Computer Science, Mathematics and Computational Science

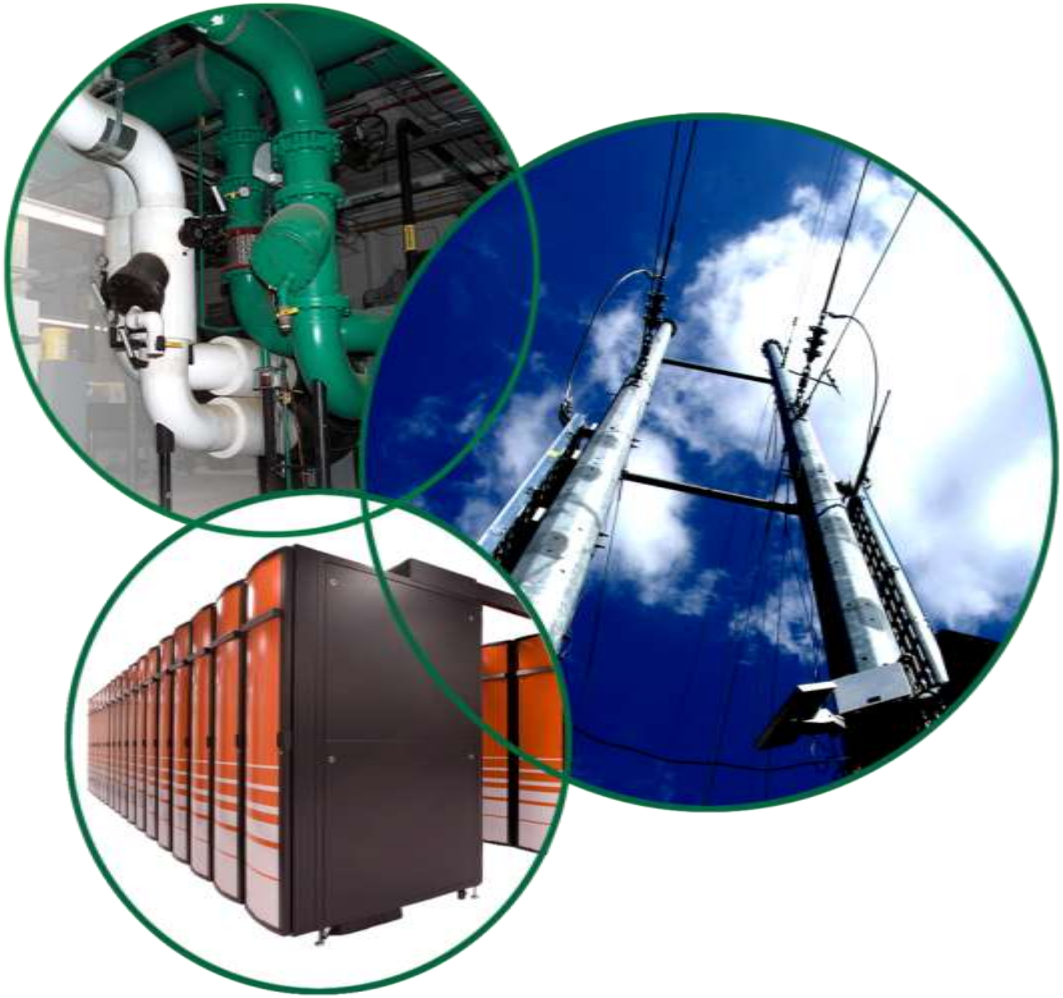

Arthur Maccabe
