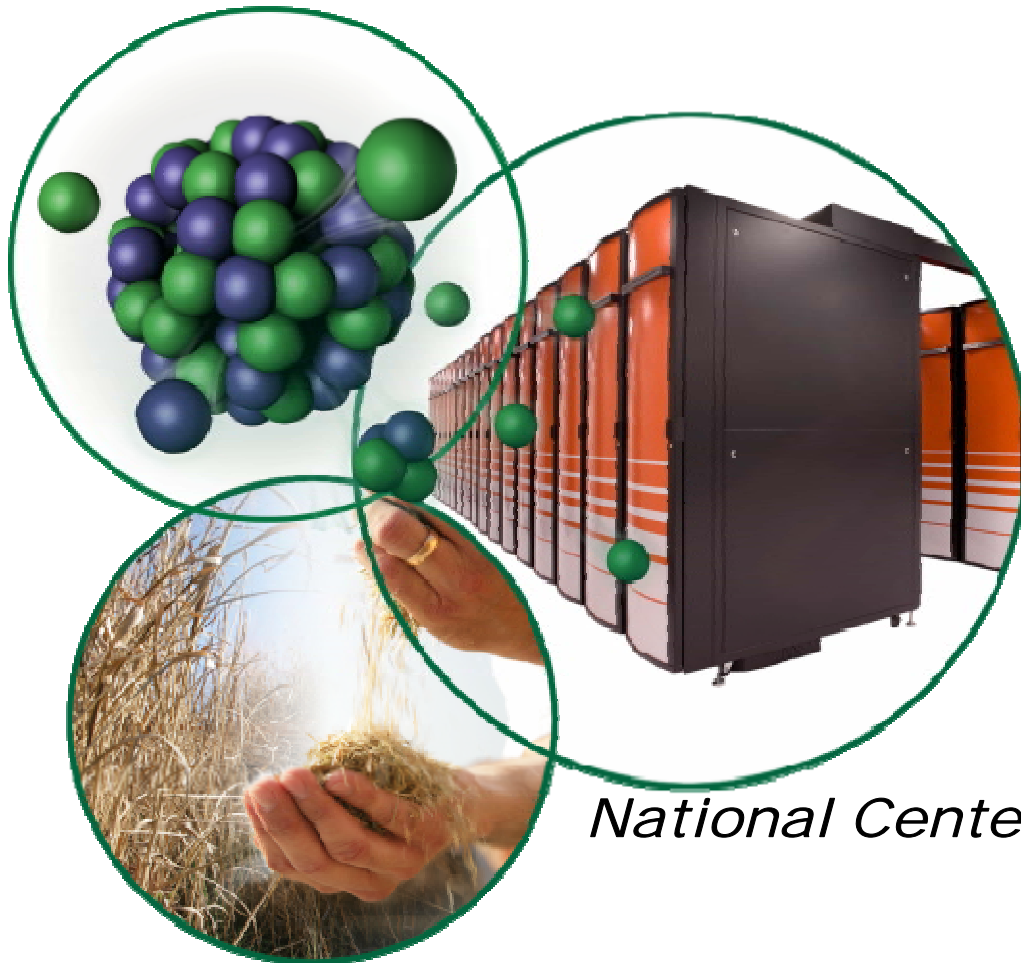


# *Science (and Engineering) at the Leadership Computing Facility at Oak Ridge National Laboratory*



*Douglas B. Kothe  
Director of Science  
National Center for Computational Sciences*

# *Leadership Computing Facility Mission*

**Our mission is to enable breakthrough science by fielding the most powerful capability computers for scientific research; building the required infrastructure to facilitate user access to these computers; working with the DOE to select a few time-sensitive problems of national importance that can take advantage of these systems; and joining forces with these teams to deliver breakthrough science.**


# *We Must Impact Science*

## Enable Breakthroughs; Accelerate Discoveries


- **Provide the nation's most powerful open resource for capability computing**
- **Follow a well-defined path for maintaining national leadership in this critical area**
- **Attract the brightest talent and partnerships from all over the world**
- **Deliver cutting-edge science relevant to the missions of DOE and key federal and state agencies**
- **Unique opportunity for multi-agency collaboration for fundamental science, applied science, and engineering**




# Advancing Scientific Discovery




Resolved decades-long controversy about validity of 2D Hubbard model in predicting behavior of high-temperature superconducting cuprate planes




300K-atom models of cellulase enzyme on cellulose substrate reveal interior enzyme vibrations that influence reaction rates converting cellulose to ethanol




Addition and intercomparison of carbon-land models in new climate model is resolving key processes for carbon sources & sinks



Turbulence chemistry revealed in study of lifted turbulent  $H_2$ /air jet flames in ignitive coflow relevant to diesel engines and gas turbines



Instability of supernova shocks was discovered directly through simulation and core collapse pulsar mechanism was explained




Providing increasing assurance that RF power will effectively heat ITER

# Some Science Drivers

Science Domains	Science and Engineering Driver
Accelerator Physics	Optimize a new low-loss cavity design for the ILC
Astrophysics	Explosion mechanism of core-collapse supernovae and Type Ia supernovae
Biology	Can efficient ethanol production offset the current oil and gasoline crisis?
Chemistry	Catalytic transformation of hydrocarbons; clean energy & hydrogen production and storage
Climate	Predict future climates based on scenarios of anthropogenic emissions
Combustion	Developing cleaner-burning, more efficient devices for combustion.
Fusion	Plasma turbulent fluctuations in ITER must be understood and controlled
High Energy Physics	Find the Higgs particles thought to be responsible for mass, and find evidence of supersymmetry
Nanoscience	Designing high temperature superconductors, magnetic nanoparticles for ultra high density storage
Nuclear Energy	Can all aspects of the nuclear fuel cycle be designed virtually? Reactor core, radio-chemical separations reprocessing, fuel rod performance, repository
Nuclear Physics	How are we going to describe nuclei whose fundamental properties we cannot measure?

**The OLCF Transition to Operations plan is accelerating readiness while emphasizing the science case for Leadership Systems.**

# We Have an Aggressive HPC Roadmap

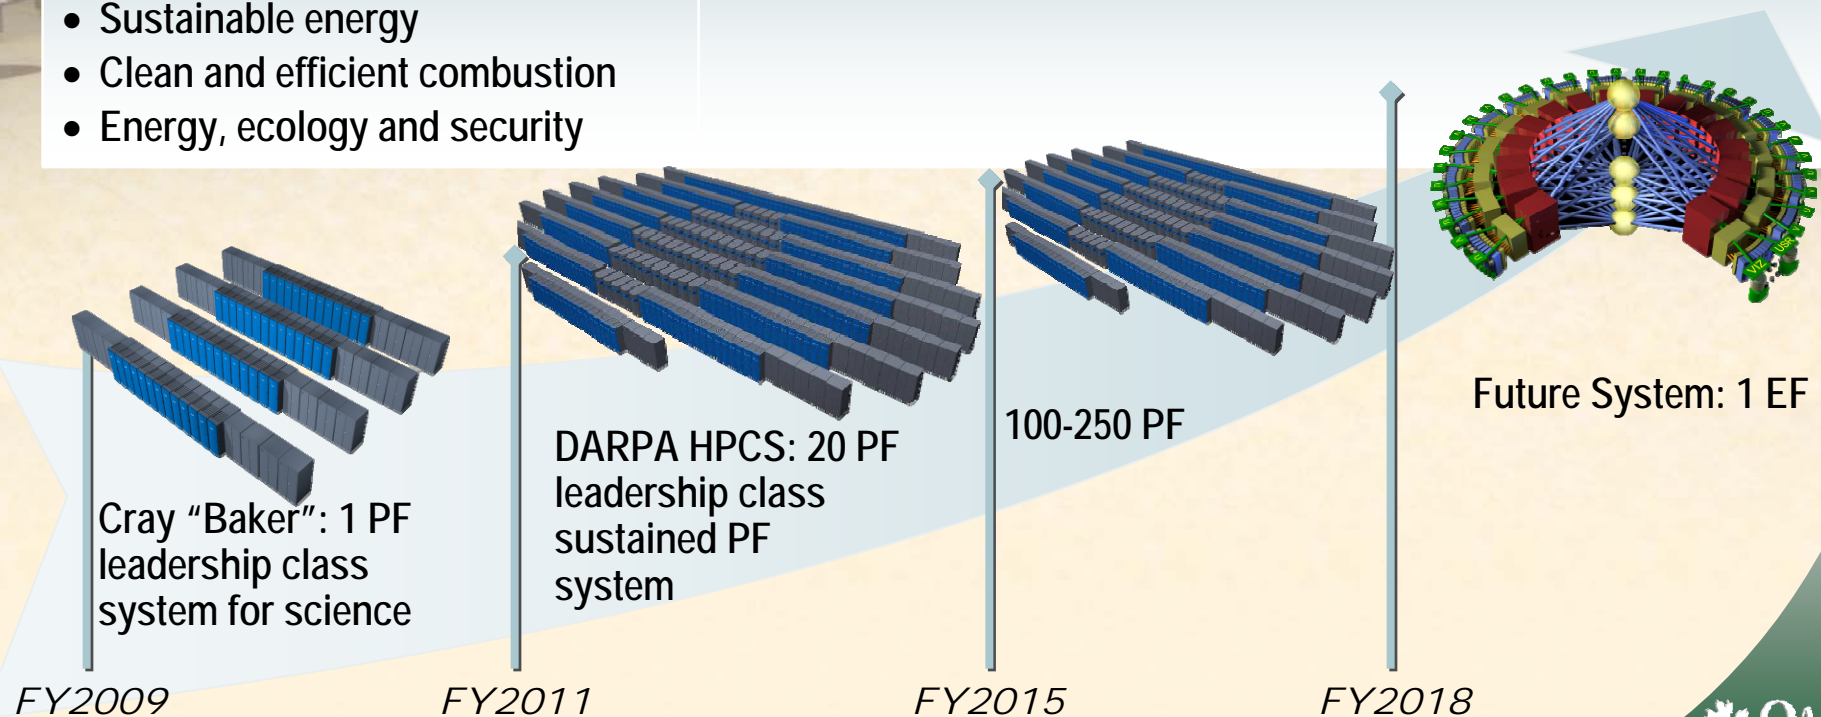


**Mission:** Deploy and operate the computational resources needed to tackle global challenges

- Understanding the universe
- Materials and nanoscience
- Climate change and terrestrial sequestration of carbon
- Sustainable energy
- Clean and efficient combustion
- Energy, ecology and security

