Scientific Computing Group

Presented by

Ricky A. Kendall

Scientific Computing
National Center for Computational Sciences

Scientific Computing

Scientific Computing facilitates the delivery of leadership science by partnering with users to effectively use computational science, visualization, and workflow technologies on OLCF resources to:

Port, tune, augment, and develop current and future applications at scale	Provide visualizations to present scientific results and augment discovery processes	Automate the scientific computational method

Visualization and data analytics

Visualization

Once users have completed their runs, the Visualization task group helps them make sense of the sometimes overwhelming amount of information they generate

- Viewing at a 30 ft × 8 ft PowerWall
- Cluster with GPUs for remote visualization

End-to-End Solutions

Researchers must analyze, organize, and transfer an enormous quantity of data. The End-to-End task group streamlines the workflow for system users so that their time is not eaten up by slow and repetitive chores

- Automate routine activities (e.g., job monitoring at multiple sites)

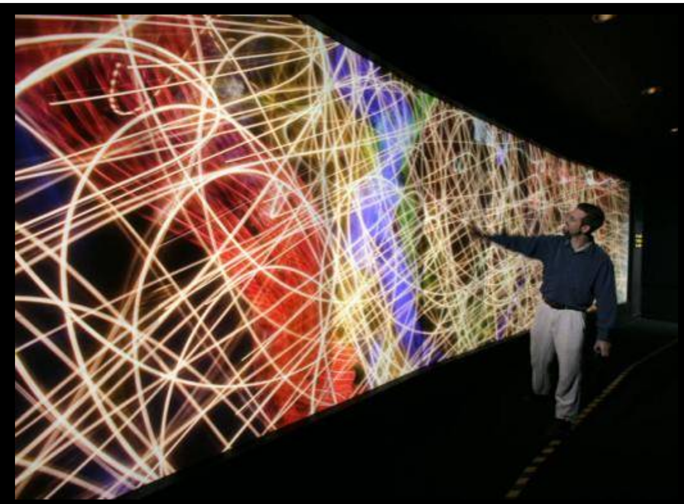

Scientific computing user support model

- "Whatever it takes" is the motto
- Share expertise in algorithms and application-development strategies
- Provide porting, tuning, optimization
- Help users in running applications, using application development tools and libraries
- Ensure application readiness by partnering with users to develop current next-generation applications
- Represent users' needs in OLCF planning and reporting exercises
  – Application requirements
  – Scientific progress and highlights
  – Issues with current resources

Expertise

The OLCF provides experts in user support, including PhD-level liaisons from fields such as chemistry, climate, physics, astrophysics, mathematics, numerical analysis, and computer science, who are also experts in developing code and optimizing it for the OLCF systems

Large projects are assigned liaisons to maximize opportunities for success on the leadership computing resources

Partnership with projects on LCF resources

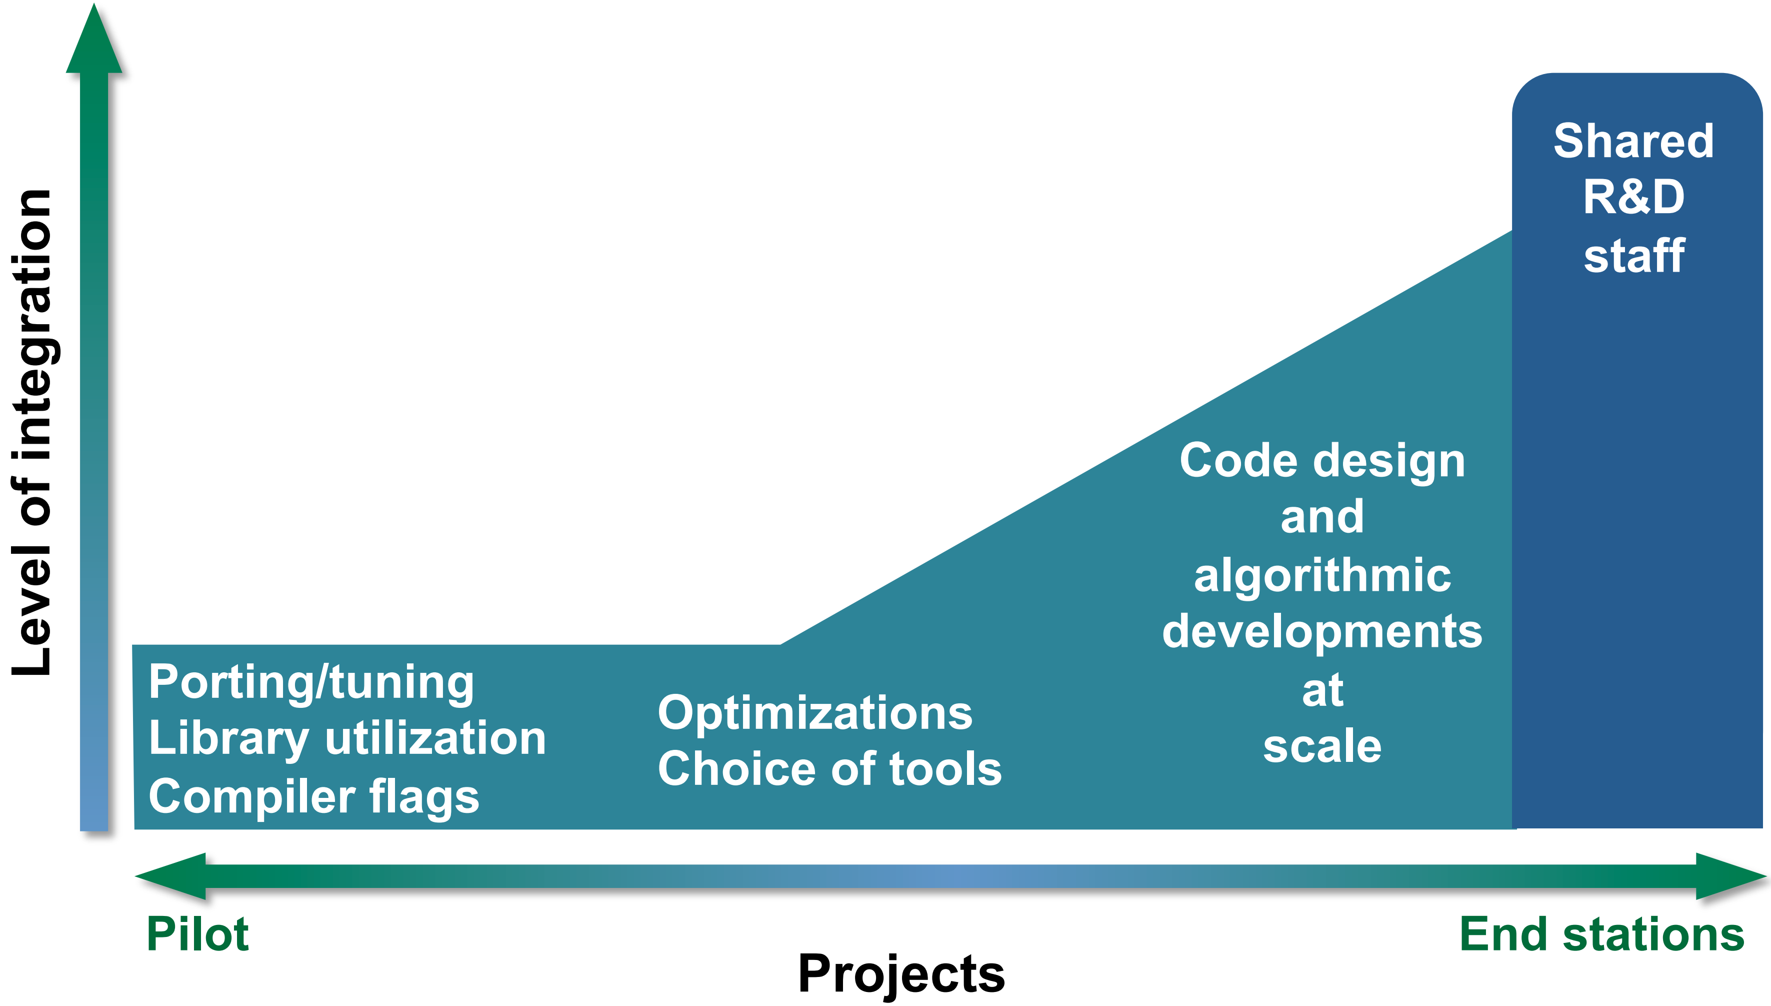

Liaison model helps maximize science

Project

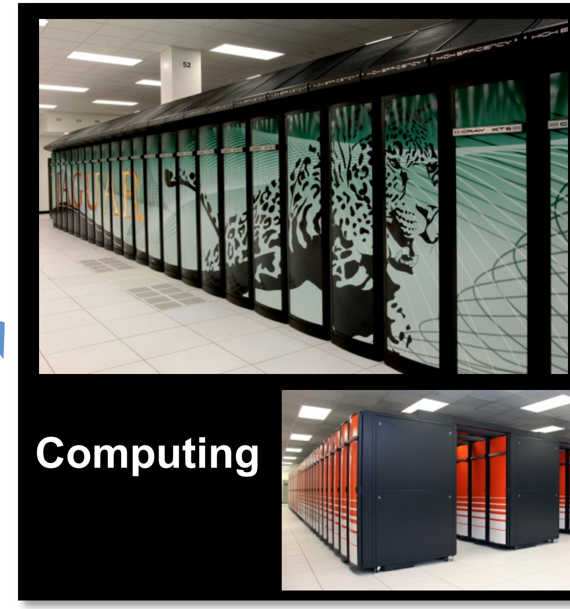

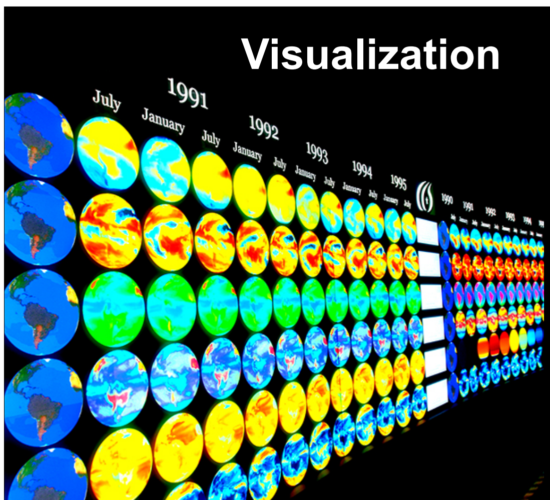

Whatever It Takes!

