.

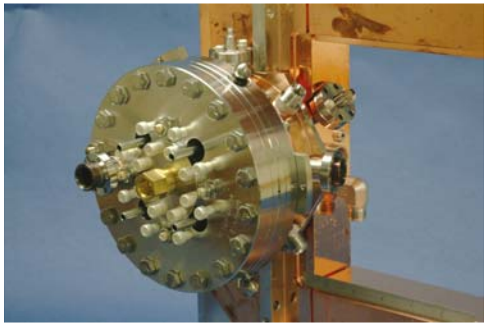

RF gun from SLAC (parallel simulation uses Zoltan)

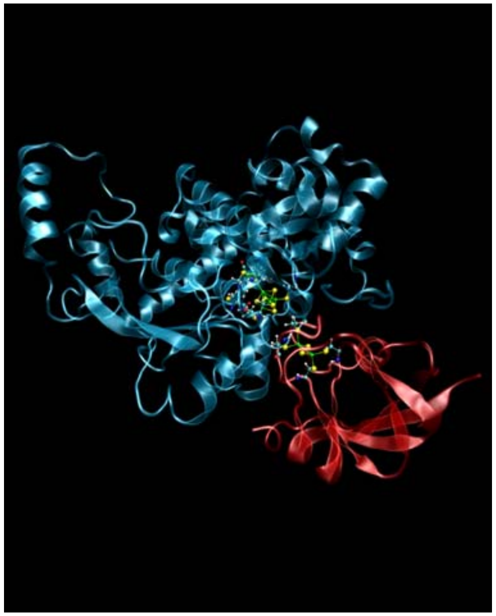

Protein from green energy project (courtesy Chris Chang, NREL)

CSCAPES
Graph Coloring for Computational Science

- "Matrix compression" in efficient derivative computation and concurrency discovery in parallel scientific computing call for the solution to various graph coloring problems
- CSCAPES is developing novel algorithms and software for a range of coloring problems.
- Examples of recent applications where CSCAPES coloring technology enabled orders of magnitude reduction in execution time and storage:
  - Computational model for chromatographic separation (Simulated Moving Bed process). Examples of industrial applications of SMB include:
    - Purification of biochemicals in pharmaceuticals
    - Purification of hydrogen from refinery gases
  - Optimization of power flow in electric network
- Collaborating institutions outside SciDAC:
  - Georgia Tech and Carnegie Mellon University (applications in chemical engineering); Technical University of Dresden, Germany (ADOL-C); University of Bergen, Norway (parallel algorithms).

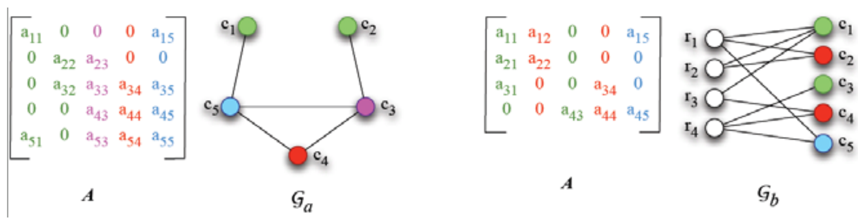

***Structurally orthogonal partition** of a matrix and its representation as **distance-2 coloring** of an appropriate graph*

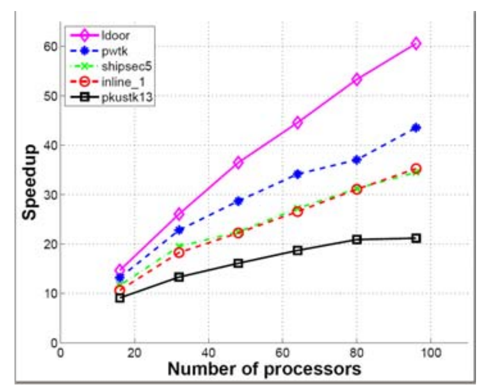

Performance of a recently developed parallel distance-2 coloring algorithm

SciDAC 2
Computer Science

Centers:

- **CScADS – Center for Scalable Application Development Software**
  - **Research on technical challenges associated with emerging multi-core architectures**
  - **Development of open-source shared software infrastructures**
- **TASCS – Technology for Advanced Scientific Component Software**
  - **Enhance core Common Component Architecture software environment with an emphasis on improving usability**
  - **Build component ecosystem to provide more off-of-the shelf**

TASCS
Techology for Advanced Scientific Component Software

Institutes:

- **PDSI – Petascale Data Storage Institute**
  - **Identify, and implement solutions for storage capacity, performance, concurrency, reliability, availability and manageability problems arising from petascale bytes of data**
- **PERI – Performance Engineering Research Institute**
  - **Focused on**
    - **Performance modeling and prediction**
    - **Automatic performance optimization**
    - **Performance engineering of high profile applications.**

pdsi

PERI

PERI
"Tiger Teams"

The Performance Engineering Research Institute (PERI) project selected two **Joule Metric** science projects for focused collaborative code tuning efforts:

S3D, led by Jacqueline Chen at LLNL:

- Compressible reacting Navier-Stokes equations.
- High fidelity numerical methods -- 8th order finite difference, 4th order explicit R-K integrator.
- Hierarchy of molecular transport models.
- Received 2007 INCITE award (6M CPU-hrs).

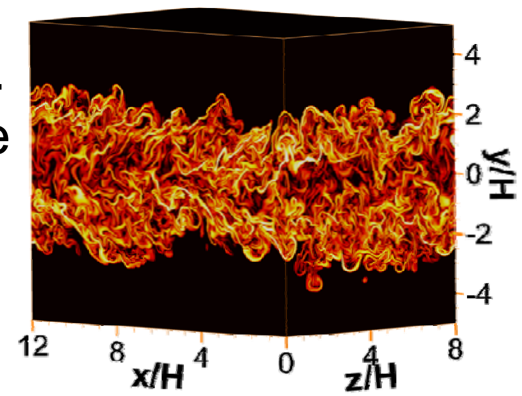

GTC, led by Zhihong Lin at UC Irvine:

- 3D particle-in-cell code to study microturbulence in magnetically confined fusion plasmas.
- Solves the gyro-averaged Vlasov equation.
- Gyrokinetic Poisson eqn solved in real space.
- Low noise δf method.
- Global code (full torus as opposed to only a flux tube).

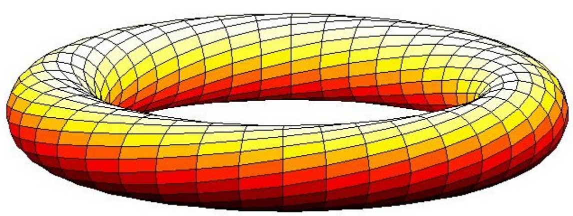

Graphics: Thanks to J. Chen, W. W. Lee and Z. Lin.

PERI-Application Team
Collaborations Results

S3D:

- In one key section of code, performance tools found that the compiler unrolled by the second index, whereas the first gave a higher L1 cache hit rate. This yielded a 7.5% overall performance increase.
- A change to “getrates” enabling pseudo-vectorization of an inner loop with exp calls resulted in 10% overall performance increase on IBM P4.
- Several changes to the loop structure of another section of code, using performance tools, resulted in a 7% overall improvement.
- Overall, performance increased by 13% on Cray XT3/4.

Thanks to B. de Supinski, J. Mellor-Crummey, S. Moore, S. Shende, N. Wright, G. Marin and numerous others.

CScADS & PERI @ Rice

Optimization of GTC

- GTC: simulates turbulent plasma in tokamak reactors
  - 3D particle-in-cell code; 1D decomposition along toroidal direction
    - charge: deposit charge from particles to grid points
    - solve: compute the electrostatic potential and field on grid points
    - push: compute the force on each particle from nearby grid points

- Data locality optimization
  - restructured program data and loops
  - adaptively reorder ions at run time
    - at run time, locality degrades gradually as ions in the plasma become disordered
    - periodic particle reordering restores locality and performance

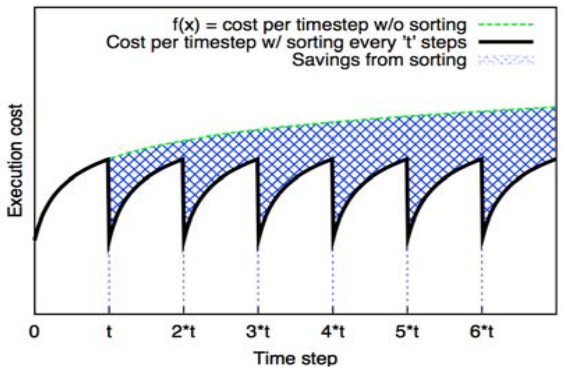

- Reduces GTC execution time by 21% on Cray XT4 systems

HPCToolkit Performance Tools

- **Collaborative between**
  - **PERI**
  - **CScADS**
- **Performance tools for measurement, analysis, and attribution of performance problems on petascale systems**
- **FY08 Accomplishments**
  - **Ability to measure performance of fully-optimized parallel codes on both Cray XT and Blue Gene systems**
  - **Pinpoint and quantify inefficiencies and scalability bottlenecks**
  - **Helping SciDAC teams analyze their code performance**
    - **Atmospheric modeling, lattice quantum chromodynamics, molecular dynamics, many Fermion dynamics, turbulent combustion, structured adaptive mesh refinement, among others**