**Vision:** Maximize scientific productivity and progress on the largest scale computational problems

- Providing world class computational resources and specialized services for the most computationally intensive problems
- Providing a stable hardware/software path of increasing scale to maximize productive applications development




We have articulated a 10-year plan with required H/W, S/W, and apps "disruptive technologies"



# *Peta/Exa-Scale HPC Technology Collaborations with the DoD*

- **System design, performance and benchmark studies**
- **Wide-area network investigations**
- **HPCS prototypes and full systems**
- **Extreme scale software center**
  - Focused on widening usage and improving productivity of the next generation of “extreme scale” supercomputers
    - Application scaling
    - HPCS languages (X10, Chapel, Fortress)
  - Systems software, tools, environments, and applications development
  - Large scale system reliability, availability, and serviceability (RAS)



## **Multiprogram Research Facility (MRF)**

- 32,000 ft<sup>2</sup> raised floor on two floors
- 20 MW technical power
- 6,000 ton chiller plant

# *Univ. of Tennessee & ORNL Partnership*

## *National Institute for Computational Sciences*

- **UT is building a new NSF supercomputer center from the ground up**
  - Building on strengths of UT and ORNL
  - This center (NICS) goes operational by May 2008
- **Series of computers culminating in a 1 PF system in 2009**
  - Initial delivery (April 2008)
    - 4512 quad-core Opteron processors (170 TF)
  - Cray “Baker” (2009)
    - Multi-core Opteron processors; 100 TB; 2.3 PB of disk space





# Current and Planned Data Centers

## Open Science Center (40K ft<sup>2</sup>)

- ▶ Upgrading building power to 14 MW
- ▶ Deploying a 6,600 ton chiller plant
- ▶ Tripling UPS and generator capability



## Multiprogram Research Facility (30K ft<sup>2</sup>)

- ▶ Capability computing for national defense
- ▶ Expanding to 20 MW of power and 6,000 ton chiller



## Multiprogram Data Center (200K ft<sup>2</sup>)

- ▶ 100K ft<sup>2</sup> classified; 100K ft<sup>2</sup> unclassified
- ▶ Shared mechanical & electrical infrastructure
- ▶ Build out 25K ft<sup>2</sup> on each side as needed
- ▶ Lights out facility



# LCF Computer Systems: Late 2007

*Jaguar: Cray XT4 – 119 TF*

- # 7 on Nov 2007
- 11,706 Dual-core Optrons
- 46 TB memory (2 GB/core)



*Phoenix: Cray X1E – 18 TF*

- Largest Cray vector system in the world
- 1,024 vector processors
- 2 TB shared memory

# The Jaguar Cray XT4 Leadership System



## 2007

- **11,508 compute nodes**
  - Dual-core AMD Opteron processors with 4 GB memory
  - 23,016 compute cores
- **396 service & I/O nodes**
- **~750 TB local storage**
- **3D Torus interconnect**
- **46 TB aggregate memory**
- **119 TF peak performance**


## Today

- **7,832 compute nodes**
  - Quad-core AMD Opteron processors with 8 GB memory
  - 31,328 compute cores
- **240 service & I/O nodes**
- **900 TB local storage**
- **3D Torus interconnect**
- **63 TB aggregate memory**
- **>250 TF peak performance**
- **Currently in software acceptance**
- **General availability to user community expected in May 2008**

# Cray "Baker" – 1 PF System

## FY 2009: Cray "Baker"

- 1 Petaflops system
- 37 Gigaflops processor
- 27,888 quad-core processors  
Barcelona 2.3 GHz
- 2 GB per core; 223 TB total
- 200+ GB/s disk bandwidth
- 13,944 dual-socket 8-core SMP  
"nodes" with 16 GB
- 6.5 MW system power
- 150 Cabinets, 3,500 ft<sup>2</sup>
- Liquid cooled
- Compute node linux  
operating system
- Torus interconnect



# Storage Infrastructure

## We are Deploying Center-Wide Shared File Systems

- “Spider” will be available later this year to provide a shared, parallel file system for all LCF systems
  - Based on Lustre file system
- Planned bandwidth of over 200 GB/s with multi-petabytes (8-12 PB) of capacity
- HPSS provides archival storage for all system
- HPSS has been upgraded with two additional tape libraries to add additional capacity and bandwidth



# Systems Infrastructure

## Current and Projected



Networking	2007	2008	2009
LAN B/W (GB/s)	40	200	200



Archival storage	2007	2008	2009
Capacity (PB)	4	10	20
B/W (GB/s)	4	10	19



Central storage	2007	2008	2009
Capacity (PB)	0.22	1.0	10.0
B/W (GB/s)	10	40	200

# Resources for Viz & Data Analytics

## Hardware

- **Everest Powerwall**
  - 30'x8' 35 Mpixel display wall
- **Hawk Cluster**
  - 58 Opteron nodes with NVIDIA GPUs
- **Lens Cluster**
  - 32 nodes each with 64 GB and 4 quad-core Opterons w/ GPUs
- **Everest Cluster**
  - 18 nodes to drive display wall
- **Ewok Cluster**
  - 81 node dual-core Xeons

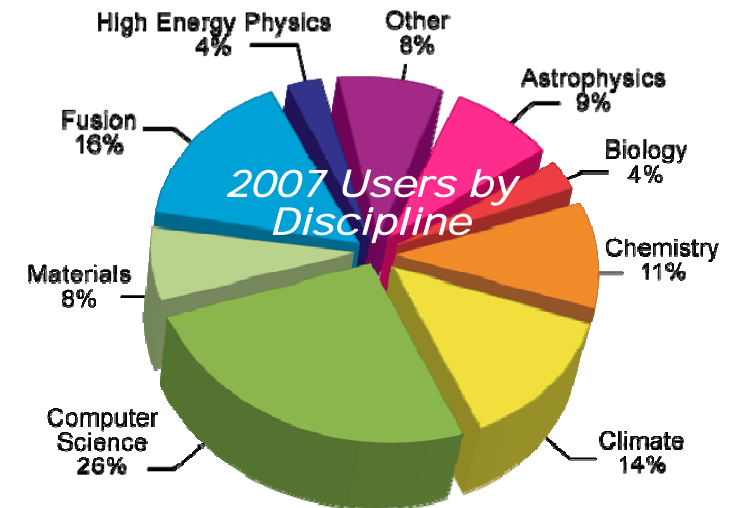
## Software

- **VisIT**
- **EnSight Gold and DR**
- **ParaView**
- **AVS/Express**
- **R MPI**
- **IDL**
- **SCIRun**
- **Xmgrace, Gnuplot, Kepler**



# Our Leadership Computing Facility Has Experienced Lots of Growing "Pains"

- **We have delivered on an aggressive path of support and collaboration with science teams in 2005-2008**
  - INCITE Projects Committed to: 5, 22, 28, 30\*
  - Users Supported: 266, 373, 685, >700\*
  - Compute Hours Used (M): 13, 38, 98, 180\*
  - Aggregate Compute Power (TF): 38, 79, 171, 327\*\*
- **We have established and are adhering to tough availability, usage, and support metrics**
  - Availability: >95% for systems in place >1 year
  - Usage: >30% of usage requires >30% of resource
  - User satisfaction: >75% (on user survey)
- **We have a prudent security posture in place**
  - Moderate Security Certification & Accreditation Level granted in 2007 allows us to fully support ITAR & proprietary work vital to our industry and government users
- **We are better enabling and driving science output**
  - Require quarterly INCITE project reports
  - Track output: publications, awards, presentations
  - Publications: 85 (2005); 129 (2006); 175 (2007 est)



"Impressive maturing of the organization over the last year."

"Doing an outstanding job at problem resolution in a rapidly evolving environment, and in user support in a uniquely challenging area of scientific endeavor."

"Impressed by scientific output – the program of consultants for the scientific projects appears to be effective."