PI: Thomas Schulthess, ORNL

Materials and nanoscience

Role of nanoscale inhomogeneities in high-temperature superconductors

- **Investigate role of inhomogeneities on pairing mechanisms and transition temperature—petascale computing problem**
- **Develop materials-specific extensions to the Hubbard model to understand variations of transition temperatures—develop materials design tool**
- **Simulations of 2-D Hubbard model with disorder in local Coulomb interaction parameter (Ui = 1 ± dU) indicate Tc is indeed suppressed due to disorder**

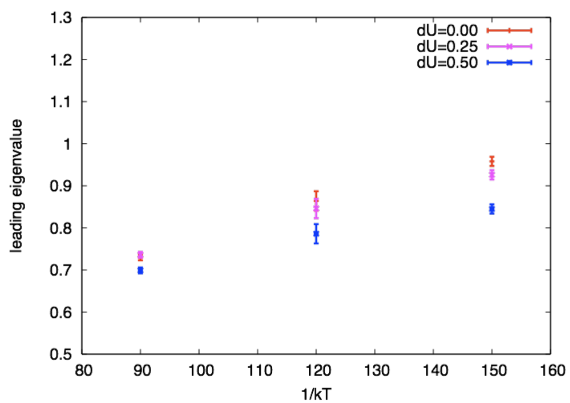

- **Further analysis of the role of inhomogeneities (substitutions, vacancies, etc.) is now possible in the simulations with 16- and 24-site clusters**

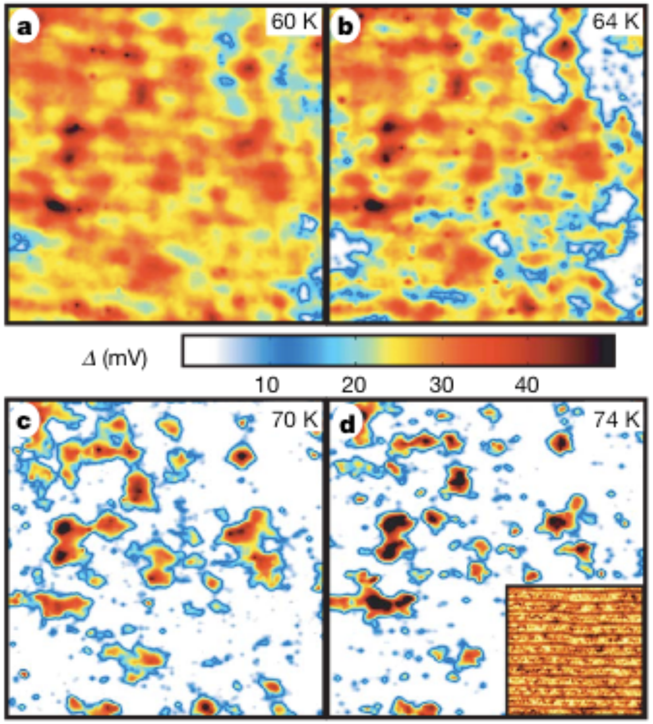

Temperature evolution of the superconducting gap taken on a 300 Å area of a cuprate with T_c = 65 K [reproduced from Gomez et al., *Nature* **447**, 569–572 (2007)]. The gap varies spatially on a scale of 1–3 nm and persists in some regions to temperatures well above T_c, as can be seen from panels *c* and *d*

Simulations at 31K cores on Jaguar perform >50% of peak (Gordon Bell submission)

PI: Jihui Yang, General Motors

Materials and nanoscience

Nanostructural features in high-performance thermoelectric materials

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

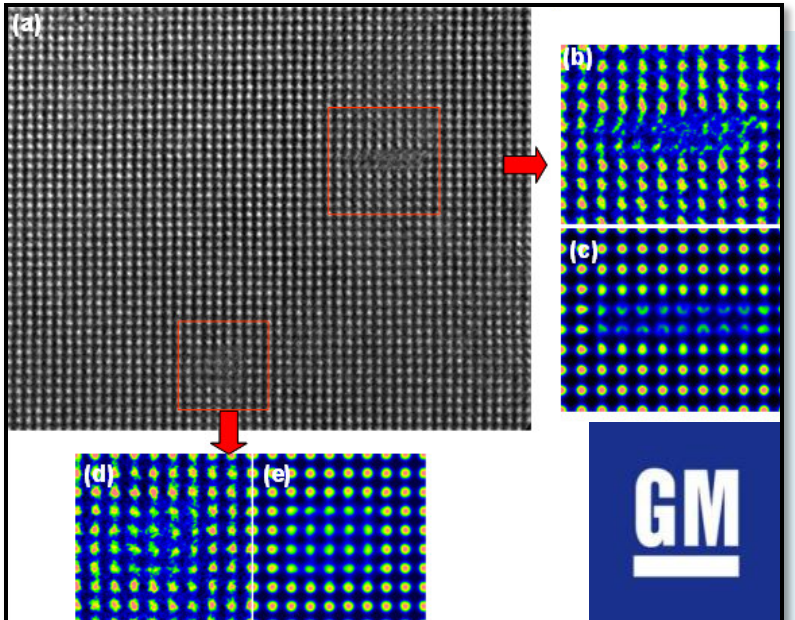

Of broad interest: The atomic structure of this material and the physical origin of its high thermoelectric nature are still under debate. This study will eventually explain these problems and lead to the discovery of new novel thermal-electric materials

- **Only 25% of auto fuel energy is used 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**
- **Largest-ever simulation—1728-atom supercell—made possible at LCF**
- **GM team identified the atomic configurations of the recently discovered $AgPb_mSbTe_{m+2}$ with DFT total energy calculations**
  - **Recent simulations of the $AgSbTe_2$ nanocluster embedded in the PbTe bulk material agree well with experimental measurements and open the way to explaining the origin of their outstanding thermal-electric performance**

OLCF liaison contributions

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

Astrophysics

Three-dimensional supernova explosions simulations with high-fidelity physics

- The world's first 3-D supernova simulations incorporating multifrequency neutrino transport are currently under way on Jaguar
- 304×152×76 spatial mesh running on 11,552 CPUs
- Initial results are promising, but full evolution of the explosion epoch will require millions of hours

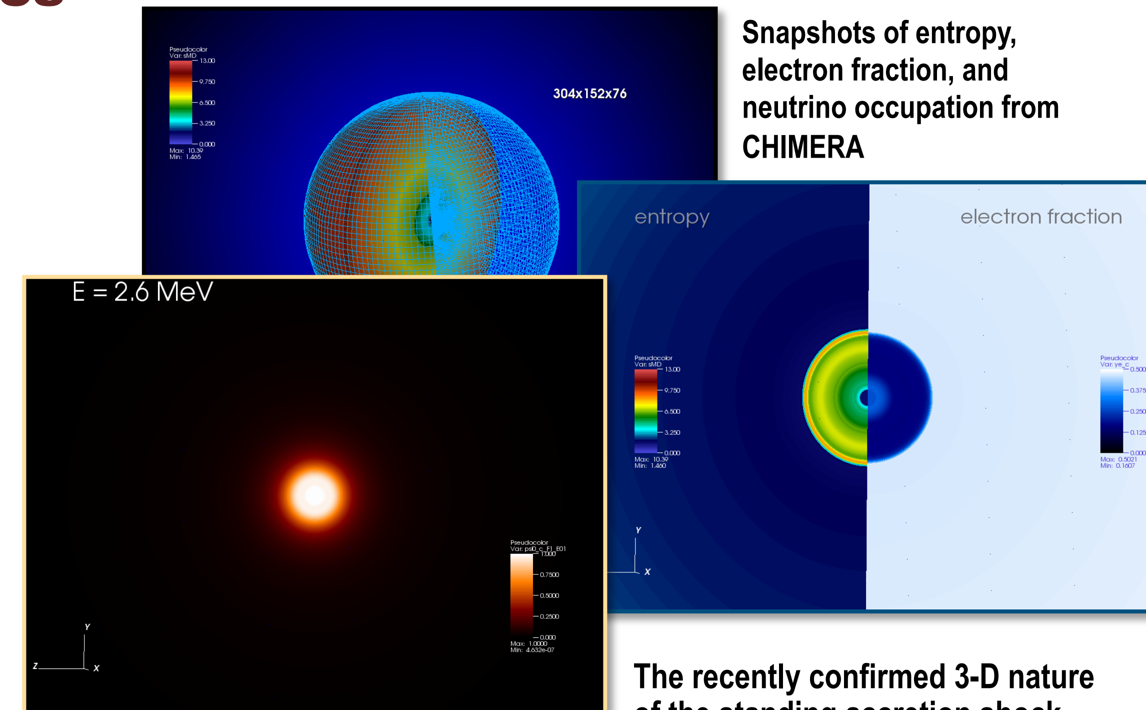

Snapshots of entropy, electron fraction, and neutrino occupation from CHIMERA

The recently confirmed 3-D nature of the standing accretion shock instability and the known 3-D nature of convective overturn make performing supernova simulations in 3-D essential