CScADS

Summer Workshops

- Goals
  - Foster collaboration between computer scientists and application teams
  - Identify important open problems and challenges for leadership computing
  - Brainstorm on promising approaches to open problems
  - Identify infrastructure needs to address key challenges
  - Extend existing software components to address new challenges
  - Collaborate on design of sharable software components
- Topics for July 2008 Workshops in Snowbird UT
  - Application engagement
    - Leadership class machines, Petascale Applications, and Performance
    - Scientific Data Analysis and Visualization for Petascale Computing
  - Technology development
    - Autotuning for Petascale Systems
    - Performance Tools for Petascale Computing

TASCS Babel

HPC Middleware
for Large Scientific Applications

R&D 100
2006

"The world's most rapid communication among many programming languages in a single application."

Performance (in process)

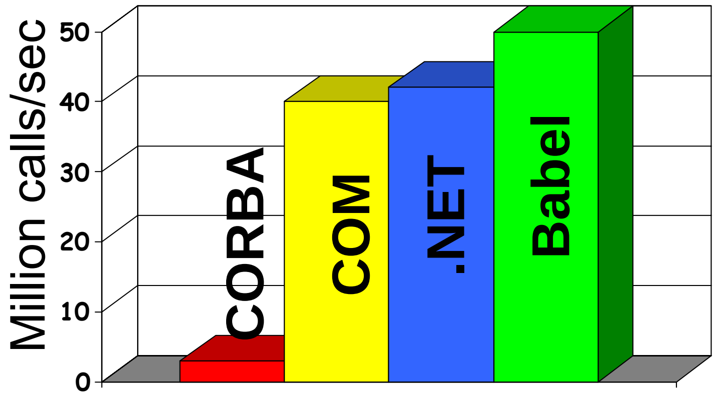

Recent Highlights

- Babel RMI and Coop solve multi-scale material science problems using adaptive sampling (Barton 2008, LLNL)
- New Babel Fortran 2003 binding for structured data types (Muszala, Epperly, and Wang 2008, Tech-X & LLNL)
- FACETS SciDAC makes first tightly integrated core-edge Tokamak fusion reactor simulation using Babel (Cary *et al* 2008)

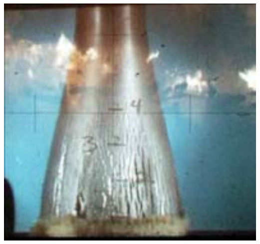

HE driven pipe

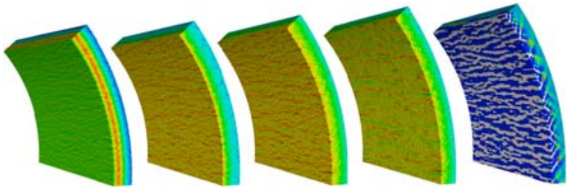

Computational fragmentation calculations qualitatively match experimental findings (SciDAC PI Meeting 2008, Barton)

TASCS
Techology for Advanced Scientific Component Software

BABEL

TASCS

Creating a Community Groundwater Simulation Model using Components

TASCS
Techology for Advanced Scientific Component Software

- Simulation created from pluggable components enables community participation.
- CCA's Bocca tool incorporates existing code into a flexible component environment.
- Bocca gets the component “glue code” out of the way of the scientist/developer.

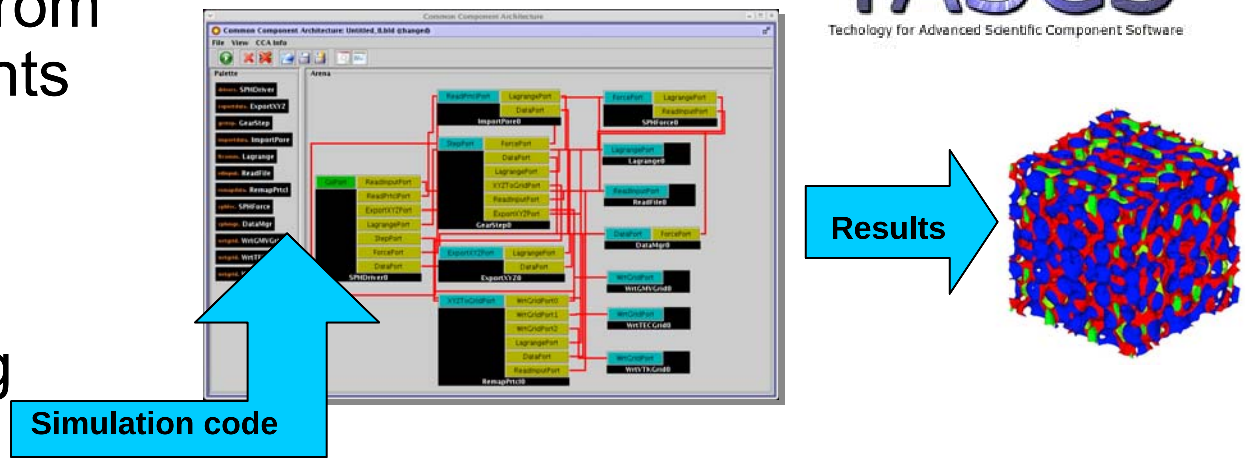

GWACCAMOLE Groundwater Model

- Uses a particle-based algorithm scalable to Leadership Class machines
- Tracking pollutants entrained in subsurface groundwater.
- Models for pore-scale transport and reaction of contaminants in groundwater

Thanks to Bruce Palmer (PNNL) and the GWACCAMOLE team.

SciDAC 2

Data Management/Visualization

Centers:

- **SDM – Scientific Data Management**
  - **Enhance and extend scientific data management framework for more interactivity and fault tolerance; better parallelism in data analytics and greater functionality interactions with local and remote parallel file systems**
- **VACET – Visualization and Analytics Center for Enabling Technologies**
  - **Creation and deployment of scientific visualization and analytics software technology to increase scientific productivity and scientific insight**

Institutes:

- **UVIS – Ultrascale Visualization**
  - **Develop comprehensive a portable parallel visualization suite to enable scientific discovery at the petascale**

U.S. Department of Energy

Office of Science

SDM
Dynamic Real-Time Framework for Monitoring of Simulations

- Problem
  - Large scale simulations may not converge or maybe unstable
  - By finding out instabilities in real time, it is possible to stop useless simulations, saving many hours of compute resources
- Achievements
  - Developed a workflow-dashboard monitoring framework, that generates images and movies dynamically while the simulations are running
    - Used in production in CPES fusion Project (Center for Plasma Edge Simulations)
      - Automatic code-coupling coordinates runs on two machines, one running simulation, and one checks for stability. If unstable it stops simulation, calculates new parameters and restarts
    - Adapted to new codes in combustion and astrophysics simulations

SDM CENTER

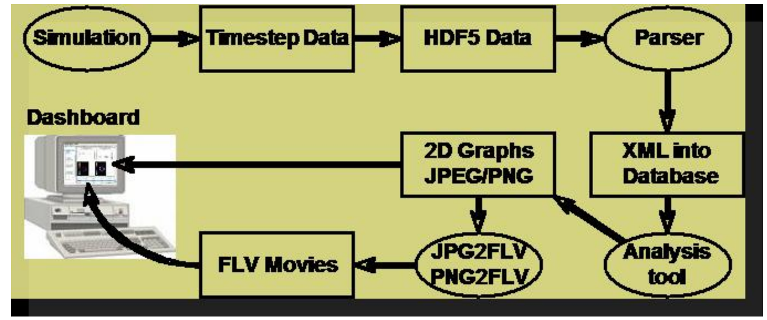

Monitoring workflow schematic using Kepler workflow system

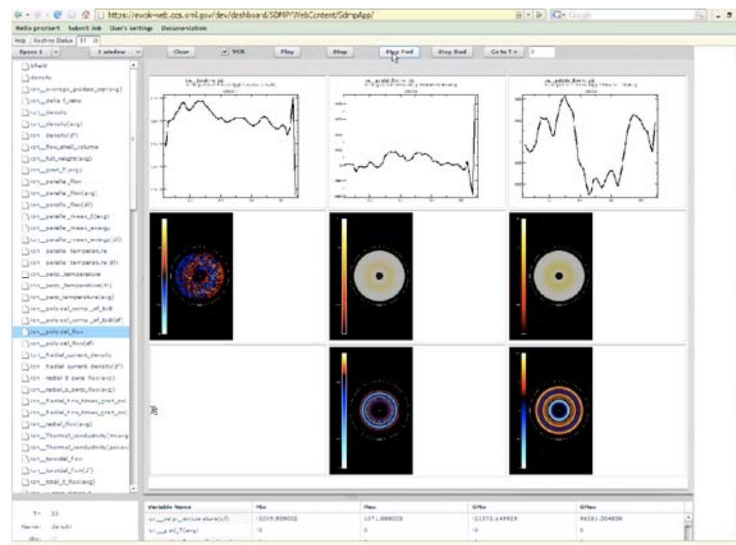

Dashboard, where dynamically generated images and graphs can be viewed and compared in multiple panels over the web