2007 DOE Operational Assessment Review

\*2008 projections


\*\*Includes current Jaguar upgrade



# *We're Organized into Four Groups*

- **HPC Operations**
  - Around-the-clock operations coverage of HPC and storage systems; systems administration; HPSS management and monitoring; configuration management; cyber security
- **Technology Integration**
  - Provides, identifies, develops and integrates new technologies to improve efficiency and support user requirements: InfiniBand, MPI, MPI-IO, HPSS, system Programming, Lustre, Linux Kernel
- **User Assistance & Outreach**
  - Delivering seamless access; accounts management; swift and effective front-line support; on-line docs; 24x7 phone response; Level I/II assistance; showcasing research; workshops and lectures; publications
- **Scientific Computing**
  - facilitates the delivery of science by partnering with users to effectively utilize computational science, visualization and workflow technologies; deeper Level II application support

# Cray System Usage at the ORNL Leadership Computing Facility in CY07



# ORNL Provides Leadership Computing to 2008 INCITE Program <sup>19</sup>

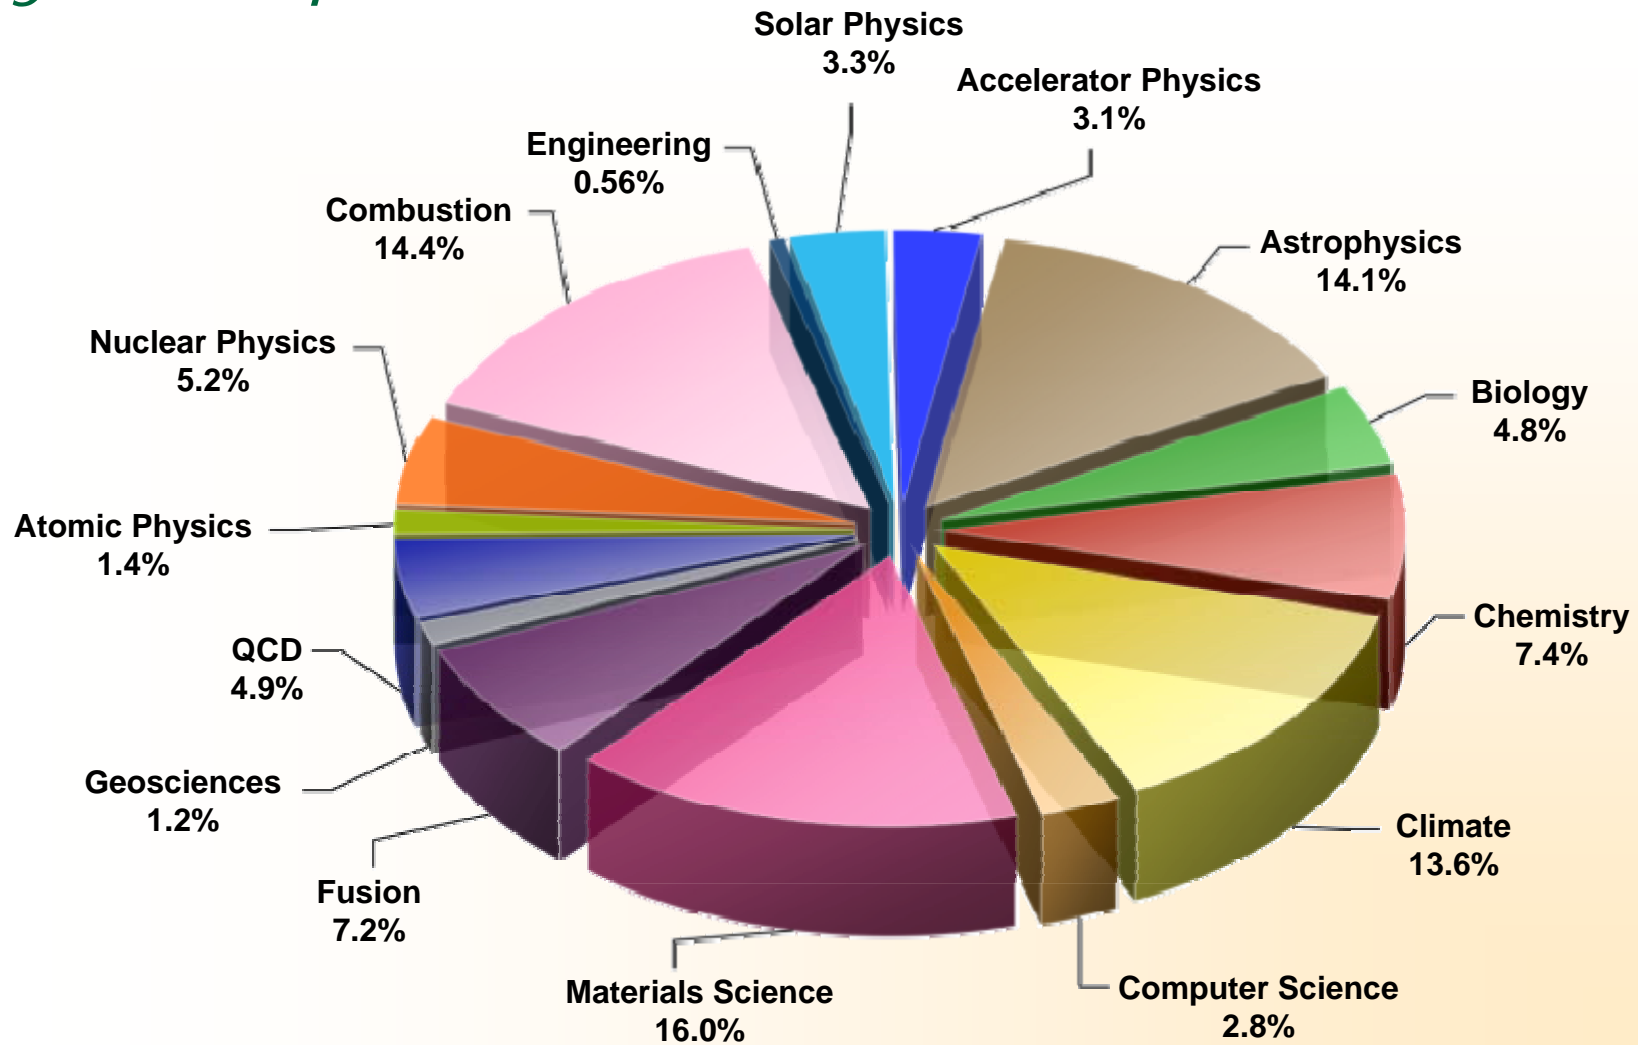
- The NCCS is providing leadership computing to 30 programs in 2008 under the DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program.
- Leading researchers from government, industry, and the academic world will explore challenges including climate change, energy and alternative fuels on the center's leadership computers.
- This year's total allotment of processing hours nearly doubles that which ORNL provided in 2007.



*Project Allocations: 145.3 million hrs*

*Industrial Allocations: 11.9 million hrs*

# ORNL INCITE 2008 Allocations by Discipline



The LCF is able to work with a wide range of science areas

# *Application Codes in 2008*


## *An Incomplete List*

- **Astrophysics**
  - CHIMERA, GenASiS, 3DHFEOS, Hahndol, SNe, MPA-FT, SEDONA, MAESTRO, AstroGK
- **Biology**
  - NAMD, LAMMPS
- **Chemistry**
  - CPMD, CP2K, MADNESS, NWChem, Parsec, Quantum Espresso, RMG, GAMESS
- **Nuclear Physics**
  - ANGFMC, MFDn, NUCCOR, HFODD
- **Engineering**
  - Fasel, S3D, Raptor, MFIX, Truchas, BCFD, CFL3D, OVERFLOW, MDOPT
- **High Energy Physics**
  - CPS, Chroma, MILC
- **Fusion**
  - AORSA, GYRO, GTC, XGC
- **Materials Science**
  - VASP, LS3DF, DCA++, QMCPACK, RMG, WL-LSMS, WL-AMBER, QMC
- **Accelerator Physics**
  - Omega3P, T3P
- **Atomic Physics**
  - TDCC, RMPS, TDL
- **Space Physics**
  - Pogorelov
- **Climate & Geosciences**
  - MITgcm, PFLOTRAN, POP, CCSM (CAM, CICE, CLM, POP)
- **Computer Science (Tools)**
  - Active Harmony, IPM, KOJAK, mpiP, PAPI, PMaC, Sca/LAPACK, SvPablo, TAU

# Preparing for the Future

## *Application requirements: Process and actionable results*

- **LCF Application Requirements Council (ARC)**
  - Stood up in 2006
  - Established ARC charter & reqms management process
- **LCF elicits requirements in many ways**
  - INCITE proposals
  - Questionnaires devised by LCF staff
  - One-on-one interviews
  - Existing publications/documentation
  - Analyzing source code
- **Application categories analyzed**
  - Science motivation and impact
  - Science quality and productivity
  - Application models, algorithms, software
  - Application footprint on platform
  - Data management and analysis
  - Early access science-at-scale scenarios
- **Results**
  - First annual 100+ page Application Requirements Document to be published 9/07
  - New methods for categorizing platforms and application attributes devised and utilized in analysis: guiding tactical infrastructure purchase and deployment
  - Best practice: Process being embraced and emulated by others