OLCF liaison contributions

- Implementing efficient, collective I/O
- New visualization tools
- Improved software management

Astrophysics

Via Lactea II: The Milky Way dark matter halo

PI: Piero Madau, UC–Santa Cruz

- For almost 80 years, astrophysicists have known that the movement of stars in galaxies must be dominated by matter they could not see (“dark matter”)
  - Less than 1/5 of matter in the universe is “visible”; the rest is dark matter!
  - Dark matter has no interaction with regular matter except through gravity
  - Yet there is so much dark matter in the universe, it controls the lives of stars and galaxies
- What would dark matter look like if we could see it?
  - Use a leadership computer to simulate dark matter evolving over billions of years to model a galaxy such as our own Milky Way; the envelope of dark matter around the spiral galaxy is called a *dark matter halo*
  - Follow a galaxy’s worth of dark matter through nearly the entire history of the universe, dividing the dark matter into more than a billion separate parcels, where each dark matter parcel is >4000 times as massive as the Sun
- Use the PKDGRAV2 code to focus entirely on the gravitational interaction between a billion dark matter parcels
- Over time, dark matter becomes clumped as gravity pulls it together
  - The clumps enlarge and are pulled together until they form large halos of dark matter massive enough to host galaxies
- Key question: Will smaller dark matter clumps remain identifiable or will they smooth out?
  - Recent billion-parcel simulation on Jaguar provided enough resolution to verify that clumps and subclumps do indeed survive
  - “What we found is that the survival fraction is quite high,” Piero Madau (UC–Santa Cruz)
- Can the dark matter halo results be verified? Yes!
  - Gamma ray bursts: Caused when dark matter and antimatter particles bump into each other (prime target for the recently launched GLAST satellite)
  - Gravitational lensing: Gravity exerted by galaxies along a line of sight bends light

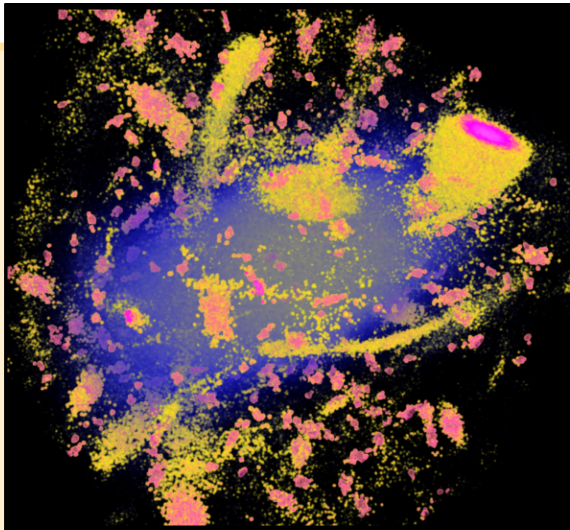

Dark matter color coded by momentum. This “phase space plot” allows researchers to trace both the aggregation of the dark matter and velocity field

OLCF liaison contributions

- Determining best processor count/wall-clock time ratio
- Data movement and visualization

PI: Stan Woosley, Stanford

Astrophysics
First principles studies of Type Ia supernovas

How do thermonuclear runaways begin in the heart of a white dwarf?

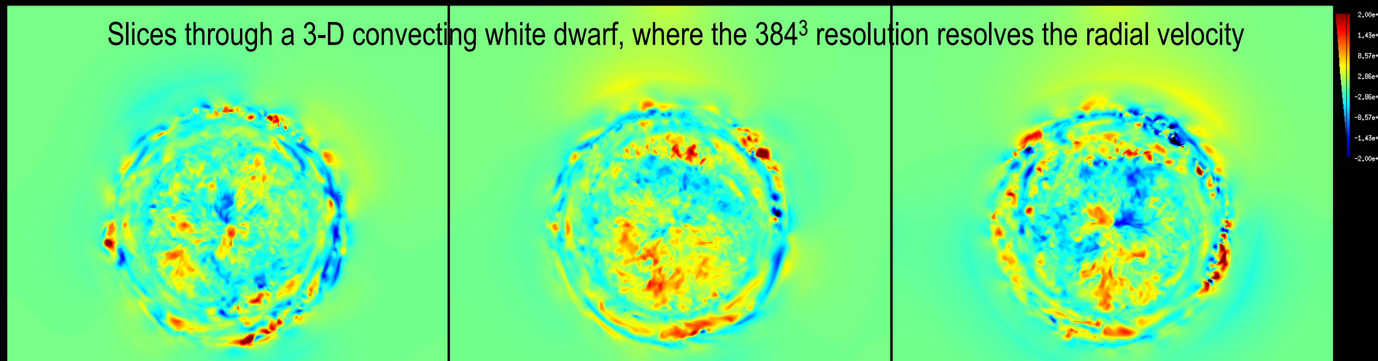

Slices through a 3-D convecting white dwarf, where the 384^3 resolution resolves the radial velocity

- Study all the key stages of Type Ia supernovas
  - The long-time convection that leads to ignition of the first flames
  - Propagation of the resultant flame(s) through the star leading to the explosion itself
  - And finally, the radiation-dominated phase at the end of the explosion
- By acting as standard candles, Type Ia supernovas have been at the forefront of a revolution in modern cosmology, leading to the discovery that the expansion rate of the universe is accelerating
- Perform long-time evolution simulations necessary to follow the convection preceding the explosion and the ignition of the thermonuclear flame
  - Low Mach number code MAESTRO has been used to model the long convective epoch that leads to ignition of a thermonuclear flame that completely disrupts the star in a Type Ia supernova
- Knowledge of the initial convective overturn is necessary to determine the ultimate evolution of the flame and the outcome of the explosion, including the element production

OLCF liaison contributions

- Methods for archival data storage

Climate

Modeling the full Earth system

New standard and low-emission climate scenarios at higher resolution

Final stage of ocean spin-up for full Earth System Model

100 years of ocean spin-up

20 years of carbon-cycle spin-up

Approaching steady state for sulfur cycle

Approaching release of CCSM4, the first Earth System Model

OLCF liaison contributions

- Advanced visualizations of new climate couplings
- Tuning of LCF infrastructure for climate requirements: special queues, reservations, dedicated file systems
- Weekly conference calls with developers to triage issues
- Development of timing libraries to assess performance and balance components

Simulated time evolution of the atmospheric CO_2 concentration originating from the land's surface