SDM's FastBit

Advancing the Frontier of Data Search

- Attributes
  - innovative bitmap indexes for searching billions of scientific data values
  - improves the search speed by 10x – 100x of times with
    - Uses *patented* fast compression techniques
  - handles large scale numeric data for scientific applications
    - breaks the limitation on the number of distinct values by using binning and efficient compression
- Successfully used in many applications
  - Searching for individual records: Grid Collector for high-energy data analysis – needles in a haystack
  - Searching for collective features: Query-Driven Visualization has proven to be very effective for analysis of methane jet flame shown on the right
  - Other application domains: real-time analysis of network intrusion attacks, fast tracking of combustion flame fronts over time, accelerating molecular docking in biology applications, and more.

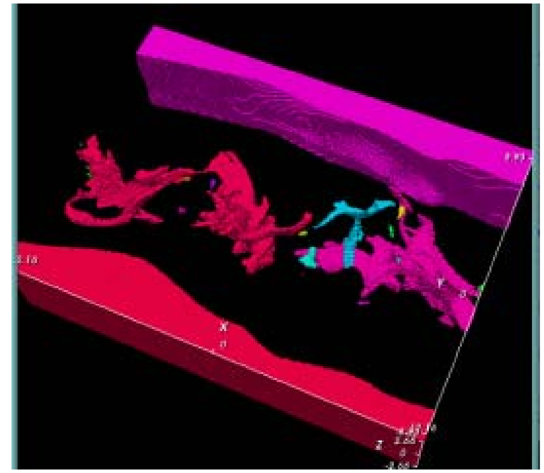

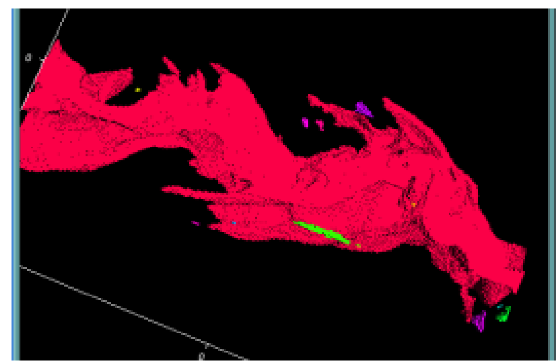

2008 R&D 100

Example where FastBit makes it possible
to explore flame jet (over million of cells) in real-time

VACET

Production Visual Data Analysis Software

- VACET provides production quality, parallel capable visual data analysis software infrastructure
- Adopted as the “de-facto” standard for project-wide use by multiple SciDAC projects: APDEC, CAC (more to come).
- Benefits:
  - “Buy rather than build.” ADPEC no longer needs to expend its resources on AMR visualization software development, maintenance.
  - VisIt makes effective use of parallel computing infrastructure for performing visual data analysis – replaces older, slower and serial tools with one that runs on all modern HPC platforms.
  - VACET technologies deployed at DOE’s Open Computing facilities: a good example of successfully bridging the gap across research, development and production deployment.

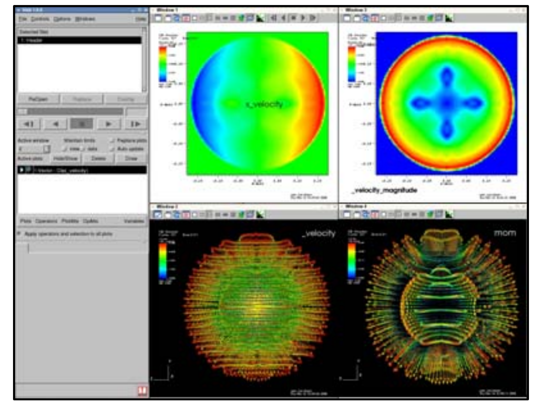

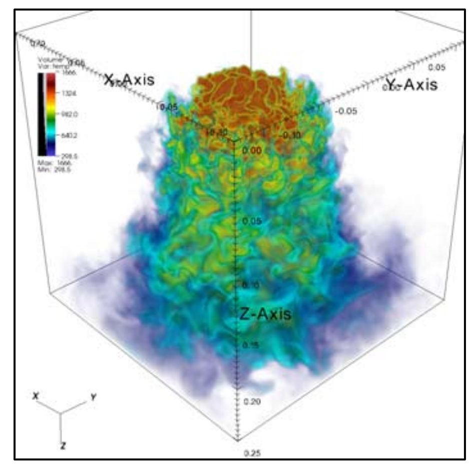

VACET

VACET
Research Enables Scientific Discovery

- Topological analysis, an award-winning part of VACET's research portfolio, enables new insight into scientific phenomena not possible with just an image or a movie.
- It provides new insights for understanding the mechanisms of extinction and reignition in turbulent, non-premixed flames (J. Chen, SNL-CA).
- It provides the ability to quantify and analyze the relationship between turbulence and the stability of the combustion process (J. Bell, LBNL).

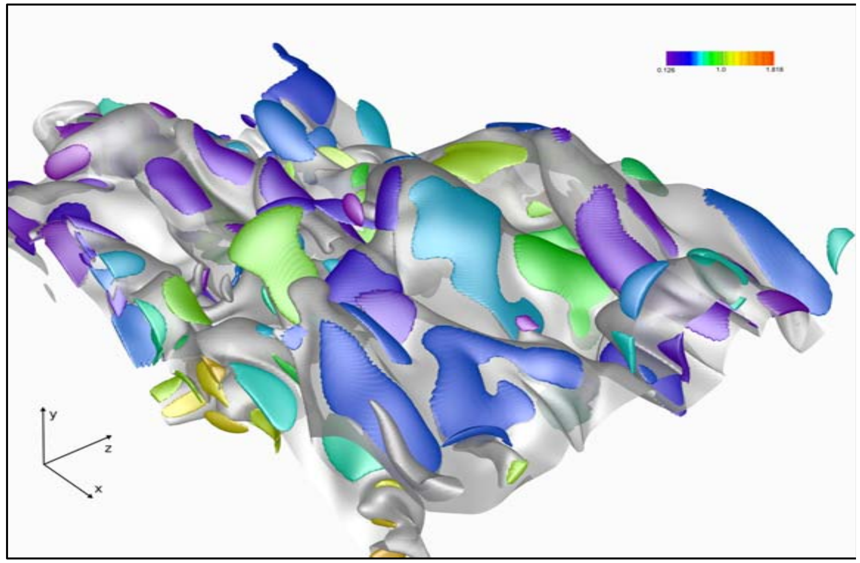

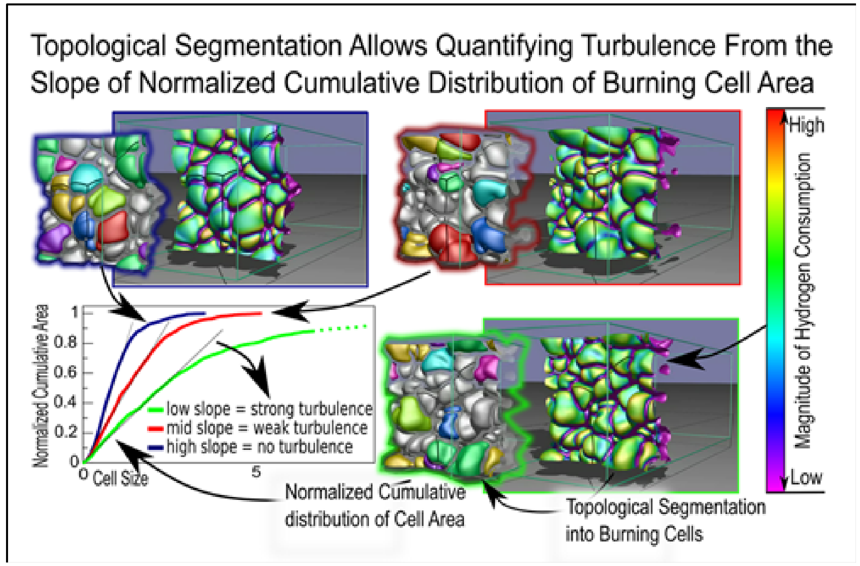

VACET

VACET

Remote Visualization Software Infrastructure

- New capability: run unmodified visualization applications on distributed memory parallel computing platforms, deliver image results to one or more remote viewers.
- Benefits:
  - Allows remote access to centrally located data, parallel visual computational infrastructure via iPhone, workstation, etc.
  - Increasingly important as data grows larger – must be visualized/analyzed “in place”
  - General purpose technology applicable to virtually all existing visualization technology
  - Open Source software
  - Being deployed at NERSC/LBNL, LLNL, CCS/ORNL
- Project participants: Tungsten Graphics, Inc. (SBIR), LLNL, ORNL, LBNL.