## Planned Pioneering Application Runs in Summer 2008

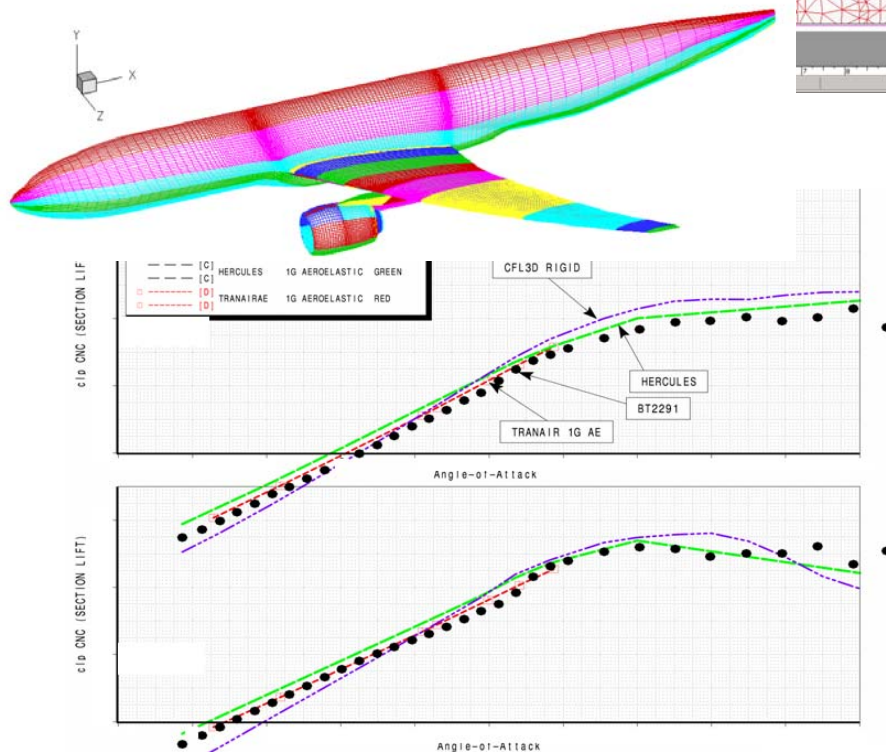


### Cursory Look at the Simulation Specs

Code	Quad-Core Nodes	Global Memory Req'm (TB)	Wall-Clock Time Req'm (hours)	Number of Runs	Local Storage Req'ms (TB)	Archival Storage Req'ms (TB)	Resolution and Fidelity
CHIMERA	7824 4045	16 8	100 100	1 1	13	50	256x128x256 or 256x90x180 20 energy groups, 14 alpha nuclei
GTC-S GTC-C	3900 3900	40 60	36 36	2 2	350	550	600M grid points, 60B particles 400M grid points, 250B particles
S3D	7824	10	140	1	50	100	1B grid points, 15 $\mu$ m grid spacing 4 ns time step, 23 transport vars
POP	2500	1	400	1	1	2	3600x2400x42 tripole grid (0.1° ) 20-yr run; partial bottom cells; first with biogeochemistry at this scale
MADNESS	7824	48	12 2	10 12	5	50	600B coefficients
DCA++	2000 6000	16 48	12 to 24	20	1	1	Lattices of 16 to 32 sites 80 to 120 time slices O(10 <sup>2</sup> -10 <sup>3</sup> ) disorder realizations

**Astrophysics — Fusion — Combustion — Climate — Chemistry — Materials Science**

# Application: The Boeing Company

Development and Correlations of Large Scale Computational Tools for Flight Vehicles



Evaluate the inclusion of unsteady flow physics (Detached Eddy Simulations / DES) to improve analysis:

- Thrust reverser efflux flows
- High Mach number flows
- High angle-of-attacks separated flows

Ref: Moeljo Hong (Boeing)



# Application Example


## Cartesian Methods for FSI

### Problem:

- Dynamic coupling of explicit shock-capturing methods for strongly driven problems requires an efficient and scalable approach for mesh modification

### Concept:

- Use a level-set-based approach to couple Lagrangian solid mechanics solvers to inviscid Eulerian fluid mechanics solvers
- Specialized algorithm to update distance information on the fly
- Use dynamic blockstructured mesh adaptation to mitigate boundary approximation errors
- Eulerian-Lagrangian inter-solver communication library synchronizes the boundary data exchange between distributed solver modules



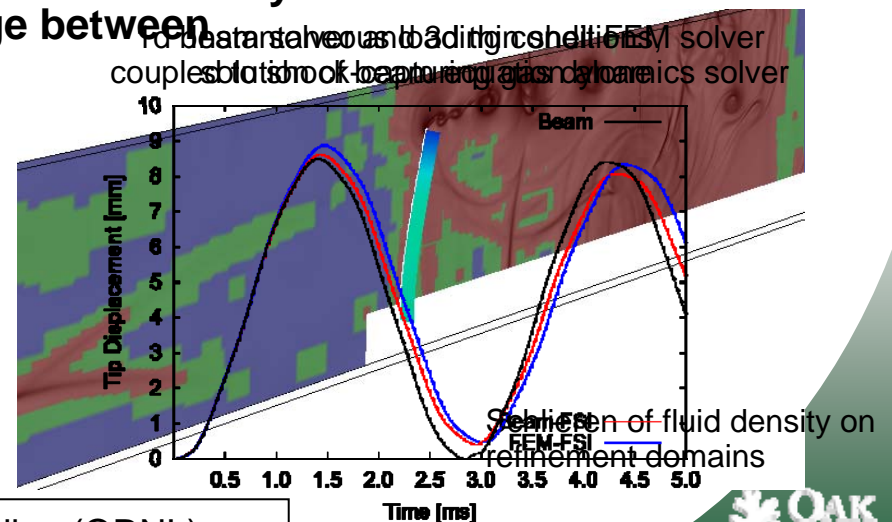
Internal boundary cells (gray) and fluid value interpolation in a level set method

### Verification example:

Elastic motion of a thin steel plate (thickness 1mm, length 50mm) hit by a planar Mach 1.21 shock wave in air

Forward facing step geometry, reflective boundaries everywhere except inflow at left side, panel 1.5cm behind start of step

Fluid base mesh 320x64, 2 additional level with factors 2, 4




Ref: Ralf Deiterding (ORNL)

HPC User Forum, Norfolk, VA, April 16, 2008

# CFD Application Example

## Turbulent Structure & Interface Dynamics in Steep Breaking Waves

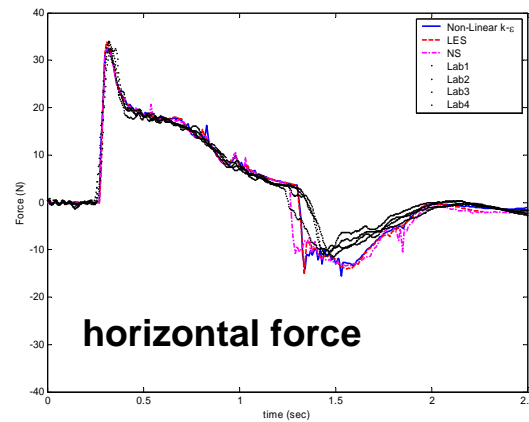
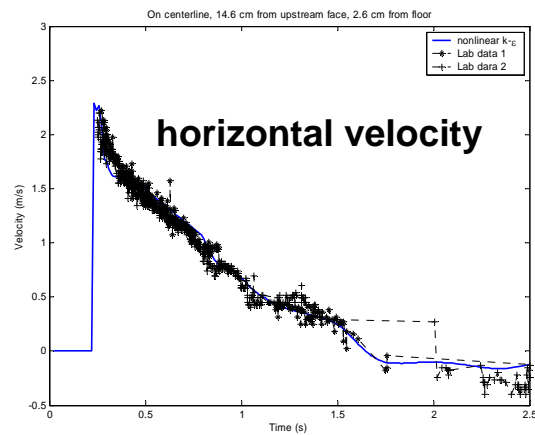
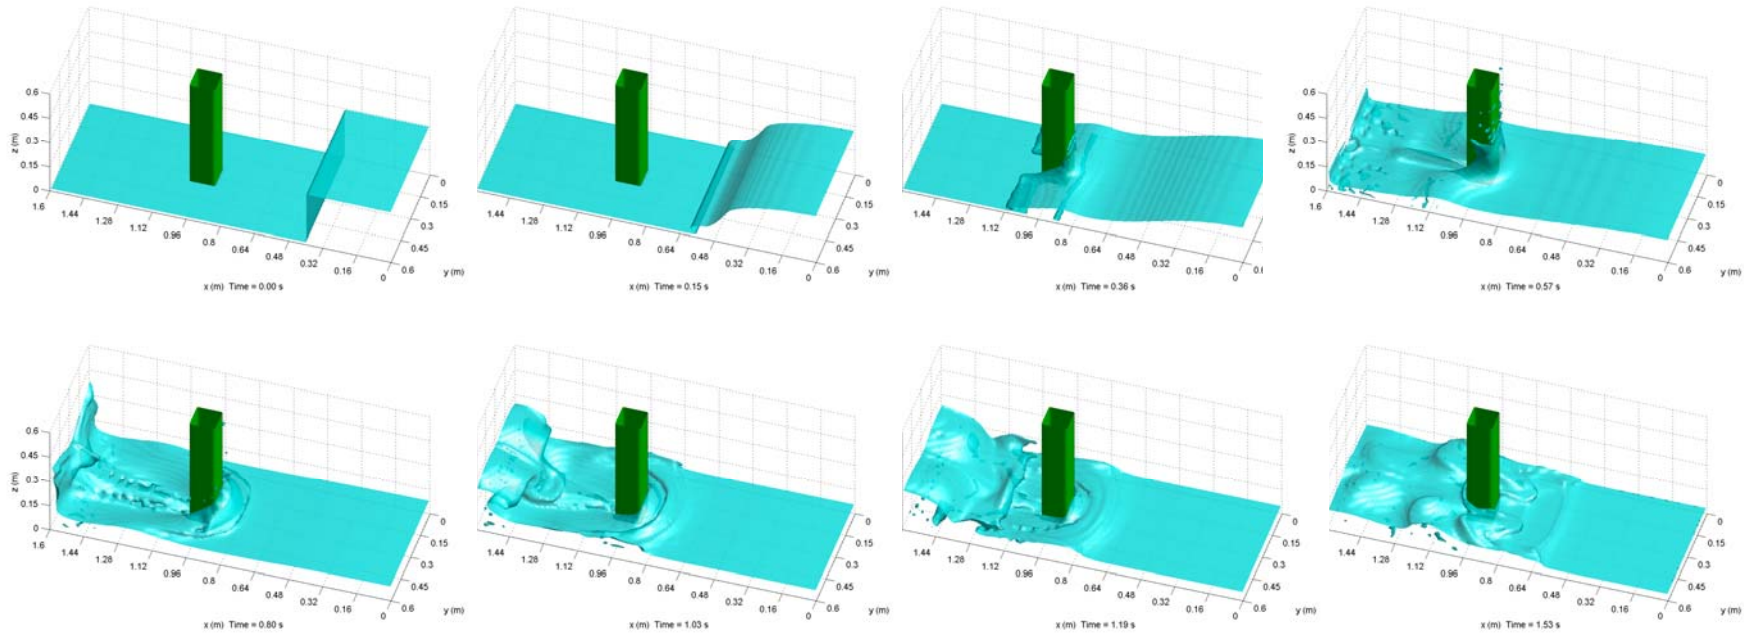