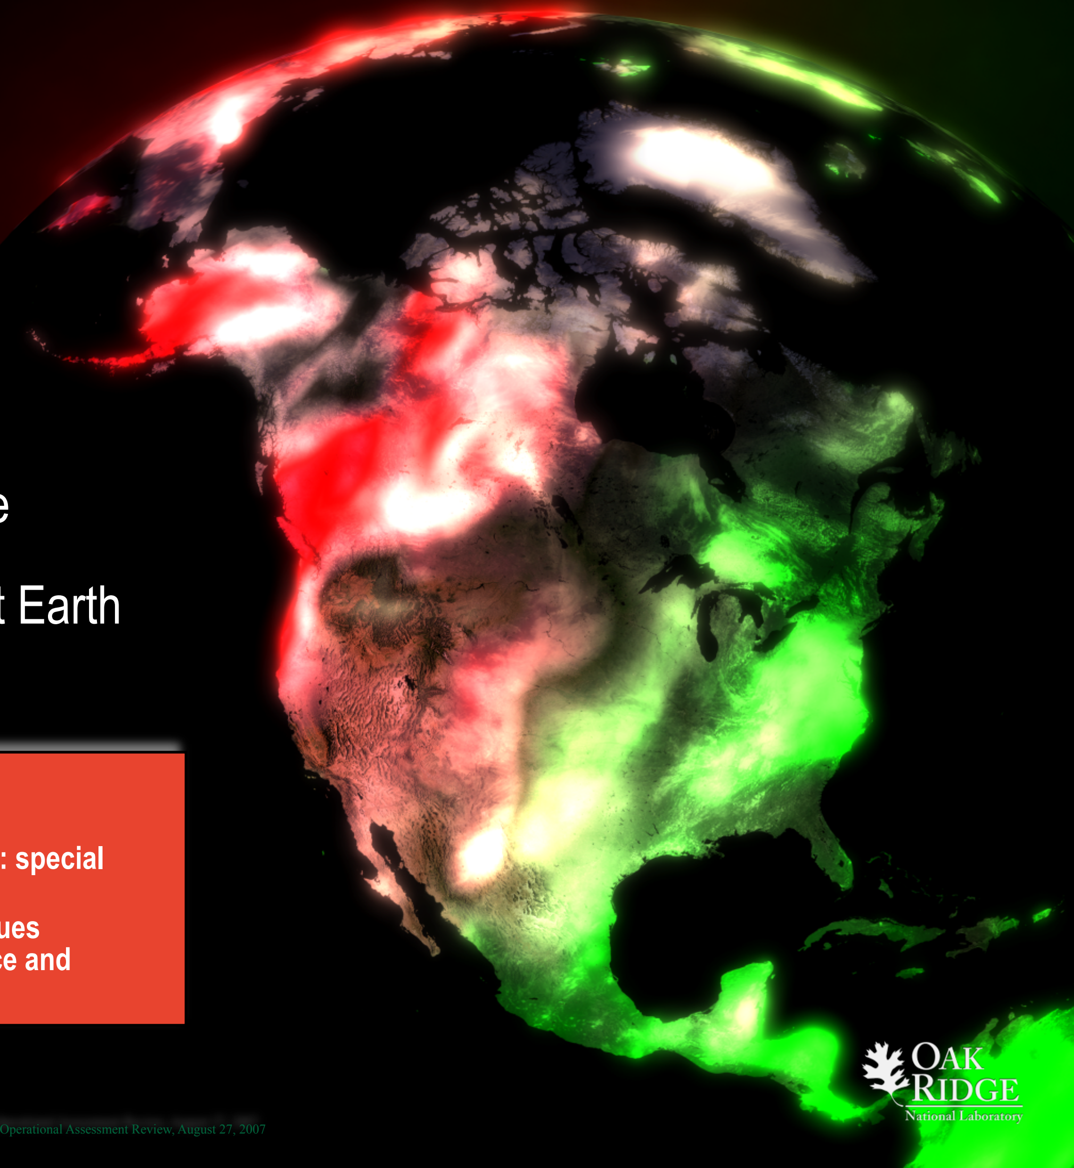

Max Suarez, William Putnam, NASA/GSFC
Shian-Jiann Lin, GFDL NOAA

Weather and climate

Non-hydrostatic weather forecasts using the GEOS-5 Cubed-Sphere FV core

Science Objectives and Impact

- **Strategy:** Predict high-impact weather events with ultra-high resolution
- **Driver:** What is the impact of weather events on climate variability (and vice versa)?
- **Objective:** Undertake two high-resolution 10-year Atmospheric Model Intercomparison Project simulations and compare with coarser climate resolutions
- **Impact:** Provide a capability for dynamical regional downscaling of climate change predictions

Science Results

- Difference between current operational and Jaguar/XT5 supportable resolutions is remarkable
  - von Karman and cloud streets now resolved
- 5-day forecasts of Gordon and Helene hurricanes in Sep 06 show @ 7 km resolution (c1440) and compare well against NHC data for track *and* intensity
  - "These results have put us in a different mindset for planning our future directions in modeling!" – Suarez
- **OLCF contribution**: Compiler utilization & scheduling assistance

Current standard for operational forecasts

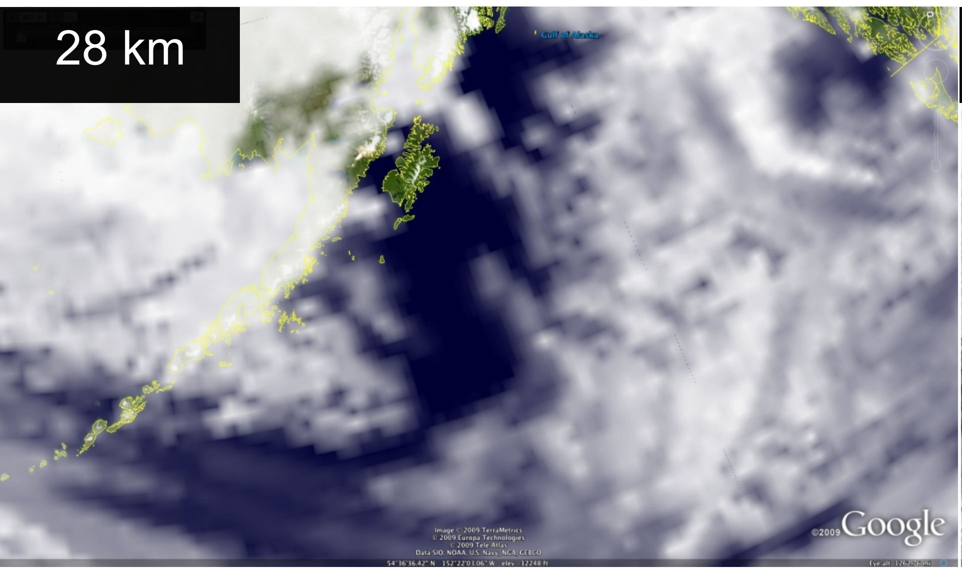

Routine on Jaguar/XT5: 8 times more resolution!

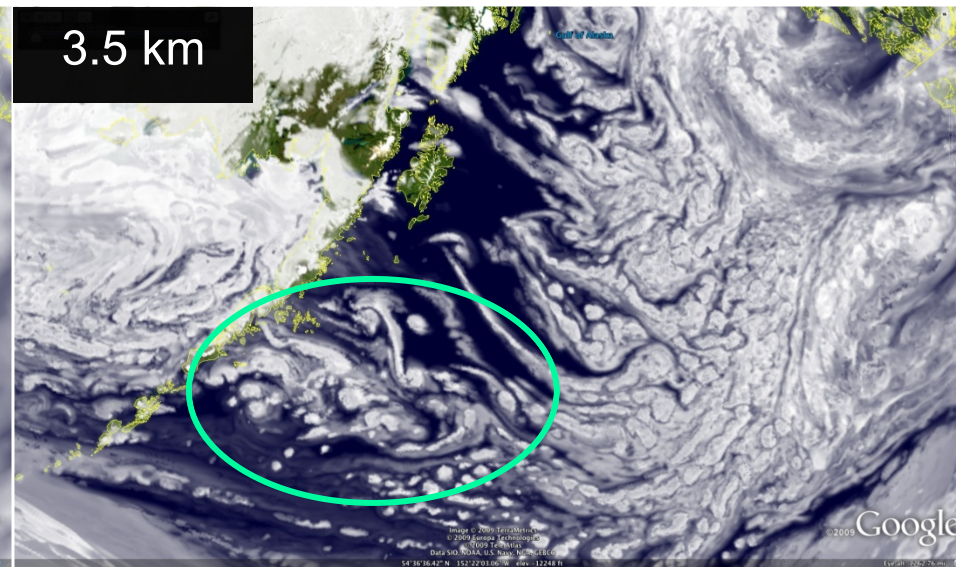

Enough resolution now to support von Karman vortex streets:
an important internal forcing mechanism for both weather and climate

V. Balaji, NOAA/GFDL

Climate

Coupled high-resolution modeling of the Earth system

Science Objectives and Impact

- **Strategy:** The goal of decadal prediction requires long-time integrations of state-of-the-art models at unprecedented resolution
- **Driver:** What is the possibility of abrupt climate change on the timescale of decades rather than centuries?
- **Objective:** Use CM 2.5 for century simulations of the scientific issues associated with resolving mesoscale features in the atmospheric and ocean circulations, and its implications for understanding of forced/natural climate variability
- **Impact:** Provide reliable input for public policy relating to increasingly destructive storms in a warming world

Science Results

- CM2.5 vs CM2.1: increased ability to capture atmosphere-ocean covariability on mesoscale
- Recent CM2.5 runs on Jaguar/XT5 show agreement with observations of winds being stronger over warmer waters, weaker over cooler
- Finer-scale couple physics leading to qualitatively better depiction of tropical convection and climate in CM2.5
- **OLCF contribution:** Tuning and optimization, scheduling on resources (special queues)

Arabian sea winds and sea surface temperature

(Figure courtesy Whit Anderson, Rusty Benson, Tom Delworth, Tony Rosati and Gabe Vecchi, NOAA/GFDL)

Sea Surface Temperature and Surface Zonal Wind Stress
NOAA/GFDL Coupled Climate Models

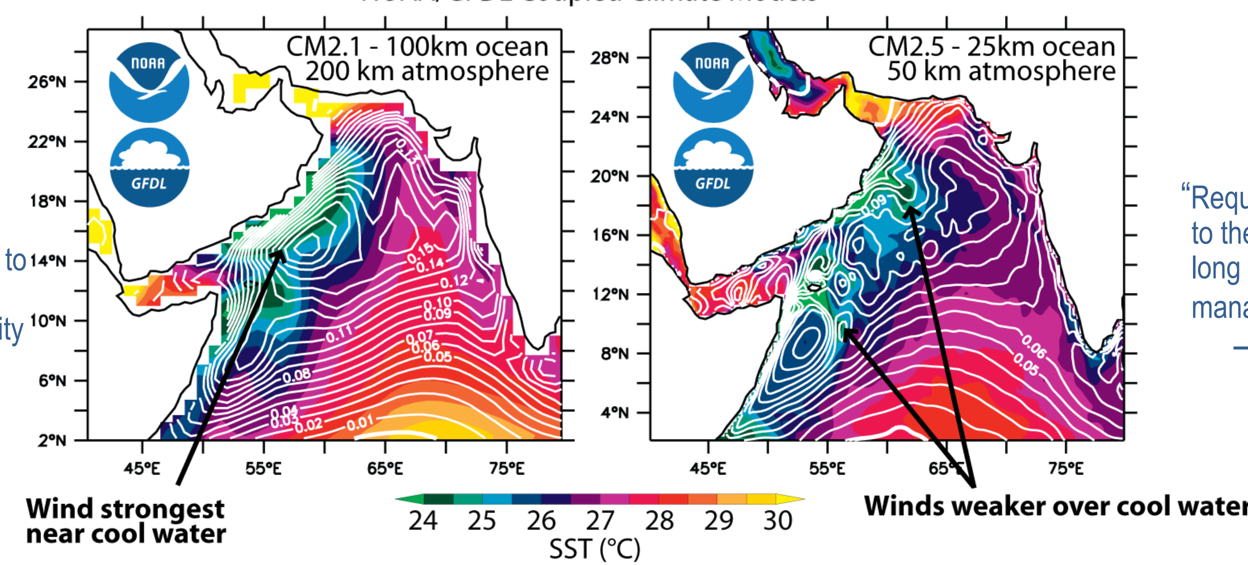

The figure shows "the key feature of what increased resolution buys us: the ability to capture in a model atmosphere-ocean covariability on the mesoscale."
-V. Balaji, NOAA/GFDL