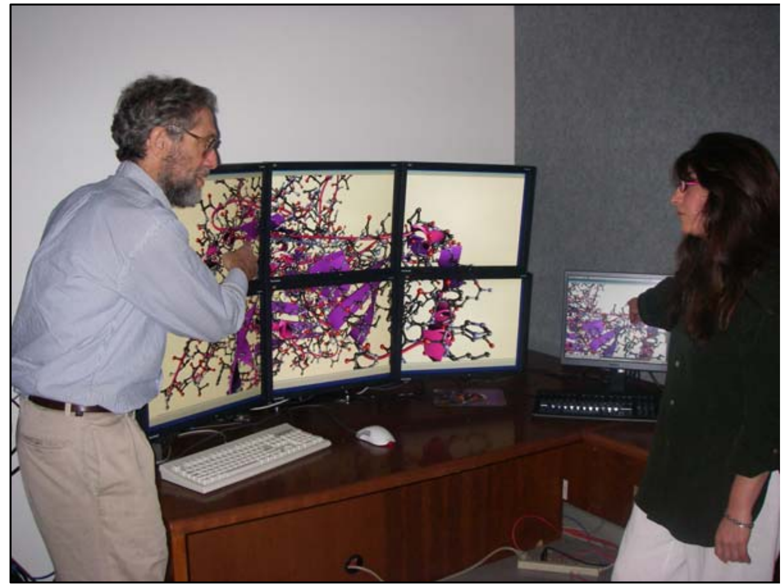

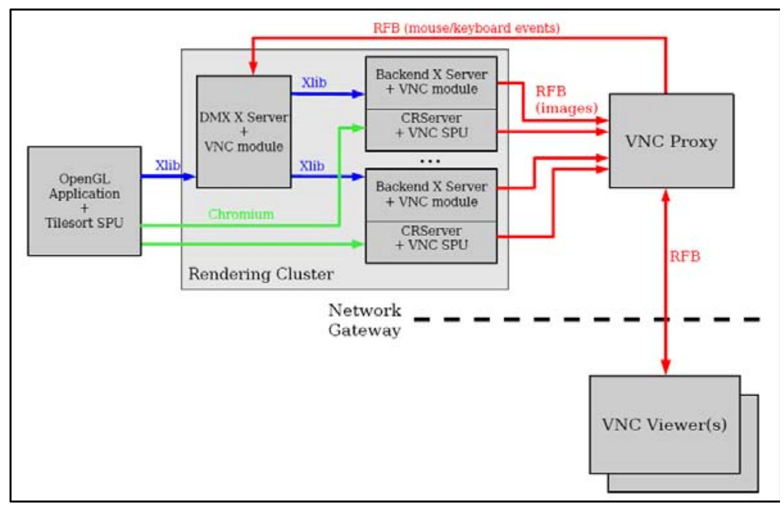

VACET

Accomplishments:

The Ultravis team has:

- **produced basic research results advancing algorithms, interfaces, and architectures for visualization in terms of visual clarity and richness, interactivity, and scalability;**
- **organized tutorials, panels, and workshops along with web-based shared space fostering interaction within the community and between visualization experts and application scientists;**
- **offered undergraduate and graduate students, postdocs, and junior faculty members access to the most challenging data analysis and visualization applications and leading-edge computing facilities to solve problems at unprecedented scale and complexity.**

U.S. Department of Energy

Office of Science

UltraVis

SciDAC Institute for Ultrascale Visualization

Scalable Visualization on Leadership Systems

- The Ultravis team has demonstrated scalable performance on the Argonne Leadership Computing Facility Blue Gene/P system using up to 32,000 processing units, largest ever used in a study.
- This new level of scalable visualization is achieved by the introduction of a new parallel rendering algorithm and the effective employment of parallel I/O support
- Scalable visualization is key to capturing hidden, complex structures in peta/exa-scale
- Publications: SC08 and 2008 Parallel Graphics and Visualization Symposium

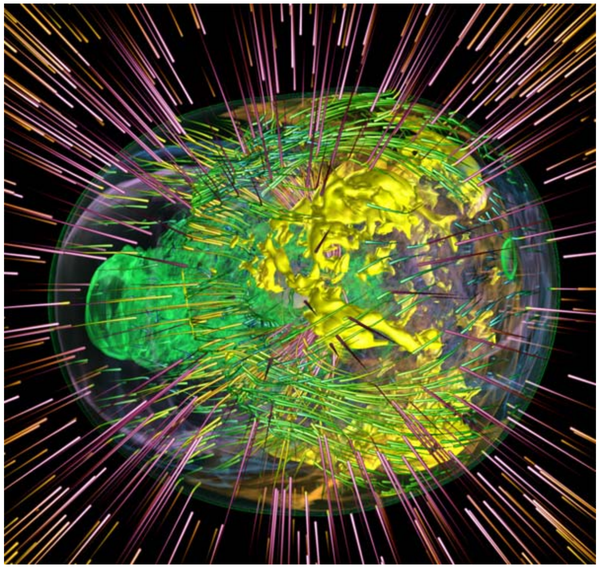

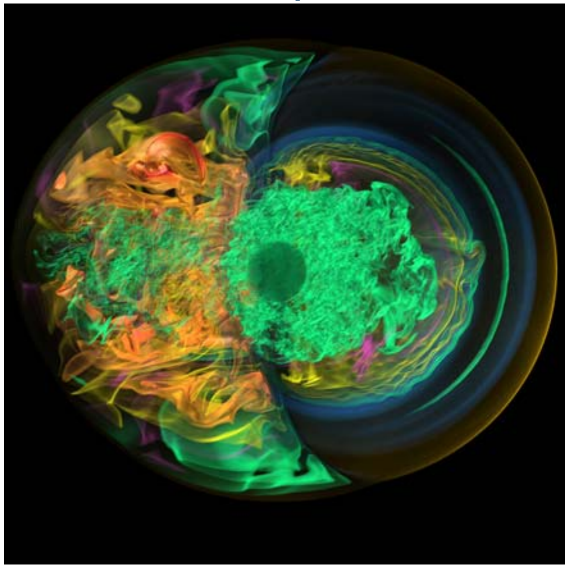

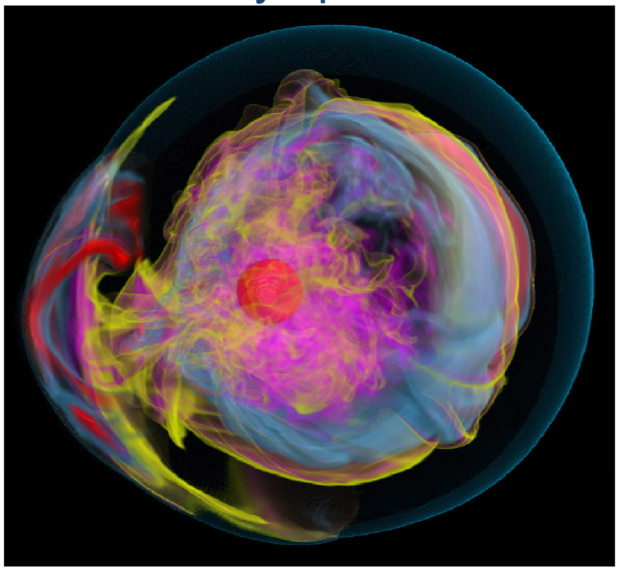

Supernova simulation, John Blondin, North Carolina State University

U.S. Department of Energy
Office of Science

SciDAC 2
Collaboratories

- **ESG – Earth Systems Grid**

  - Develop and deploy multiple types of model and observational climate data
  - Provide more ESG access and analysis services
  - Enhance interoperability between common climate analysis tools

- **CEDPS – Center for Enabling Distributed Petascale Science**
  - Design, develop and deploy services and tools for data placement within a distributed high-performance environment
  - Construct and test scalable services for processing of computation and data analysis requests from remote clients

- **Open Science Grid**
  - Maintain and operate a Petascale nationwide distributed facility that can grow to provide thousands of users at universities and DOE laboratories throughout the U.S.
  - Engage train and include new researchers
  - Integrate and deploy new IT technologies in response to explicit needs and evaluate in a real-life setting.