- High-order finite difference formulation
- 2<sup>nd</sup>-order, planarity-preserving volume tracking algorithm
- Interface topology: iterative method for unit normal; height function for curvature
- Advanced LES model with reconstructed distance function near free surface
- Capacity parallel simulation (100s of processors) but need lots more!
- The simulated flow behavior is consistent with theory and experiment
  - Wave steepening is 2D, whereas breaking is inherently 3D
- Application: variety of industrial, natural, and defense-related flows

Ref: Petar Liovic (Univ. of Melbourne), Djamel Lakehal (ETH Zurich)

# CFD Application Example

## Dam-Break Bore Interacting with Square Cylinder



Ref: Philip L.-F. Liu (Cornell), Doug Kothe (ORNL)

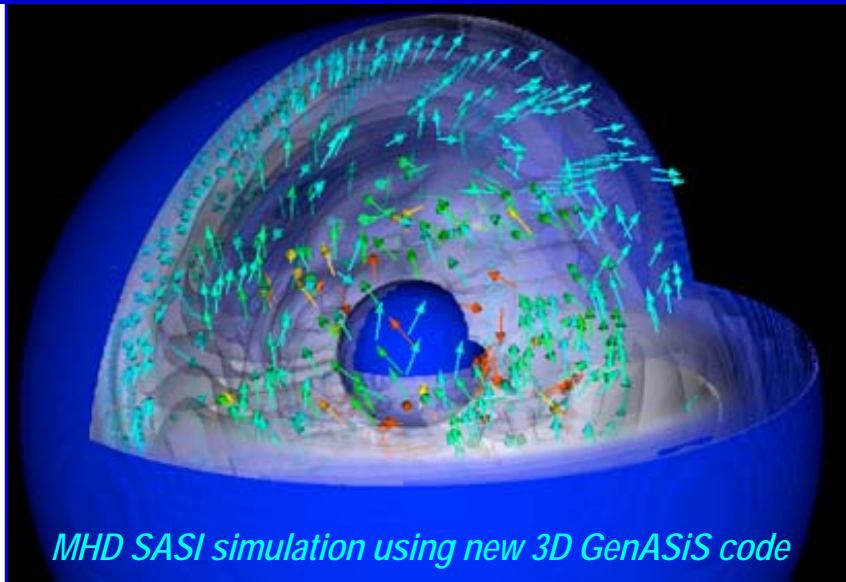
# Barriers in Ultrascale Climate Simulation

## Attacking the Fourth Dimension: Parallel in Time

- **Problem**
  - Climate models use explicit time stepping
  - Time step must go down as resolution goes up
  - Time stepping is serial
  - Single-process performance is stagnating
  - More parallel processes do not help!
- **Possible *complementary* solutions**
  - Implicit time stepping
  - High-order in time
  - “Fast” bases: curvelets and multi-wavelets
  - “Parareal” parallel in time
- **Progress**
  - Implicit version of HOMME for global shallow-water equations: 10x speedup for steady-state test case
  - High-order single-step time integration
  - Single-cycle multi-grid linear solver for 1D
  - Pure advection with curvelets and multi-wavelets
- **Near-term plans**
  - Scale, tune, and precondition implicit HOMME
  - Single-cycle multi-grid linear solver for 2D
  - “Parareal” for Burgers’ (1D nonlinear)

# Discovering the Elusive Core Collapse Supernova Explosion Mechanism

**Researchers glean unprecedented insight into the shock waves that blow apart a 10- to 20-solar mass star**



Researchers can now simulate ~1 second after 'post-bounce'. Petascale systems will allow longer simulations: tens of seconds after the explosion and will allow inclusion of neglected yet important physics such as magnetic fields.

Ref: Tony Mezzacappa (ORNL)


- Achieved longer run simulations and, 0.8 seconds after explosion, saw the initial shock wave revived by turbulence of in-falling material
- CHIMERA used to investigate multiple stellar models, effect of both Newtonian and Einsteinian gravity, and impact of recently discovered subatomic physics
  - >12K cores used in current 3D simulations
- Current 3D spatial resolution
  - 78x156x312 (Chimera)
  - 256x256x256 (Genasis)

## LCF liaison contributions

- Implementing efficient, collective I/O
- Pencil decomposition of 3D flow algorithm
- Preconditioning of the neutrino transport equation

# New Results in Flame Stabilization in an Auto-Ignitive Jet

- First fully-resolved simulation of a 3D lifted flame in heated co-flow with detailed chemistry
- Lifted flames occur in diesel engines and gas turbine combustors
  - Flame stabilized against fuel jet and recirculating hot gases
- Direct numerical simulation of a lifted flame in heated co-flow
  - ~1 billion grid points and 14 degrees of freedom per grid point
  - H<sub>2</sub>/Air detailed chemistry
  - Jet Reynolds number = 11,000
  - Largest DNS at the highest Reynolds number
  - 2.5M hours on Jaguar at the LCF
- Simulation reveals source of stabilization
  - Upstream auto-ignition
  - Vorticity generation at flame base due to baroclinic torque



**Instantaneous OH radical concentration on a stoichiometric mixture fraction iso-surface shows flame lift-off**

Ref: Jackie Chen (ORNL)

## LCF liaison contributions

- Cray X1E loop vectorization of S3D
- Identified and fixed X1E MPI bottleneck
- Lagrangian tracers; I/O rework with NW University
- Jaguar scaling studies helped to identify processors burdened by memory corrections

# LCF 2008 Director Discretion Program

## LCF Director Discretion Program

Jaguar – 17.8M hours  
Phoenix – 0.5M hours

Application  
Performance and  
Data Analytics

Strategic  
Partnerships

Leadership  
Computing  
Preparation

We welcome more collaboration and presence on our systems

# *Supplementary Material*



# INCITE 2008 Projects



## Astrophysics

Multi-Dimensional Simulations of Core-Collapse Supernovae  
**Anthony Mezzacappa (Oak Ridge National Laboratory)**

First Principles Models of Type Ia Supernovae  
**Stan Woosley (University of California, Santa Cruz)**

Numerical Relativity Simulations of Binary Black Holes and Gravitational Radiation  
**Joan Centrella (National Aeronautics and Space Administration)**

## Biology

Cellulosic Ethanol: Physical Basis of Recalcitrance to Hydrolysis of Lignocellulosic Biomass  
**Jeremy Smith (Oak Ridge National Laboratory)**

Gating Mechanism of Membrane Proteins  
**Benoit Roux (Argonne National Laboratory & Univ. of Chicago)**

## Chemistry

Molecular Simulation of Complex Chemical Systems  
**Christopher Mundy (Pacific Northwest National Laboratory)**

An Integrated Approach to the Rational Design of Chemical Catalysts  
**Robert Harrison (Oak Ridge National Laboratory)**

## Computer Science

Performance Evaluation & Analysis Consortium End Station  
**Patrick Worley (Oak Ridge National Laboratory)**

## Climate

The Role of Eddies in the Meridional Overturning Circulation  
**Paola Cessi (University of California, San Diego)**

Assessing Global Climate Response of the NCAR-CCSM3: CO<sub>2</sub> Sensitivity and Abrupt Climate Change  
**Zhengyu Liu (University of Wisconsin, Madison)**

Eulerian and Lagrangian Studies of Turbulent Transport in the Global Ocean  
**Synte Peacock (ASC/Alliance Flash Center, University of Chicago)**

Climate-Science Computational End Station Development and Grand Challenge Team  
**Warren Washington (National Center for Atmospheric Research)**

Modeling Reactive Flows in Porous Media  
**Peter Lichtner (Los Alamos National Laboratory)**

## Engineering

High-Fidelity Simulations for Clean and Efficient Combustion of Alternative Fuels  
**Jacqueline Chen (Sandia National Laboratories)**

Clean and Efficient Coal Gasifier Designs Using Large-Scale Simulations  
**Madhava Syamlal (National Energy Technology Laboratory)**

Landmark Direct Numerical Simulations of Separation and Transition for Aerospace-Relevant Wall-Bounded Shear Flows  
**Hermann Fasel (University of Arizona)**

# INCITE 2008 Projects (continued)



## Fusion

Verification and Validation of Petascale Simulation of Turbulent Transport in Fusion Plasmas  
**Patrick Diamond (University of California, San Diego)**

Fluctuation Spectra and Anomalous Heating in Magnetized Plasma Turbulence  
**William Dorland (University of Maryland)**

Gyrokinetic Steady State Transport Simulations  
**Jeff Candy (General Atomics)**

High Power Electromagnetic Wave Heating in the ITER Burning Plasma  
**Fred Jaeger (Oak Ridge National Laboratory)**

## Materials (continued)

Bose-Einstein Condensation vs. Quantum Localization in Quantum Magnets  
**Tommaso Roscilde (Max-Planck Gesellschaft)**

Linear Scale Electronic Structure Calculations for Nanostructures  
**Lin-Wang Wang (Lawrence Berkeley National Laboratory)**

## Materials

Predictive and Accurate Monte Carlo Based Simulations for Mott Insulators, Cuprate Superconductors, and Nanoscale Systems  
**Thomas Schulthess (Oak Ridge National Laboratory)**

Electronic, Lattice, and Mechanical Properties of Novel Nanostructured Bulk Materials  
**Jihui Yang (GM R&D Center)**

Development and Correlations of Large-Scale Computational Tools for Flight Vehicles  
**Moeljo Hong (The Boeing Company)**