"Requires sustained access to the petascale for very long runs … currently being managed by a heroic effort"
–V. Balaji, NOAA/GFDL

PI: Fred Jaeger, ORNL

Fusion

Producing new insights for RF heating of ITER plasmas

Three-dimensional simulations of RF heating in the ITER fusion reactor as well as in present tokamaks (NSTX) shed new light on the behavior of superheated ionic gas

- When deuterium (D) or helium-3 (^{3}He) are used to damp the launched waves, they are accelerated to high energies, forming supra-thermal tails that significantly affect the wave propagation and absorption
- Energetic minority D ions enhance the fusion reaction rate
- Energetic ^{3}He ion tails form on both the tritium and ^{3}He distributions

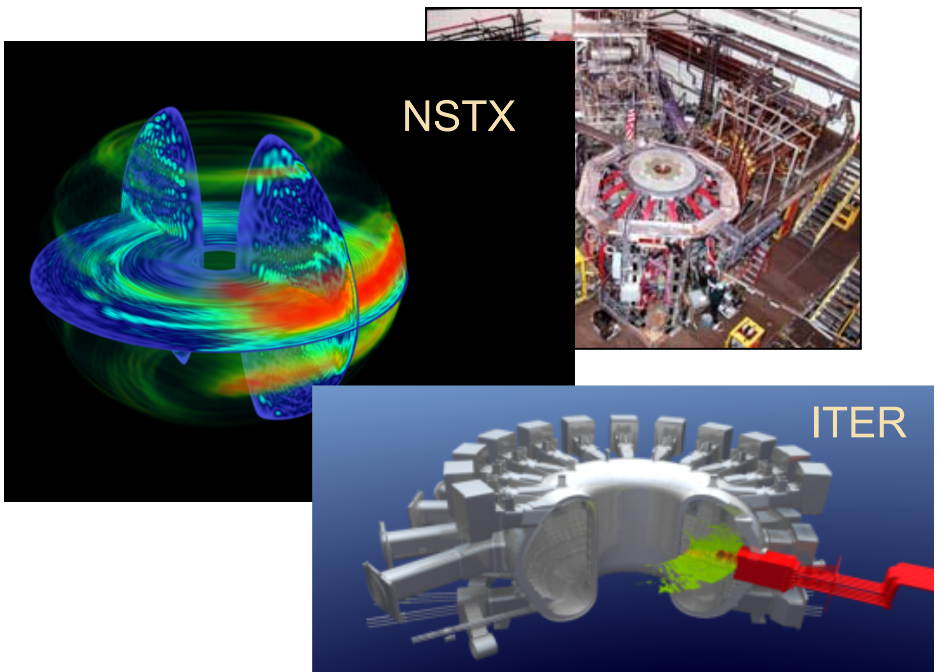

- **3-D simulations reveal new insights**
  - **"Hot spots" near antenna surface**
  - **"Parasitic" draining of heat to the plasma surface in smaller reactors**
- **Work pushing the boundaries of the system (28,900 cores, 154 TF) and demonstrating**
  - **Radial wave propagation and absorption**
  - **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**

OLCF liaison contributions

- **Converted HPL from double real to double complex and replaced ScaLAPACK**
- **Acquired new version of BLAS from TACC**
- **Net performance gain of a factor of two**

PI: Patrick Diamond, UCLA

Fusion

Breakthrough simulation of turbulent transport in ITER fusion plasmas

- **Transition to favorable scaling of confinement for both ions and electrons now observed in GTC simulations for ITER plasmas**
  - GTC simulation of electron turbulence used 28,000 cores for 42 hours producing 60 TB of data in a dedicated run on Jaguar as part of early-access 263 TF transition to operations
  - *[Yong (UCI), "Researchers Conduct Breakthrough Fusion Simulation,"* HPCWire *(July 14, 2008)]*
- **GTC simulation critically tested quasilinear theory used in transport model**
  - *[Z. Lin et al.,* Phys. Rev. Lett. *99, 265003 (December 2007)]*
- **GTC simulation predicted properties of turbulent transport of fusion products**
  - *[W. Zhang, Z. Lin, and L. Chen,* Phys. Rev. Lett. *101, in press]*
- **GTC collaboration team**
  - Led by Z. Lin of UC–Irvine and S. Klasky of ORNL

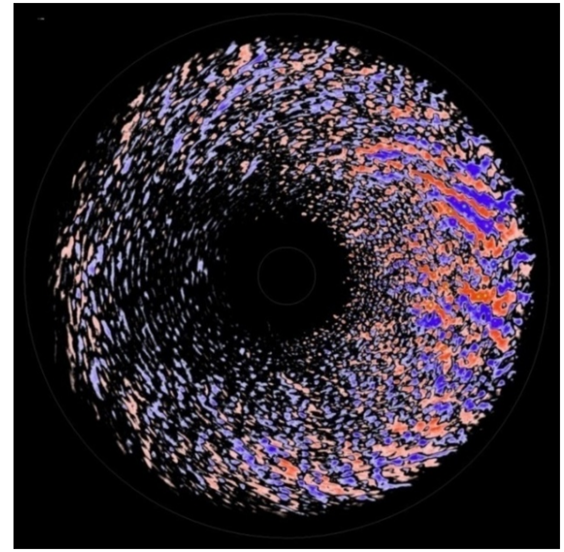

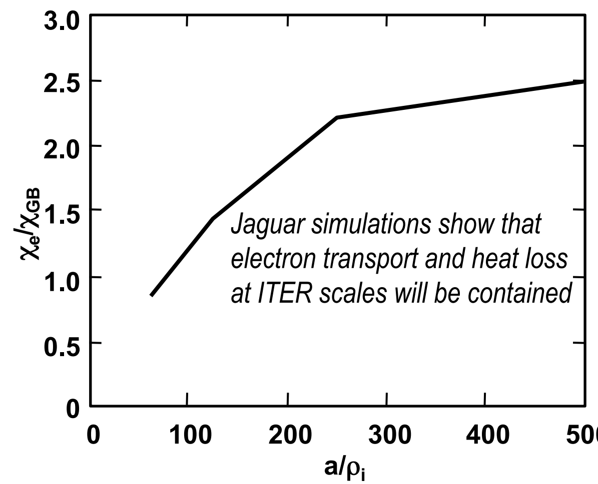

OLCF liaison contributions

- ADIOS for high-performance I/O. Wrote >60 TB of data at over 20 GB/second on 29K cores

C.S. Chang, New York Univ.

Fusion energy science

Core-edge nonlocal profile formation of tokamak plasma

Science Objectives and Impact

- **Strategy:** Obtain first-principles understanding on the fusion efficiency of tokamak reactor plasma using extreme scale HPC
- **Driver:** Edge conditioning, which may be the only mean to control the reactor plasma, has been ubiquitously observed to improve the core plasma profile dramatically in experiments
- **Objective:** Multiscale nonlocal core-edge simulation of the combined ion-temperature-driven turbulence dynamics and the background ion-temperature profile evolution in realistic DIII-D device geometry in a day
- **Impact:** Open up a door to practical first-principles predictive simulation capability for ITER and DEMOs