ESG:
The World's Source for Climate Modeling Data

U.S. Department of Energy
Office of Science

Main ESG Portal

198 TB of data at four locations

- 1,150 datasets
- 1,032,000 files
- Includes the past 6 years of joint DOE/NSF climate modeling experiments

8,000 registered users

Downloads to date

- 60 TB, 176,000 files

ESG usage: over 500 sites worldwide

CMIP3 (IPCC AR4) ESG Portal

35 TB of data at one location

- 74,700 files
- Generated by a modeling campaign coordinated by the Intergovernmental Panel on Climate Change
- Data from 13 countries, representing 25 models

2,000 registered projects

Downloads to date

- ~1/2 PB
- 1,300,000 files
- 500 GB/day (average)

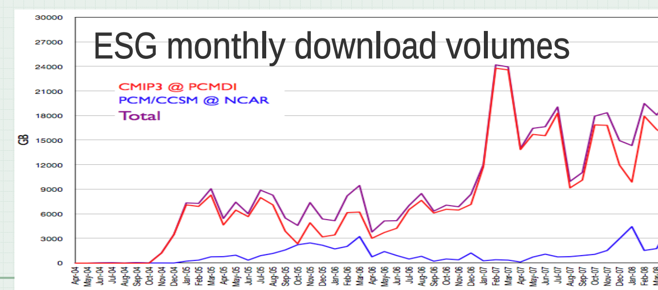

400 scientific papers published to date based on analysis of CMIP3 (IPCC AR4) data

Open Science Grid

Delivered to physics simulation & analysis goals in last "Data Challenge" before start of data taking at CERN:

- Data transfers across the globe, averaging up to 120 Terabytes a day over a month (>100 links)

- Job throughputs of >40,000/day.

- Supported 450 active users running 1.5M jobs a month across 30+ Tier2 sites.

- Simultaneous Cosmic Ray running of detector and fully distributed data analysis.

- Simultaneous simulation production of >10Mevents/week.

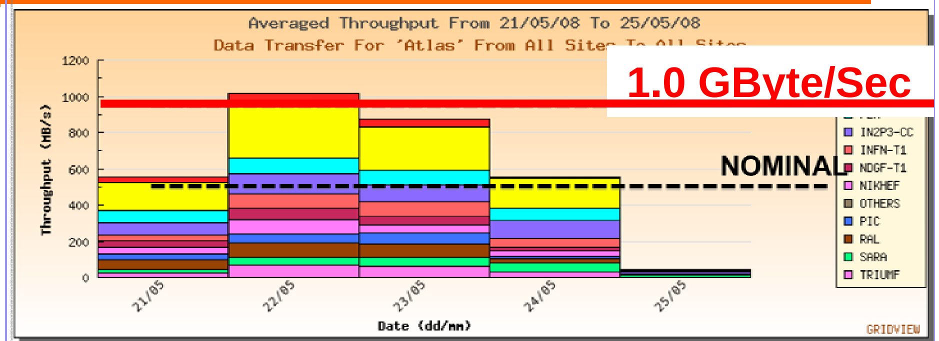

End-to-End Data Transfer Throughput (ATLAS)

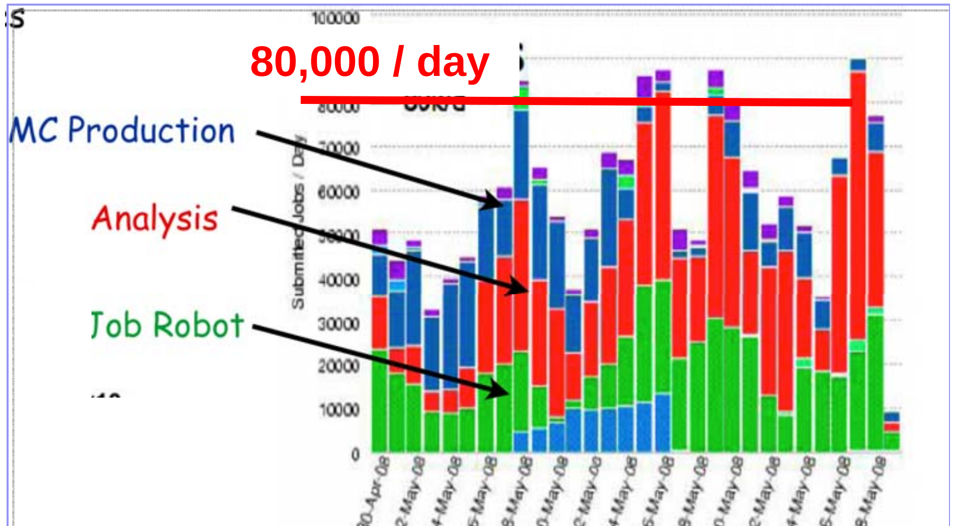

Mix of Production and Analysis Jobs(CMS)

OSG
Value from Integrated Software Distributions

2 major new versions of integrated software suite deployed for all OSG Science groups (>15).

- ✓ Used by >70 DOE Labs and Universities;
- ✓ Adopted by >4 Campus-wide Grids;
- ✓ Being evaluated by Structural Biology Grid at Harvard;
- ✓ Subset supported for LIGO Data grid, TeraGrid, UK grid and EU EGEE.

Contributing to broader SciDAC impact through hardening and maturing of software tools developed by the SciDAC and other ASCR programs

Interoperation across distributed systems in support of Global science: LIGO, LHC, Pragma.

Increased customer satisfaction and stability (immediate adoption, reduced number of problems).

What it takes to Release Software for Production Grids

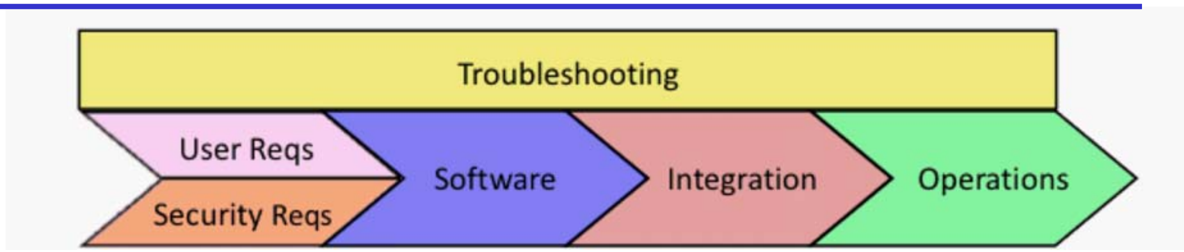

- ✓>65 components in Virtual Data Toolkit contributed from >15 software development groups.
- ✓ Automated nightly builds on >15 Linux OS variants, MacOSX, AIX;
- ✓ 2 month iterative system and application testing across >13 sites.

CEDPS

Connecting Science Facilities to the World

SciDAC-supported GridFTP enables secure, reliable, high-speed data movement over wide-area networks—an essential capability for the effective use of unique US science facilities

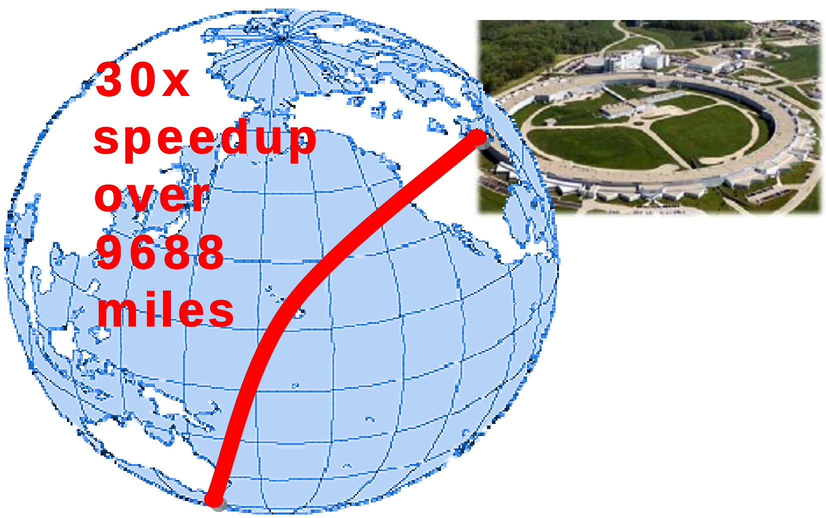

One terabyte moved from an Advanced Photon Source tomography beamline to Australia, at a rate 30x faster than standard FTP