## Physics

Computational Nuclear Structure  
**David J. Dean (Oak Ridge National Laboratory)**

Petascale Computing for Terascale Particle Accelerator: International Linear Collider Design and Modeling  
**Lie-Quan Lee (Stanford Linear Accelerator Center)**

Lattice QCD  
**Robert Sugar (University of California, Santa Barbara)**


Computational Atomic and Molecular Physics for Advances in Astrophysics, Chemical Sciences, and Fusion Energy Sciences  
**Michael Pindzola (Auburn University)**

Modeling Heliospheric Phenomena with an Adaptive, MHD-Boltzmann Code  
**Nikolai Pogorelov (University of California, Riverside)**

# Producing New Insights for RF Heating of ITER Plasmas

**Fully 3-dimensional simulations of plasma shed new light on the behavior of superheated ionic gas in the multibillion-dollar ITER fusion reactor**

**“Until recently, we were limited to two-dimensional simulations. The larger computer [Jaguar] has allowed us to achieve three-dimensional images and validate the code with observations.” – Fred Jaeger, ORNL**



- 3D simulations reveal new insights
  - “Hot spots” near antenna surface
  - “Parasitic” draining of heat to the plasma surface in smaller reactors
- Work pushing the boundaries of the system (22,500 processor cores, 87.5 TF) and demonstrating
  - Radial wave propagation and rapid absorption
  - Efficient plasma heating
- AORSA’s predictive capability can be coupled with Jaguar power to enhance fusion reactor design and operation for an unlimited clean energy source

## LCF liaison contributions


- Converted HPL from double real to double complex and replaced Scalapack
  - 45 TF to 75 TF!
- Acquired new version of BLAS from TACC
  - 75 TF to 87 TF!
- Net performance gain of almost a factor of 2

# Accelerating Climate Science

- **First-ever control runs of CCSM 3.5 at groundbreaking speed**

*“[On Jaguar,] we got 100-year runs in three days. This was a significant upgrade of how we do science with this model. 40 years per day was out of our dreams.”*

Peter Gent of NCAR, Chairman of CCSM Scientific Steering Committee, during keynote at CCSM Workshop, June 19, 2007



**Instantaneous net ecosystem exchange (NEE):** eastern half is in sunlight and the terrestrial ecosystems are taking up carbon (negative NEE, shown in green to bright white). Meanwhile, the sun has not yet risen in the western half of the image where the ecosystems are only respiring (positive NEE, shown in red)


- **Major improvements in CCSM 3.5**
  - Arctic and Antarctic sea ice: Will the Arctic be ice free in summer of 2050?
  - Surface hydrology of land, critical for predictions of drought
- **Positioned to test full carbon-nitrogen cycle**

## LCF liaison contributions

- New preconditioner for barotropic solver
- Visualization of CO<sub>2</sub> transport
- Contributed bug fixes to POP 2.0
- Represented needs at OBER/ESNET meeting

# Nanostructural Features in High Performance Thermoelectric Materials

**Researchers simulate materials that can transform automobile waste heat directly into electricity**



- Only 25% of automobile fuel energy is utilized for vehicle mobility and accessories
- Team led by GM is working to develop waste heat recovery technology (a “thermoelectric alternator”) for fuel economy improvement
  - Goal: 3-5% fuel economy increase with no emission increase
- Largest-ever simulation – 1700 atom supercell – made possible at LCF
- Exploring low thermal conductivity properties of silver and antimony-doped lead-telluride-based material

**“Only at a place like the LCF can such an expensive calculation be done. We’re very lucky that LCF has been very supportive.”**


– Jihui Yang, General Motors

## LCF liaison contributions

- Worked with GM researchers to evolve a more scalable and capable version of VASP ab-initio code

# Gaining Understanding of Cause and Effect of Core Plasma Turbulence

- Simulation of experimental discharges (NSTX and DIII-D) has shown the behavior of microturbulence to be intimately related to geometry and shaping
- Recent improvements to GTC-S allowed more realistic simulations of electron temperature gradient (ETG) drift instabilities, ion temperature gradient (ITG) drift instabilities with non-adiabatic electrons, and trapped electron modes (TEM)
  - The number of particles included in recent simulations allows this project to reduce the amount of statistical noise and explore core turbulence at levels of fidelity never seen before





## LCF liaison contributions

- Asynchronous I/O
- Automated end-to-end workflow
- Porting/scaling new shaped plasma version
- Helped with new toroidal plane decomposition

# Creation of efficient enzymes for cellulose degradation through protein engineering

- *Renewable energy:*  
**Ethanol production from cellulose**
- **Developing a detailed understanding of cellulase enzyme mechanisms from multiscale modeling**
  - 1- to 100-ns trajectories for systems with more than 800,000 atoms
- **Carrying out simulations with different substrates and mutant enzymes**



*C. thermocellum*  
growing on cellulose

Metabolism of cellulose  
by exocellulase enzyme

# Simulation Aids Development of First Coal Plants with Near-Zero Emissions

## Plants could soon produce electricity with minimal CO<sub>2</sub>

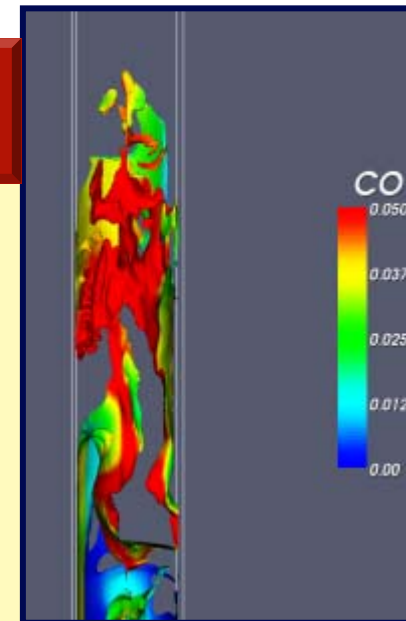
- Different sizes and densities for the fresh coal and the recycle material were considered
- 2048-core simulation was a 50x increase in size relative to standard gasifier simulations

*“An important part of NETL’s mission is to supply clean coal technology, and our research group at NETL develops computational tools and applications in support of that mission.”*

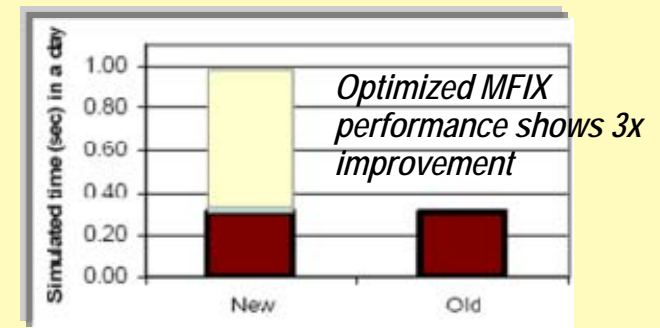
**Madhava Syamlal , NETL**

### LCF liaison contributions

- Optimized MPI settings
- Improved MFIX multi-phase code performance 3x
- Used more highly optimizing compiler (Pathscale)
- Reduced global (collective) communication without changing base algorithm



*Efficient gasifiers maximize carbon monoxide (CO) production. Particle clusters colored by CO concentration are shown. Image courtesy NETL.*





# *The Science Case*

- **Energy**
  - Climate, materials, chemistry, combustion, biology, nuclear fusion/fission
- **Fundamental**
  - Materials, chemistry, astrophysics
- **There are many others**
  - QCD, accelerator physics, wind, solar, engineering design (aircraft, ships, cars, buildings)
- **What are key system attribute issues?**
  - Peak speed
  - Memory (capacity, B/W, latency)
  - Interconnect (B/W, latency)

# Science Case: Climate

**Mitigation: Evaluate strategies and inform policy decisions for climate stabilization; 100-1000 year simulations**  
**Adaptation: Decadal forecasts & region impacts; prepare for committed climate change; 10-100 year simulations**

- **250 TF**
  - **Mitigation: Initial simulations with dynamic carbon cycle and limited chemistry**
  - **Adaptation: Decadal simulations with high-resolution ocean (1/10° )**
- **1 PF**
  - **Mitigation: Full chemistry, carbon/nitrogen/sulfur cycles, ice-sheet model, multiple ensembles**
  - **Adaptation: High-resolution atmosphere (1/4° ), land, and sea ice, as well as ocean**
- **Sustained PF**
  - **Mitigation: Increased resolution, longer simulations, more ensembles for reliable projections; coupling with socio-economic and biodiversity models**
  - **Adaptation: Limited cloud-resolving simulations, large-scale data assimilation**
- **1 EF**
  - **Mitigation: Multi-century ensemble projections for detailed comparisons of mitigation strategies**
  - **Adaptation: Full cloud-resolving simulations, decadal forecasts of regional impacts and extreme-event statistics**

**Resolve clouds, forecast weather & extreme events, provide quantitative mitigation strategies**

# Science Case: Materials

- **250 TF**
  - Multi band Hubbard model calculations that will enable the investigation of realistic models of high Tc superconductors
  - Will contribute significantly to the understanding of the mechanism that leads to highTc superconductivity.
- **1 PF**
  - Thermodynamics of magnetism in realistic sized nanoparticles using a combination of first principles for the magnetic state and classical Monte Carlo for thermodynamics
- **Sustained PF**
  - Thermodynamics of composition of alloy nanoparticles
  - Understand the synthesis of alloy nanoparticles with potential impact for the design of new catalysts
- **1 EF**
  - Perform Molecular Dynamics with forces found with Quantum Monte Carlo computations
  - This will provide a superior description of the microscopic behavior of water and improve the understanding of systems in aqueous environments (e.g. biological molecules etc.)