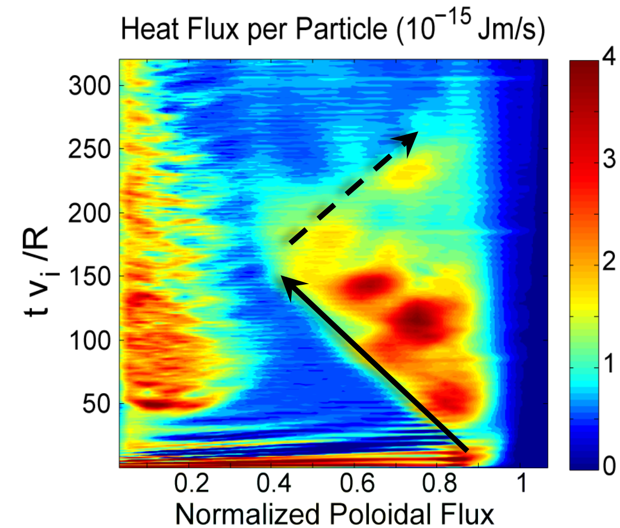

XGC1 Performance

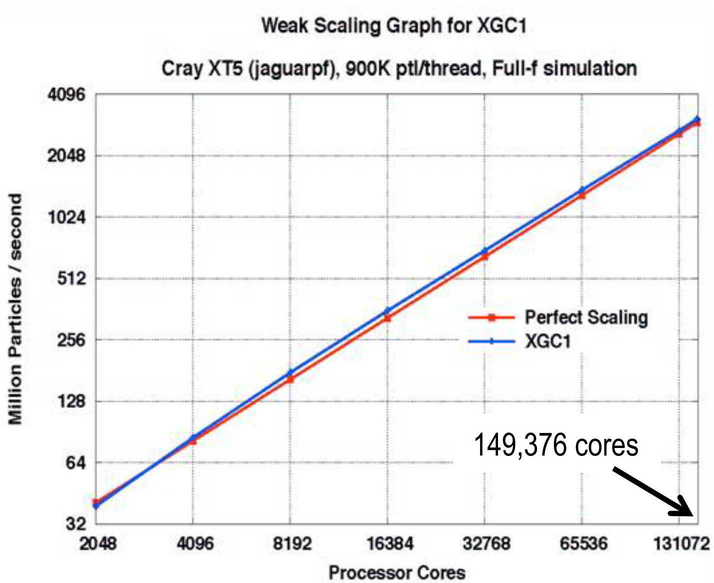

Science Results

- World's only whole volume turbulence simulation in realistic tokamak geometry
- Core turbulence is sum of the incoming intensity from edge and the ambient local fluctuations, self-organizing the temperature gradient in turbulence propagation time scale (similar to experiments)

OLCF contribution

MPI/OpenMP architecture
Improved I/O speed in ADIOS

Image caption

Shown above is the turbulent heat flux in time and radius (normalized flux). Turbulence propagates from edge to core (solid arrow), induces outward heat flux (dashes arrow), and leads to an eventual new self-organized nonlocal state

Turbulent combustion

Jackie Chen, SNL

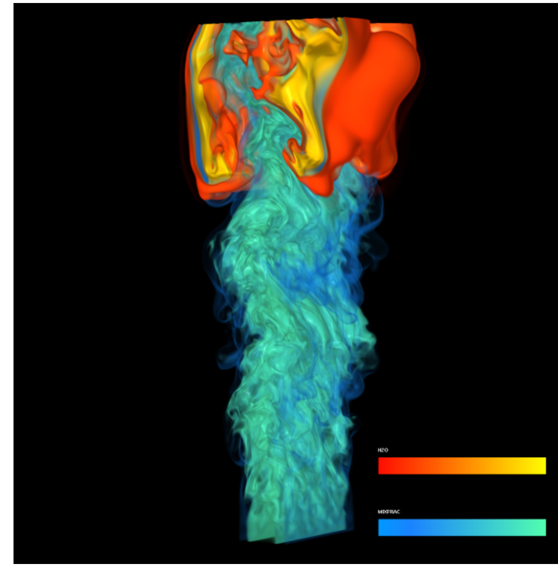

Science Objectives and Impact

- **Strategy:** Understand and model the effects of complex fuel chemistry in combustion engine conditions through high fidelity turbulence simulation
- **Driver:** Move to biofuel and coal-derived liquid fuels: Climate change and energy security
- **Objective:** Use of adaptive chemistry reduction and vectorization of new, larger chemistry models at low-temperature high pressure; efficient scaling to 140,000 cores
- **Impact:** Benchmark data for turbulent, igniting *n*-heptane and dimethyl ether flames at diesel engine pressures for model development and validation

Application Performance

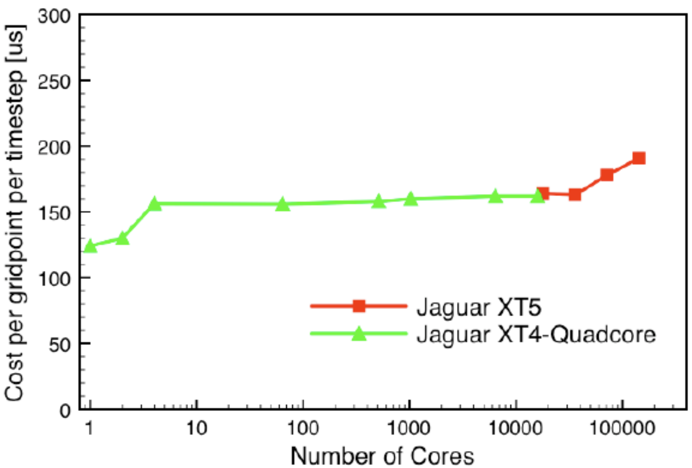

Science Results

- DNS of non-premixed *n*-heptane and dimethyl ether fuel jets issuing into high pressure, high temperature conditions to study the influence of multistage ignition chemistry on lifted flame stabilization under diesel thermochemical environment
- First 3-D simulation to fully resolve flame and ignition features including chemical composition, temperature profile, and turbulence flow characteristics

OLCF contribution

Application support and performance optimization for JaguarXT5 by R. Sankaran

Image caption

***n*-Heptane fuel jet (blue scale) and water vapor (red scale) indicating burned gas**

Biology

Cellulosic ethanol: Physical basis of recalcitrance to hydrolysis of lignocellulosic biomass

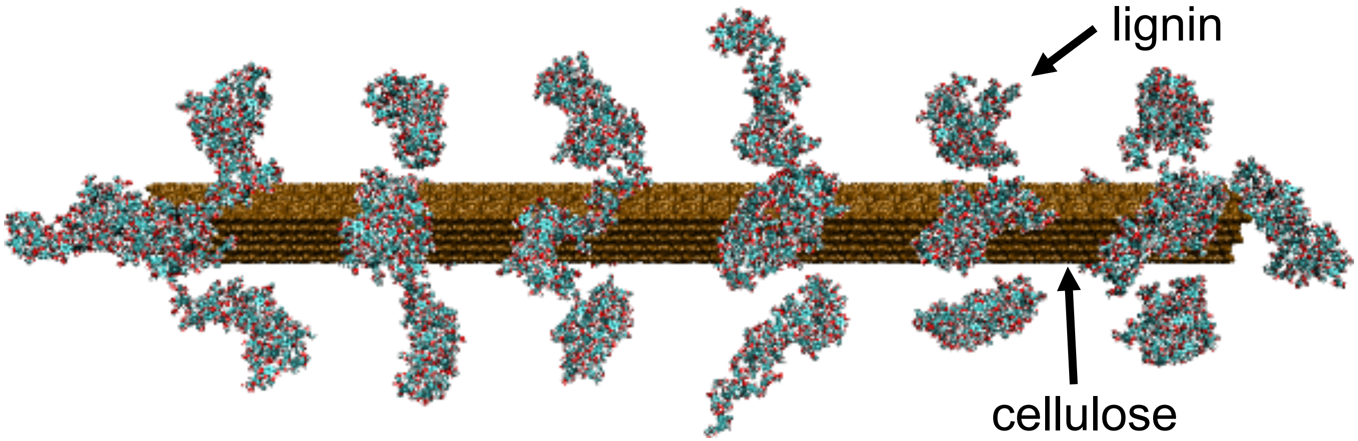

Simulations provide physical insight into physical origins of biomass recalcitrance and cellulosome organization and assembly

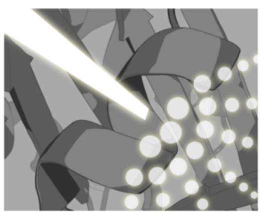

Center for Molecular Biophysics

ORNL

Loukas Petridis, Benjamin Lindner, and Jeremy C. Smith

See L. Petridis and J.C. Smith, *J. Comp. Chem*. In press.

- Understanding biomass structure is key to overcoming recalcitrance
- Large-scale molecular dynamics simulation (1–3M atoms)
  - Currently using NAMD code with 1.9M atoms for lignocellulose (on 6072 cores)
- Results to be used to interpret biophysical experiments
  - Lignin-cellulose simulations will be used to calculate small-angle neutron scattering intensities, which will be compared with experimental data obtained at the Spallation Neutron Source

Peter Lichtner, LANL

Geoscience

Modeling reactive flows in porous media

Science Objectives and Impact

- **Strategy:** Evaluate remediation strategies for cleanup of uranium plume
- **Driver:** Potentially save billions of dollars in cleanup costs
- **Objective:** Perform high-fidelity simulations of the Hanford 300 Area and demonstrate scalability of PFLOTRAN to the petascale
- **Impact:** Stakeholders, DOE, and public are concerned that cleanup proceed to remediate contaminants along the Columbia River corridor —a resource for fish, recreation, and other uses