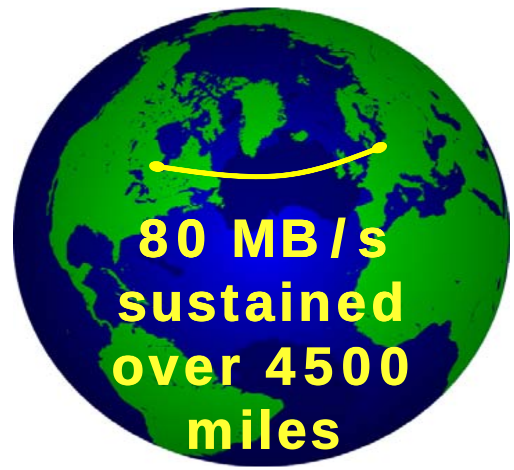

1.5 terabyte moved from University of Wisconsin, Milwaukee to Hannover, Germany at a sustained rate of 80 megabyte/sec

CEDPS
On-Demand Computing for DOE Science

Example: computing for STAR experiment

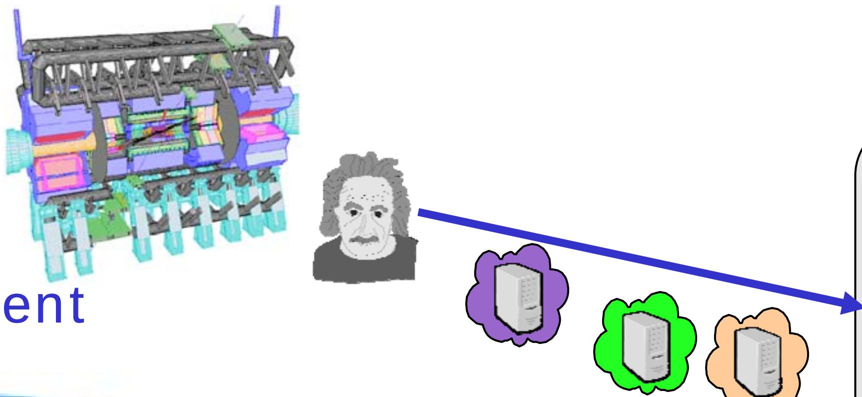

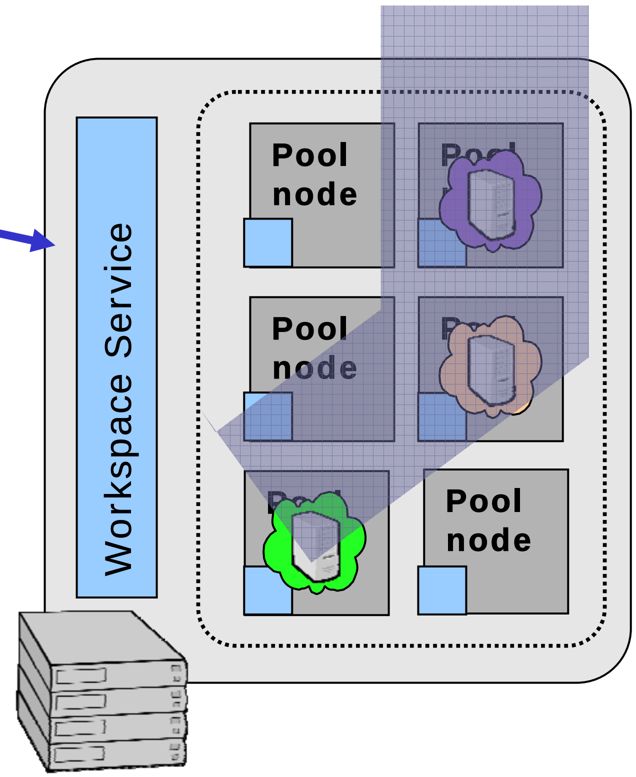

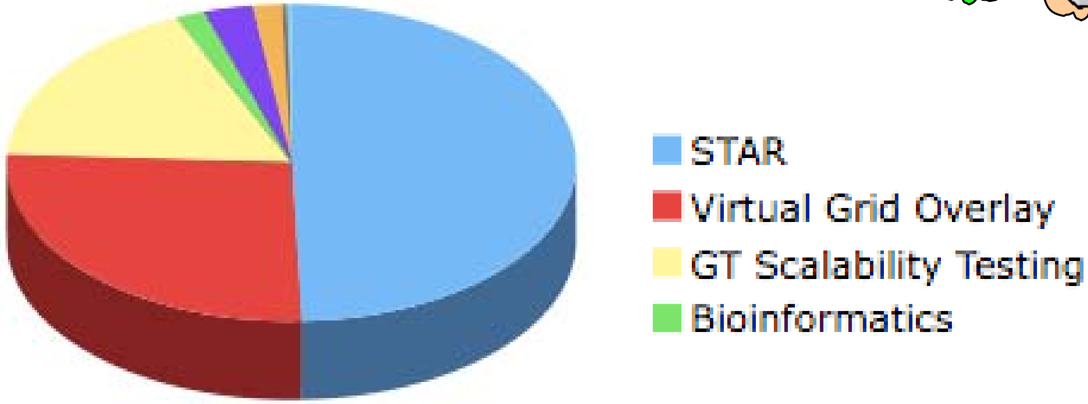

SciDAC-supported "Nimbus" enables one-click creation of virtual clusters configured with the specialized software required for scientific simulation and data analysis—"Cloud computing" for science