**Understand high Tc superconductivity, design & synthesize nanoparticles, predict biological system behavior**

# Science Case: Chemistry



- **250 TF**
  - Accurate large-scale, all-electron, density functional simulations
  - Absorption on transition metal oxide surface (important catalytic phenomenon)
  - Benchmarking of Density-Functionals (importance of Hartree-Fock exchange)
- **1 PF**
  - dynamics of few-electron systems
  - Model of few-electron systems interaction with intense radiation to a guaranteed finite precision
- **Sustained PF**
  - Treatment of absorption problem with larger unit cells to avoid any source of error
- **1 EF**
  - Extension of the interaction with intense radiation to more realistic systems containing more electrons

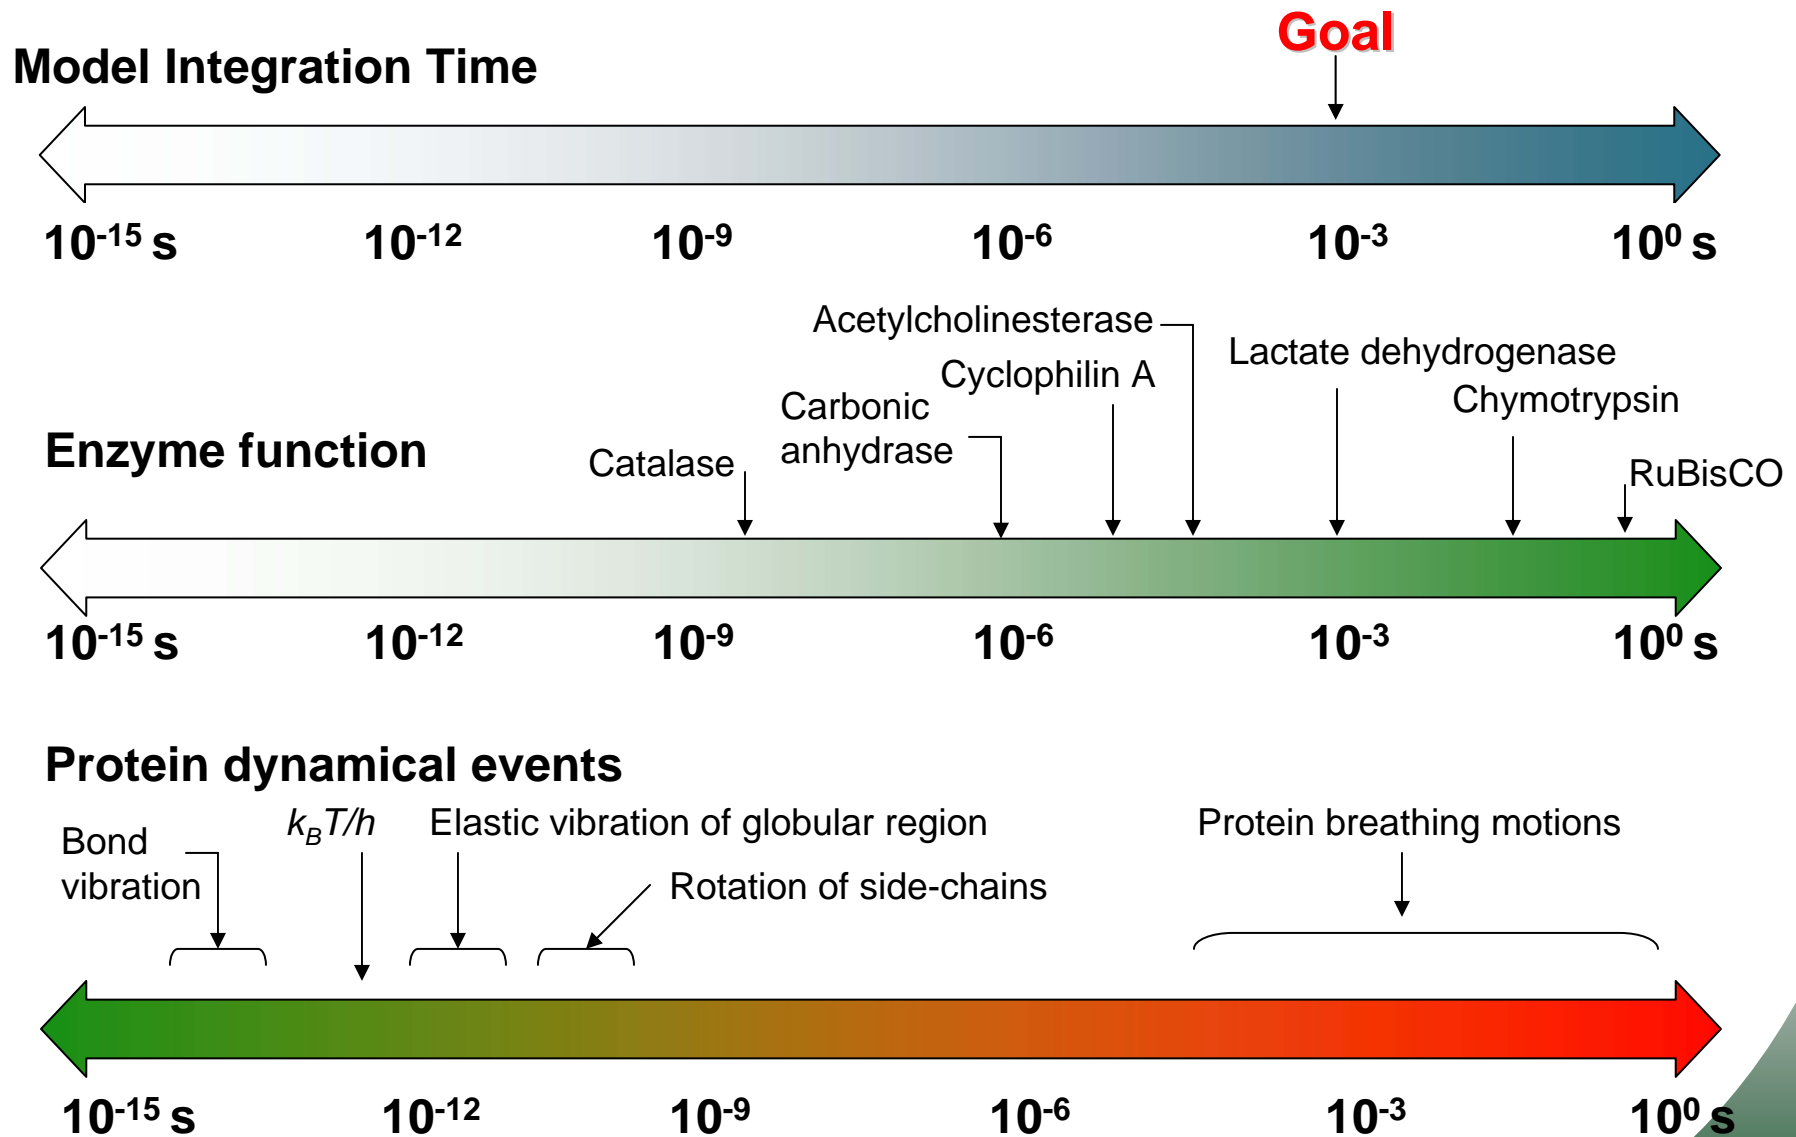
**Design catalysts, understand radiation interaction with materials, quantitative prediction of absorption processes**

# Science Case: Combustion

- **250 TF**
  - Lifted hydrocarbon flame
  - Determine the mechanism of flame stabilization and the role of auto-ignition vs. edge flame propagation
- **1 PF**
  - Stabilization of auto-igniting diesel jets
  - Predictive models for the effect of multi-stage ignition and cool chemistry in diesel combustion
- **Sustained PF**
  - Liquid jet combustion with spray dynamics, evaporation, & chemical reaction at phase boundaries
  - Active control of spray shape and injection parameters for clean combustion in flexible fuel vehicles
- **1 EF**
  - Emission in direct injection combustion using multiple statistical moments and stochastic models for particulate matter
  - Reduce soot and emissions from combustion of diesel and alternative fuels such as bio-butanol

**Clean and efficient combustion design for diesel and alternative (bio) fuels**


# Science Case: Biology



# *Science Case: Biology*

- **Ethanol production from cellulose**
  - **Creation of more efficient enzymes for cellulose degradation through protein engineering**
  - **Detailed understanding of cellulase enzyme mechanisms**
- **Carbon sequestration**
  - **Understanding the activity of the large enzyme RuBisCO, occurring on the order of seconds**

# Science Case: Astrophysics



- **250 TF**
  - The interplay of several important phenomena: hydrodynamic instabilities, role of nuclear burning, neutrino transport
- **1 PF**
  - Determine the nature of the core-collapse supernova explosion mechanism
  - Fully integrated, 3D neutrino radiation hydrodynamics simulations with nuclear burning
- **Sustained PF**
  - Detailed nucleosynthesis (element production) from core-collapse SNe
  - Large nuclear network capable of isotopic prediction (along with energy production)
- **1 EF**
  - Precision prediction of complete observable set from core-collapse SNe: nucleosynthesis, gravitational waves, neutrino signatures, light output
  - Tests general relativity and information about the dense matter equation state, along with detailed knowledge of stellar evolution
  - Full 3D Boltzmann neutrino transport, 3D MHD/RHD, nuclear burning



**Explanation and prediction of core-collapse SNe; put general relativity, dense EOS, stellar evolution theories to the test**



# Science Case: Fusion Micro-Turbulence

## GTC-C

- **250 TF**
  - Address critical physics issues in drift wave turbulence
- **1 PF**
  - Study mesoscale dynamics; structure formation and heat and particle transport in CTEM turbulence
- **1 EF**
  - Core-edge coupling for a predictive ITER capability

## GTC-S

- **250 TF**
  - Extensive code validation of NSTX plasmas with ETG data
- **1 PF**
  - Fluctuation spectra and zonal flow patterns compared with BES Data
- **1 EF**
  - ITER predictive capability for ITG ETG coupling with realistic mass ratios
  - Couple with edge code and look into late time simulations.