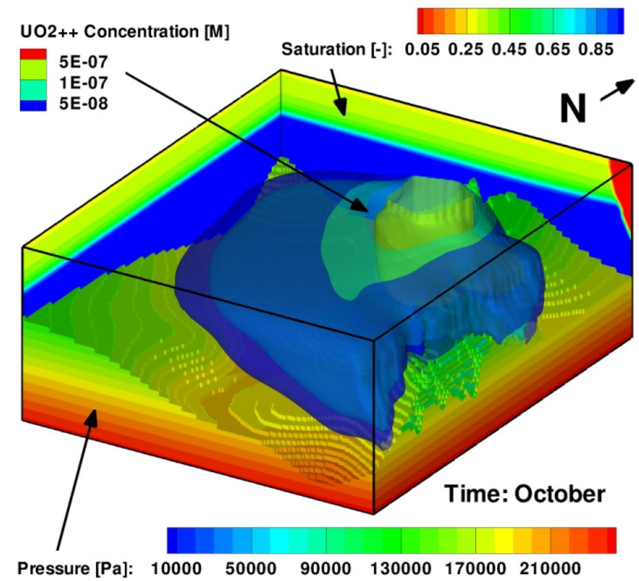

PFLOTRAN Performance

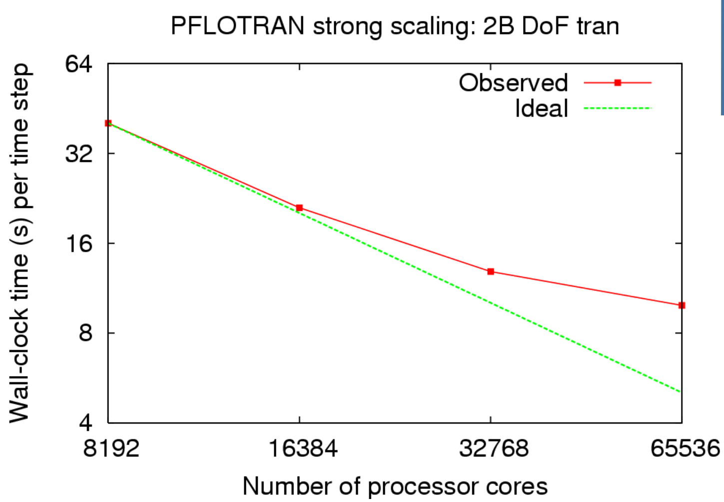

Science Results

- Demonstrated the ability of PFLOTRAN to scale to the petascale with complex nonlinear chemistry and highly transient fluid flow caused by fluctuations in the Columbia River stage
- Showed feasibility in using PFLOTRAN to model uranium dissipation at the Hanford 300 Area and evaluate remediation strategies

OLCF contribution

Ongoing development of multilevel solvers, parallel I/O performance optimization solver optimizations for full-machine runs on Jaguar

Image Caption (Top) Isopleths of the simulated uranium plume at the Hanford 300 Area after 1 year. The Columbia River is on the right face. Also shown is the pressure and saturation fields. (Left) Departure from linear scaling is due to high cost of MPI_Allreduce() operations; performance should improve with a reformulated BiCGStab implementation with only one MPI_Allreduce() per iteration

Nuclear physics

Nuclear structure/reactions: Computing nuclei from the ground up

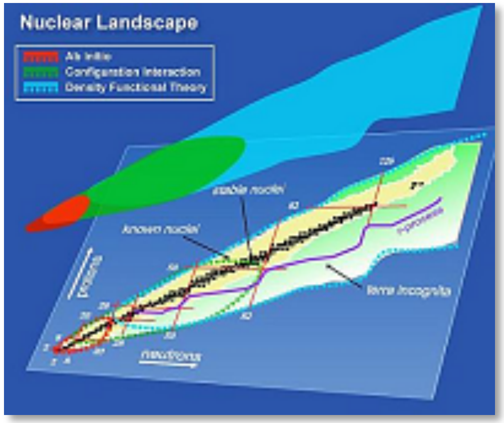

INCITE and UNEDF: The computer science part of the UNEDF SciDAC project consists of six individual collaborations between computer scientists and nuclear physicists working on various aspects of code development and scalability. The INCITE project provides cycles to UNEDF as appropriate

- Goal: Develop an understanding of all atomic nuclei, based on applying a number of computational approaches
- Leverage developments within the UNEDF (Universal Nuclear Energy Density Functional) SciDAC project to calculate nuclei using Greens Function Monte Carlo (GFMC), No Core Shell Model (NCSM), Coupled-Cluster (CC) methods, and nuclear Density Functional Theory (nDFT)

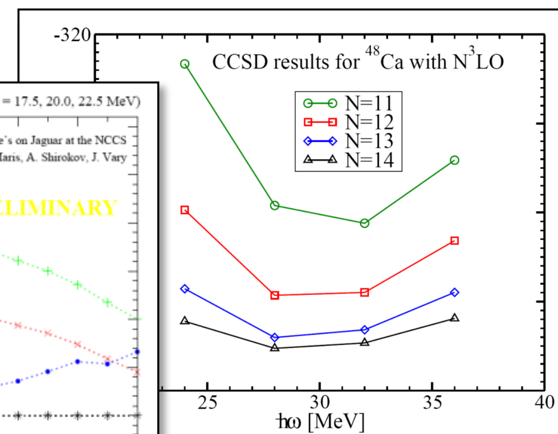

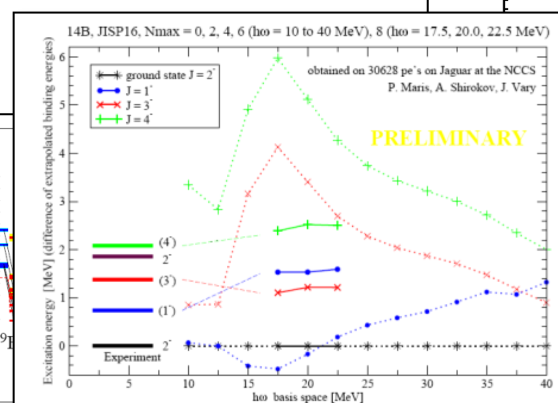

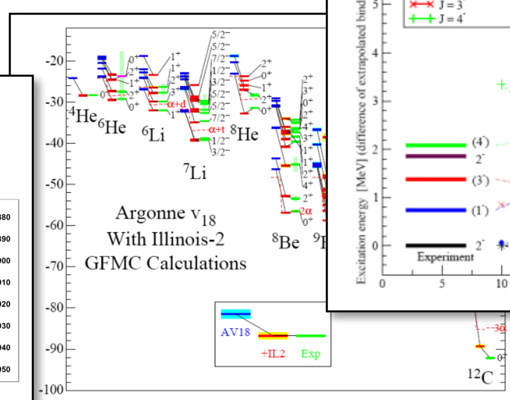

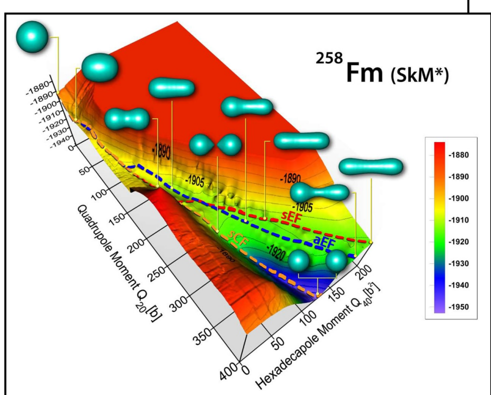

PI: Robert Harrison, ORNL

Chemistry

Organic molecular magnets from zwitterionic molecules

- Hybrid DFT calculations of the electronic structure of a new class of zwitterionic molecules (betaine derivatives) reveal an unexpected magnetic ground state
- Nanochains with a maximum length of 10.4 nm and over 1700 basis functions. Over 500 Jaguar processors used. Scaling with respect to system size is $O(N^2)$
- Origin of the magnetism is quite complex and not well understood as compared to more conventional magnetic systems
- Magnetic materials that are purely organic are rather rare, and the betaine derivatives promise to offer a new class of molecular magnetic materials with diverse potential applications including organic spintronics, information storage, and nanoscale sensors

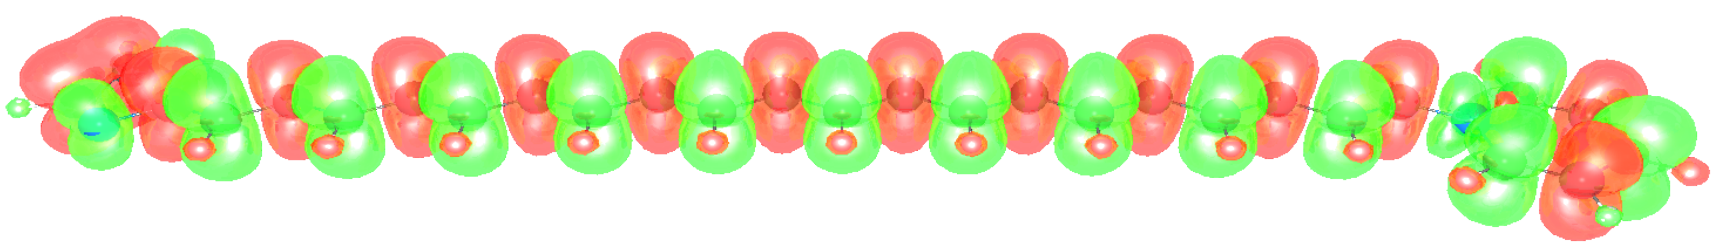

W.A. Shelton, E. Aprà, B.G. Sumpter, V. Meunier

PI: Madhava Syamlal, NETL

Engineering

Simulation of commercial-scale coal gasifier

- High-resolution simulation of commercial-scale transport gasifier using 10M cells
  - Simulation used carbonaceous chemistry for continuum modeling (C3M) module
  - C3M was modified to include multiple particle sizes
- Simulation shows new structure inside the gasifier
  - A detailed mapping of local hydrodynamics, temperature, and species distribution will allow engineers to optimize design of the gasifier
- Industrial partners are in the process of commercializing this technology and are using the results in their designs