CEDPS

Workflow Automation at DOE Facilities

Advanced Photon Source

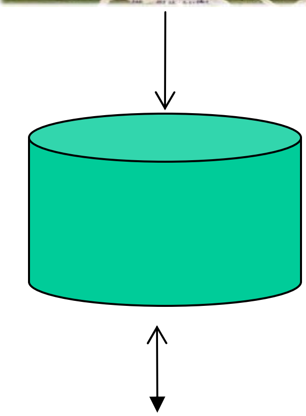

Storage
Metadata
Analysis
Visualization

Automation
Reproducibility
Security
Reusability

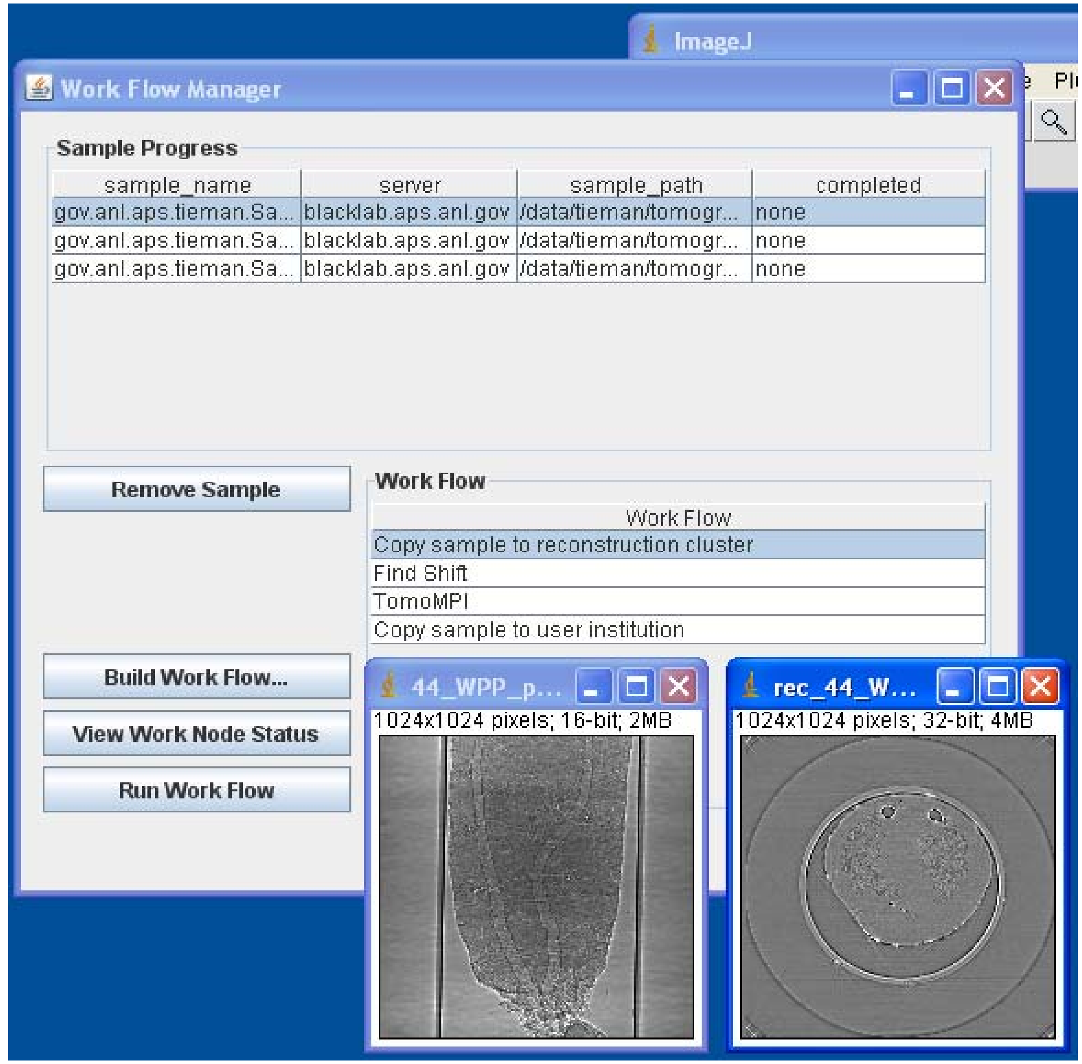

U.S. Department of Energy

Office of Science

SciDAC 1 vs SciDAC 2

Funding

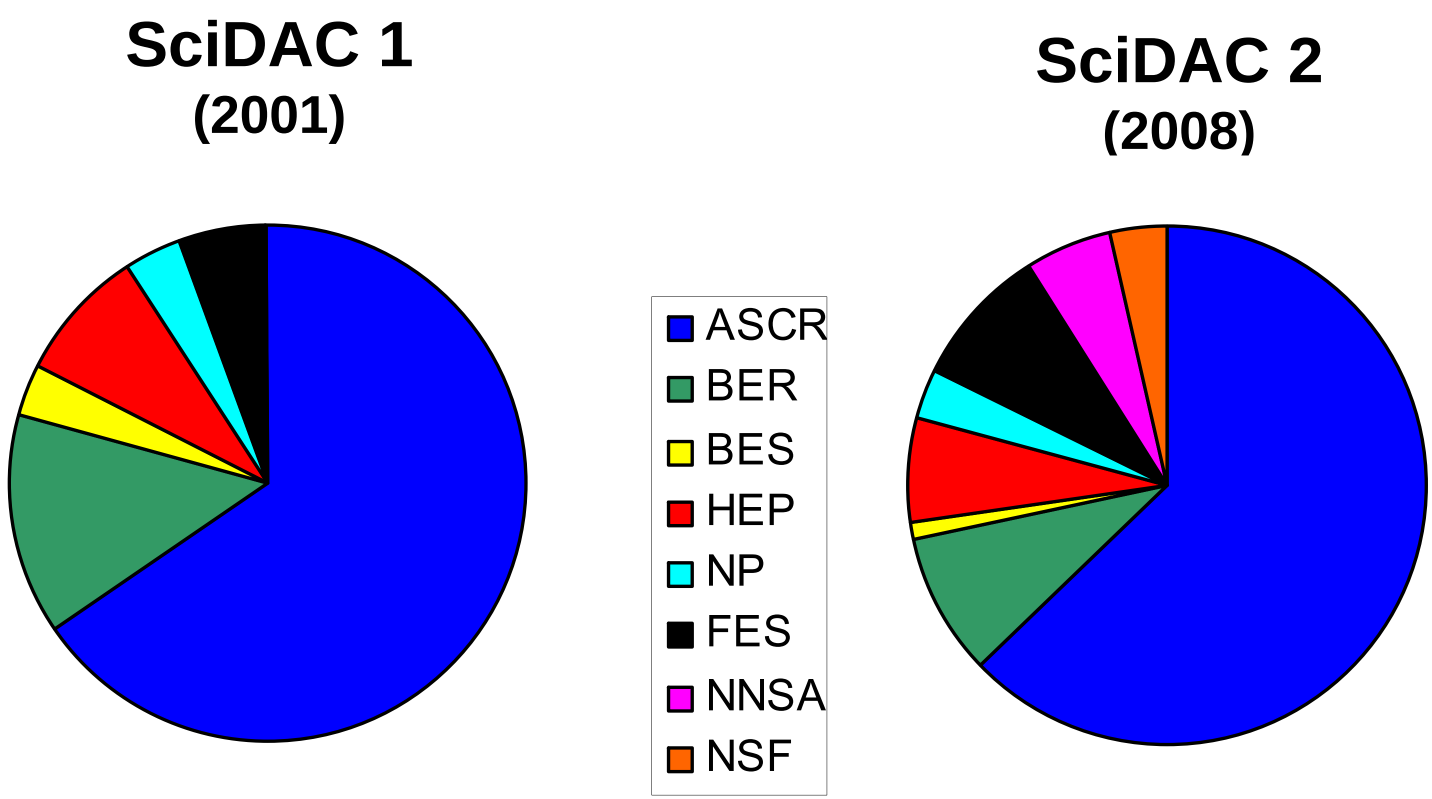

Total- $57.3 M

Total- $79.9 M

U.S. Department of Energy

Office of Science

SciDAC 2

Conferences

Registered Participants	
2005	285
2006	300
2007	310
2008	348

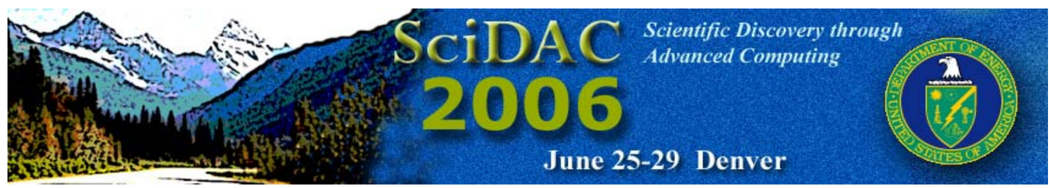

Bringing together Leading Computational Scientists

U.S. Department of Energy

Office of Science

SciDAC 2

SciDAC Review

A Premier Outreach Magazine

Highlighting Scientific Discovery and the Role of High End Computing

SciDAC 2

Outreach Center

- A central resource for inquiries and technical information about SciDAC
- Objective
  - Develop portal services
  - Expand outreach services to industry
  - Prepare for peer review
- Activities
  - Help SciDAC projects organize/develop/update/test software (major interactions with VACET, PERI, OSG & UNEDF)
  - Respond to inquiries

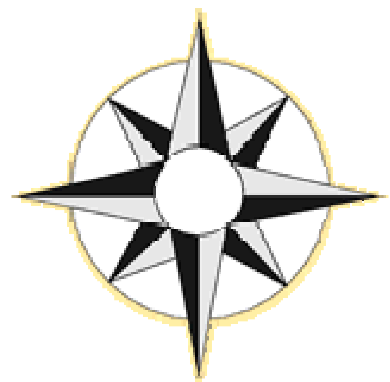

Outreach.scidac.gov

U.S. Department of Energy
Office of Science

E3 Initiative

Transforming Energy, the Environment and Science through simulations at the eXtreme Scale

http://www.science.doe.gov/ascr/ProgramDocuments/TownHall.pdf

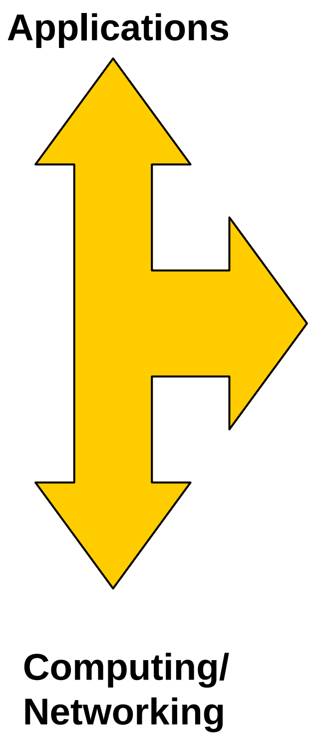

SciDAC X

- Materials, chemistry
- Nuclear physics
- Fusion science
- Combustion
- Alternative fuels & transportation
- Carbon management
- Groundwater
- Astrophysics
- Computational Biology
- High-energy physics
- Advanced Fuel Cycles
- Climate

- Institutes for applied mathematics and computer science
- Software & algorithm development centers
- Cyber environment for open, but secure science

INCITE X
NERSC Allocation

Multi-peta flop and Exaflop Compute Resources	300-1,000 PF, 20-80MW
ESnet	Best-in-class, reliable, robust, on demand bandwidth and network services