**Predictive ITER capability with core-edge coupling, realistic mass ratios, and validated turbulence models**

# System Attribute: Memory B/W

- **Application drivers**
  - Multi-physics, multi-scale applications stress memory bandwidth as they do node memory capacity
    - “Intuitive” coding styles often produce poor memory access patterns
  - Poor data structures, excessive data copying, indirect addressing
- **Algorithm drivers**
  - Unstructured grids, linear algebra
  - Ex: AMR codes work on blocks or patches, encapsulating small amounts of work per memory access to make the codes readable and maintainable
    - This sort of architecture requires very good memory BW to achieve good performance
- **Memory B/W suffers in the multi-core future**
  - Must apps just have to “get used to not having any”?
- **Methods to (easily) exploit multiple levels of hierarchical memory are needed**
  - Cache blocking, cache blocking, cache blocking
  - Gather-scatter

# System Attribute: Memory Capacity

- **Application drivers: strongly coupled multi-physics**
  - Multiple degrees of freedom at each spatial grid point lead to large, multidimensional data structures
  - Examples: chemical networks for combustion, neutrino transport with large phase space at each spatial point, neutron transport in nuclear reactors
- **The current ratio of flops to bytes**
  - *Must* be maintained or increased
- **A minimum per-process amount is needed**
  - There is a threshold below which science applications cannot perform
  - Correlated to code complexity and slowly growing
- **Can be used to reduce peak speed and memory B/W requirements**
  - Ex: redundantly store rather than re-compute

# System Attribute: Peak Speed

- **Application drivers**
  - All applications responded with high priority for FLOPs in the application requirement survey
  - Strong scaling problem domains (e.g., engr apps), unscalable algorithms
- **Algorithm drivers**
  - Linear algebra, FFT, Monte Carlo, high on-PE data volume:surface ratio
- **Path forward**
  - Can often help to compensate for other system attributes falling below optimal performance (e.g., bandwidth)
  - Computational performance can offset other deficits
    - Overprovision FLOP to minimize communication
    - FLOP-intensive algorithms will achieve faster time to solution than FLOP-conserving with memory/network intensive algorithms
  - FLOPs increase will come from deeper pipelines and vector operations
    - Need streamlined and vectorizable algorithms
  - Some applications unable to benefit from increased FLOPs is indicative of other performance problems
    - Ex: indirect addressing invokes lots of integer/logical ops

# *System Attribute: Memory Latency*

- **Application drivers**
  - Random access to data in main memory
  - Retrieval of many small pieces of information (e.g., MD)
- **Algorithm drivers**
  - Stride-one access patterns, cache-aware algorithms
  - Sparse linear algebra
- **Path forward**
  - Recompute instead of store (FLOPs can offset somewhat)
  - Reorder computation
  - Cache-happy algorithms
  - Distributed-shared-memory implementations on a single socket
  - Vectorization
  - Gather-scatter

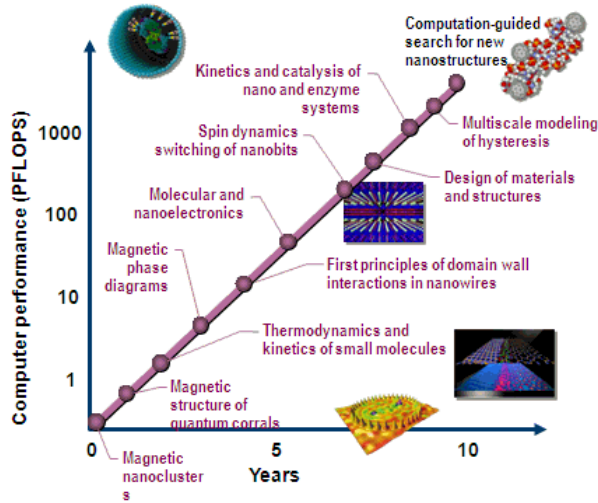
# *System Attribute: Interconnect B/W*

- **Application drivers**
  - Coupled or multi-physics (implicit) applications
  - Communicating few, large messages (e.g., FFT)
  - Most applications see the need for this attribute increasing
  - Astrophysics, nuclear energy, fusion, engineering apps
- **Algorithm drivers**
  - Multi-level methods
  - Load balancing
  - Multi-component coupling
  - Asynchronous I/O
  - Embedded analysis and visualization
- **Path forward**
  - Optical interconnects: exponential growth feasible
  - Trade this off with more-limited resources

# System Attribute: Interconnect Latency

- **Application drivers**
  - Others: chemistry, materials science
  - Biology
    - Stated goal of 1 ms simulation time per wall clock day translates to 0.1  $\mu$ s wall clock time per time step (assuming  $10^7$  integration time steps)!
- **Algorithm drivers**
  - Explicit algorithms using nearest-neighbor or systolic communication
  - Medium- to fine-grain parallelization strategies (e.g. various distributed data approaches in computational chemistry)
- **Not likely to fall below the 1  $\mu$ s barrier**
  - Yet raw compute power keeps getting faster, increasing imbalance
- **Path forward**
  - Need new parallelization languages and environments
  - The combination of SW and HW must allow the ability to fully overlap communication and computation
  - Specialized hardware for common ops?
    - Team synchronization; global and semi-global reductions
  - Vectorization/multithreading of communication?

# DOE-SC Science Drivers

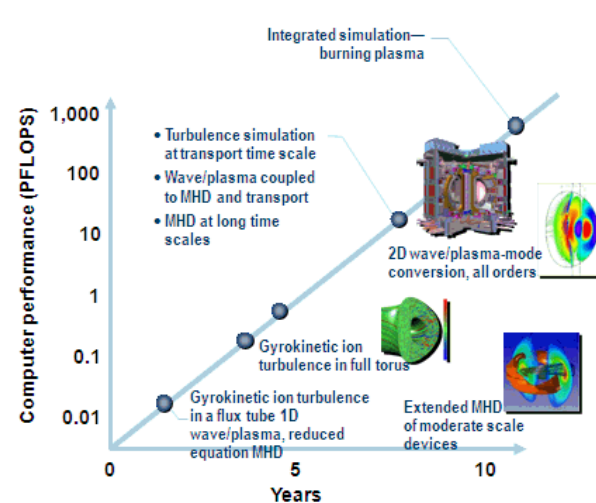


## Expected outcomes 5 years

- Realistic simulation of self-assembly and single-molecule electron transport
- Finite-temperature properties of nanoparticles/quantum corrals

## 10 years

- Multiscale modeling of molecular electronic devices
- Computation-guided search for new materials and nanostructures



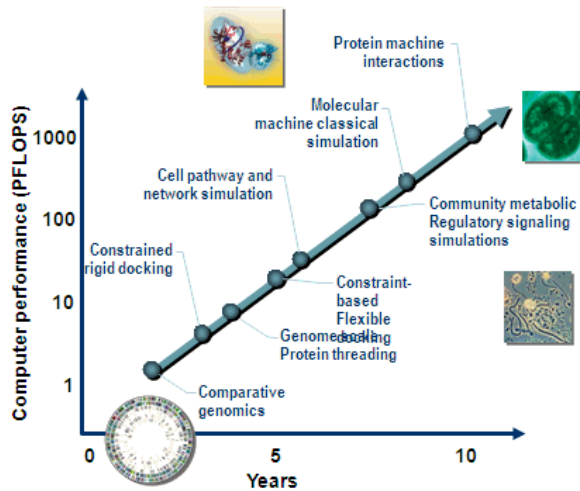
## Expected outcomes 5 Years

- Full-torus, electromagnetic simulation of turbulent transport with kinetic electrons for simulation times approaching transport time-scale
- Develop understanding of internal reconnection events in extended MHD, with assessment of RF heating and current drive techniques for mitigation

## 10 years

- Develop quantitative, predictive understanding of disruption events in large tokamaks
- Begin integrated simulation of burning plasma devices –

## Fusion



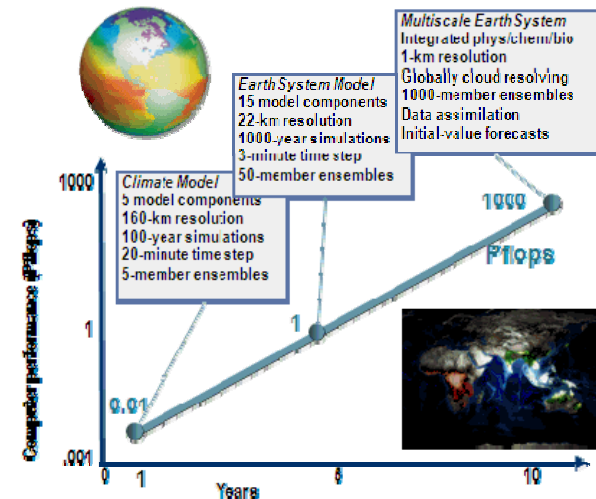
## Expected outcomes 5 years

- Metabolic flux modeling for hydrogen and carbon fixation pathways
- Constrained flexible docking simulations of interacting proteins

## 10 years

- Multiscale stochastic simulations of microbial metabolic, regulatory, and protein interaction networks
- Dynamic simulations of complex molecular machines

## Biology



## Expected outcomes 5 years

- ☐ Fully coupled carbon-climate simulation
- ☐ Fully coupled sulfur-atmospheric chemistry simulation

## 10 years

- ☐ Cloud-resolving 1-km spatial resolution atmosphere
- ☐ Fully coupled, physics, chemistry, biology earth system model

## Climate