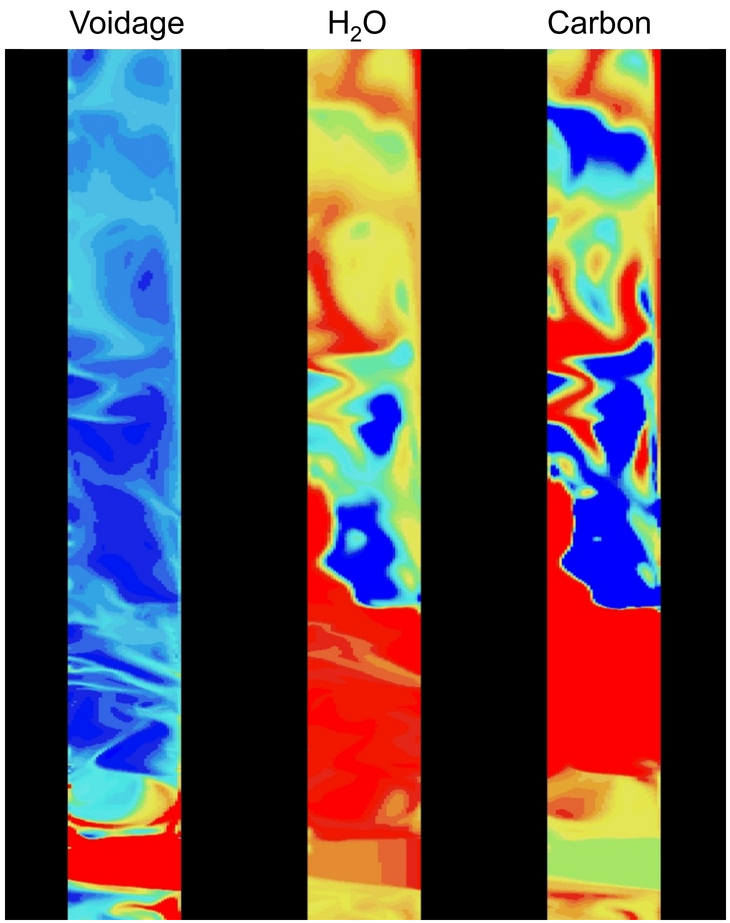

Ten-million-cell simulation of a commercial-scale transport gasifier on 1028 cores of Jaguar. Instantaneous concentrations at t ~ 0.9 second

OLCF liaison contributions

- Assisted in scaling studies
- Troubleshot run failures

Preparing for the future

Application requirements: Process and actionable results

OLCF Application Requirements Council (ARC)	• Stood up in 2006 • Established ARC charter and requirements–management process
OLCF 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 published September 2007 • New methods for categorizing platforms and application attributes devised and used in analysis: Guiding tactical infrastructure purchase and deployment • Best practice: Process being embraced and emulated by others

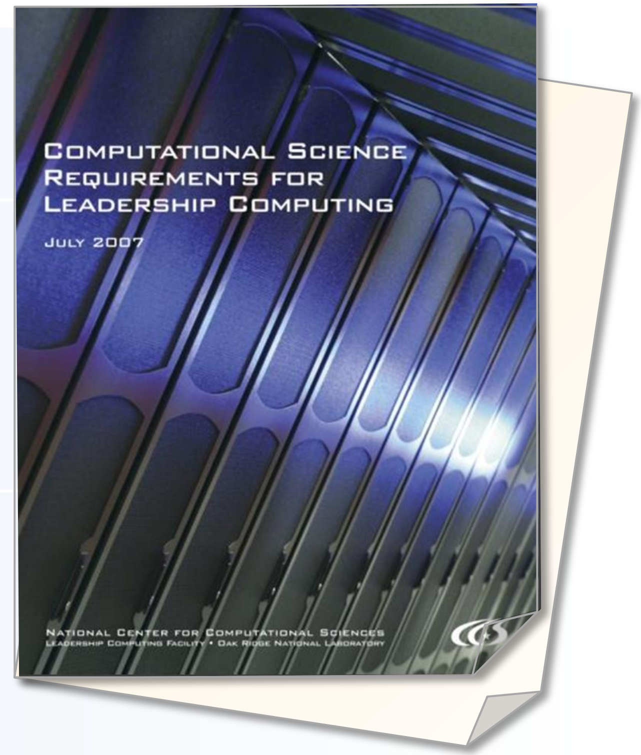

Innovation
Feedback loop for ensuring application readiness

Scientific Computing Group	Provides liaisons to application project teams
Application Requirements Council	Identifies application requirements
Technology Council	Decides how to best meet future application resource needs
Resource Utilization Council	Takes into account Science Team time constraints (e.g., upcoming meeting)

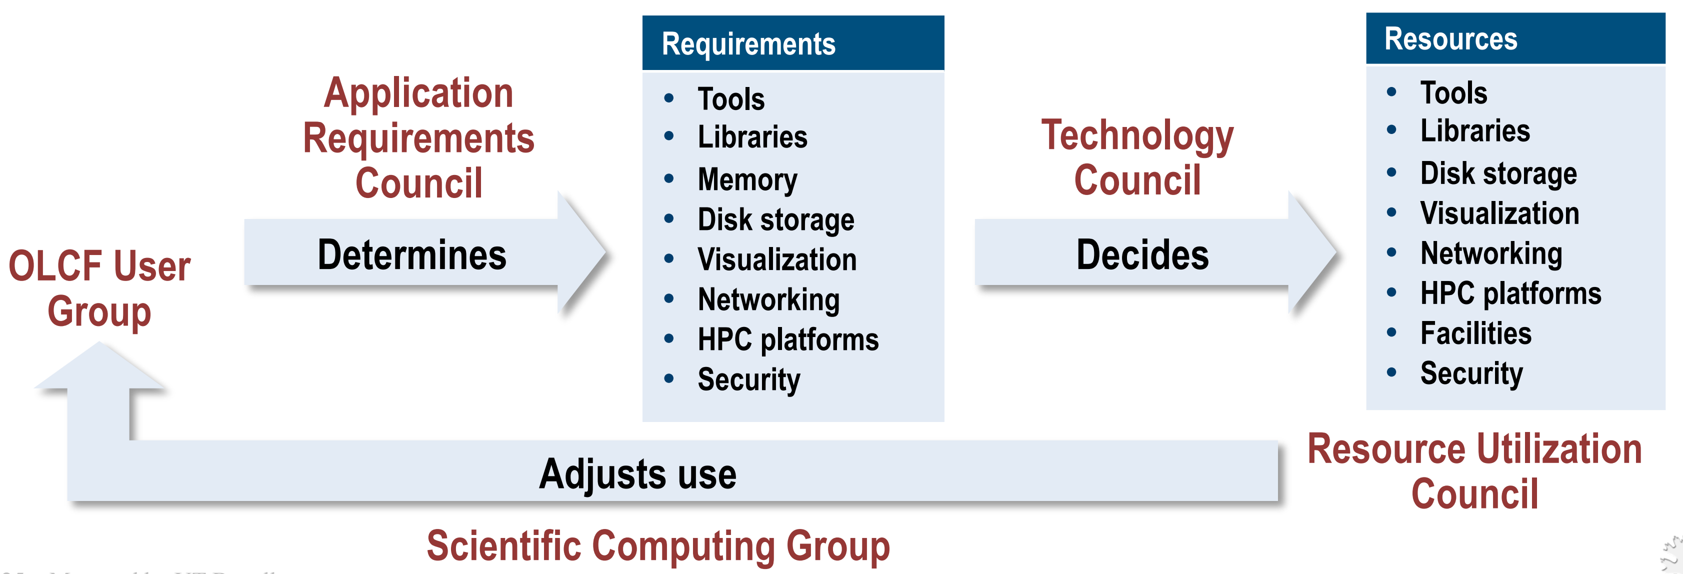

Software effectiveness on Jaguar/XT5

Improved computational science capabilities: Joule metric

- Four FY09 "Joule codes" have been exercised extensively on Jaguar/XT5 during ES period
  - They possess science drivers for both strong and weak scaling
  - A very successful set of performance improvements have been realized for all four codes (by virtue of this metric)
- Raptor
  - Chemically reacting, turbulent multiphase flows in complex geometries
  - Combustion modeling of turbulent flame dynamics and mixing processes
  - Performance improvement: MPI halo communication redesign/refactor
- VisIt
  - Open source visualization and analysis tool for processing massive data sets
  - 2005 R&D 100 award winner with >100K downloads
  - Performance improvement: MPI all-to-all communication redesign/refactor
- CAM
  - Fifth generation Community Atmosphere Model of the NCAR AGCM used in climate studies
  - Performance improvement: moving to a hybrid MPI/OpenMP node model
- XGC1
  - 5-D gyrokinetic particle-in-cell code designed to model the whole plasma dynamics in experimentally realistic device geometry
  - Performance improvement: moving to a hybrid MPI/OpenMP node model
- Much more in the final FY09 Joule report
  - Take away: pushing scalability as driven by the science led to unanticipated performance improvements and fixes
  - Scaling is hard! Even the best people working on mature are confronted with surprises

Contact

Ricky A. Kendall

Scientific Computing
National Center for Computational Sciences
(865) 576-6905
kendallra@ornl.gov