HPC Resources

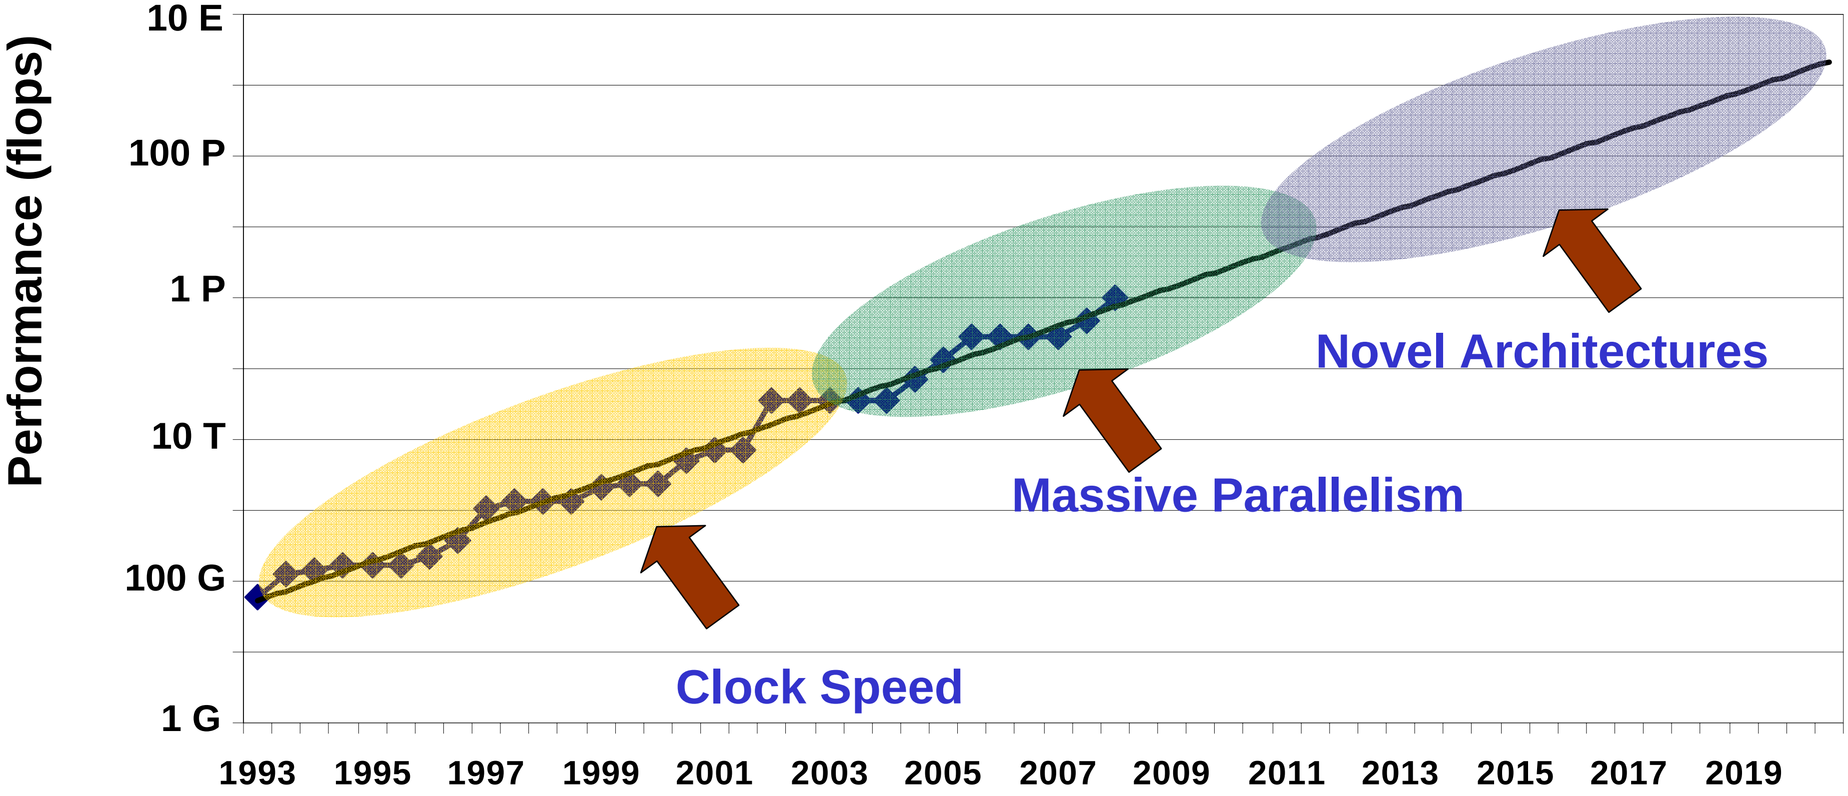

An Era of Challenge

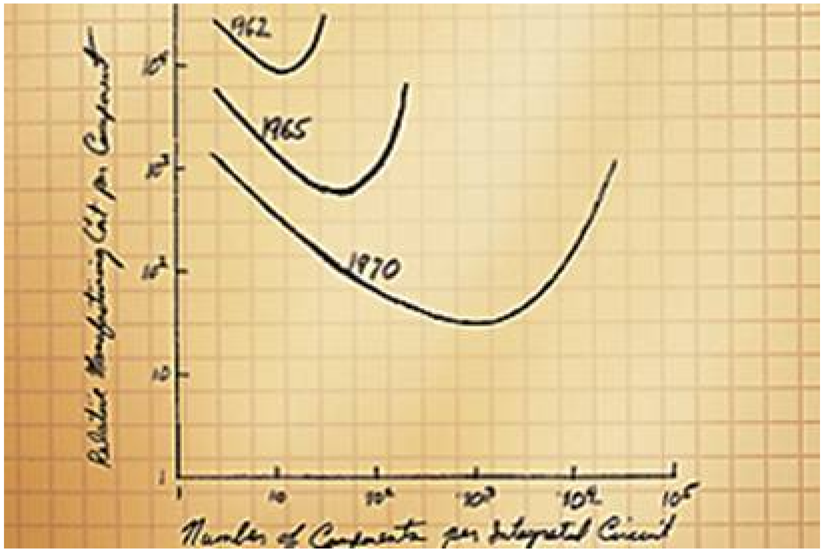

Moore's original graph predicting Moore's Law in 1965. Chip capacity will double every 2 yrs.

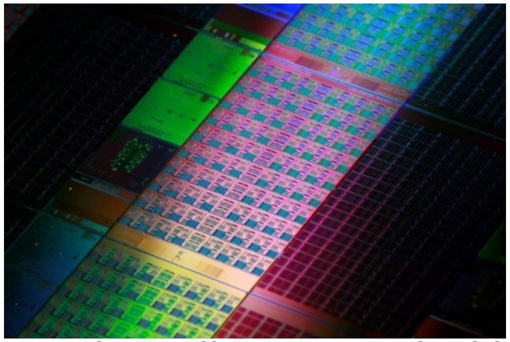

Intel Teraflops Research Chip

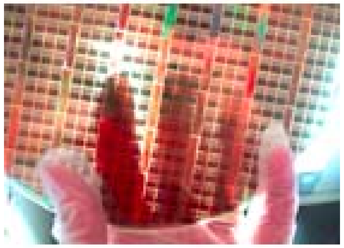

IBM Stacked Chip

- **Unpredictable evolution of hardware**
- **Multilevel and heterogeneous parallelism; memory hierarchies**
- **Programming models must work at scale (numbers of core, lines of code, numbers of components)**
- **Managing data, simulation, experimental and observed**
- **Communications: synchronous → asynchronous**
- **Reliability**

It's not just extreme scale, it's also extreme complexity

U.S. Department of Energy
Office of Science

SciDAC X

Will you be ready for extreme scale computing?

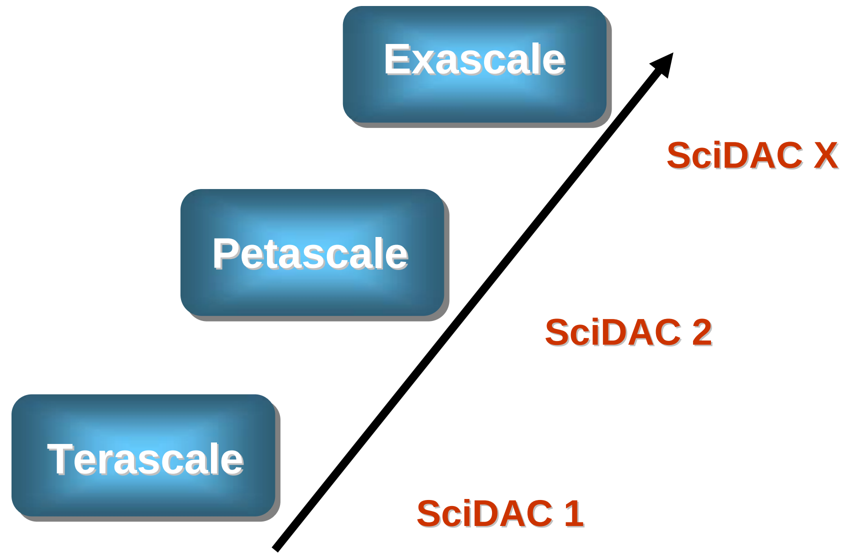

I climb the "Hill of Science,"
I "view the landscape o'er;"
Such transcendental prospect,
I ne'er beheld before!

Emily Dickinson