Director, Computer Science and Mathematics Division

September 2010

Managed by UT-Battele
for the Department of Energy

ORNL has a role in providing a healthy HPC ecosystem for several agencies

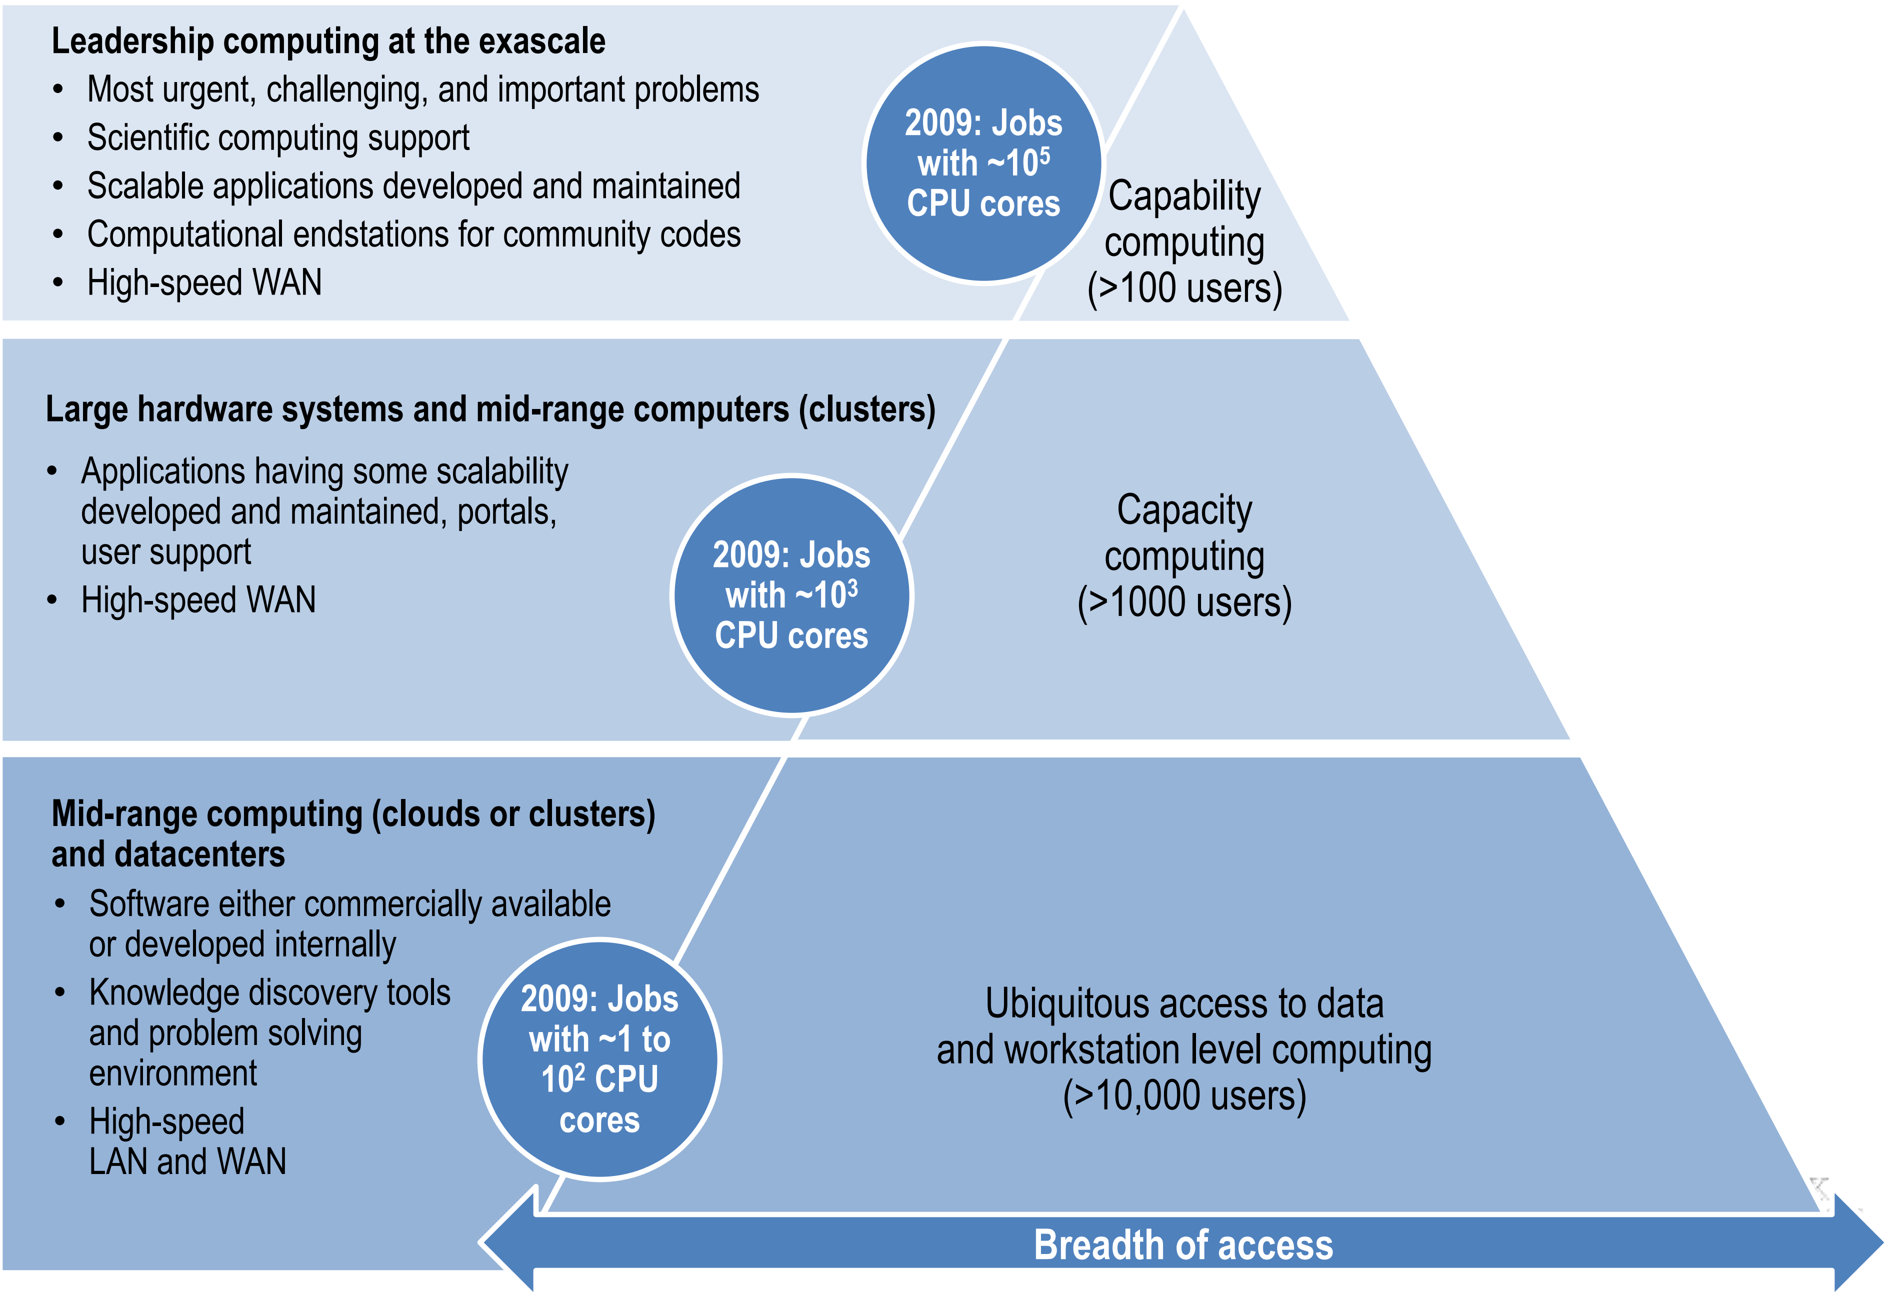

ORNL's computing resources

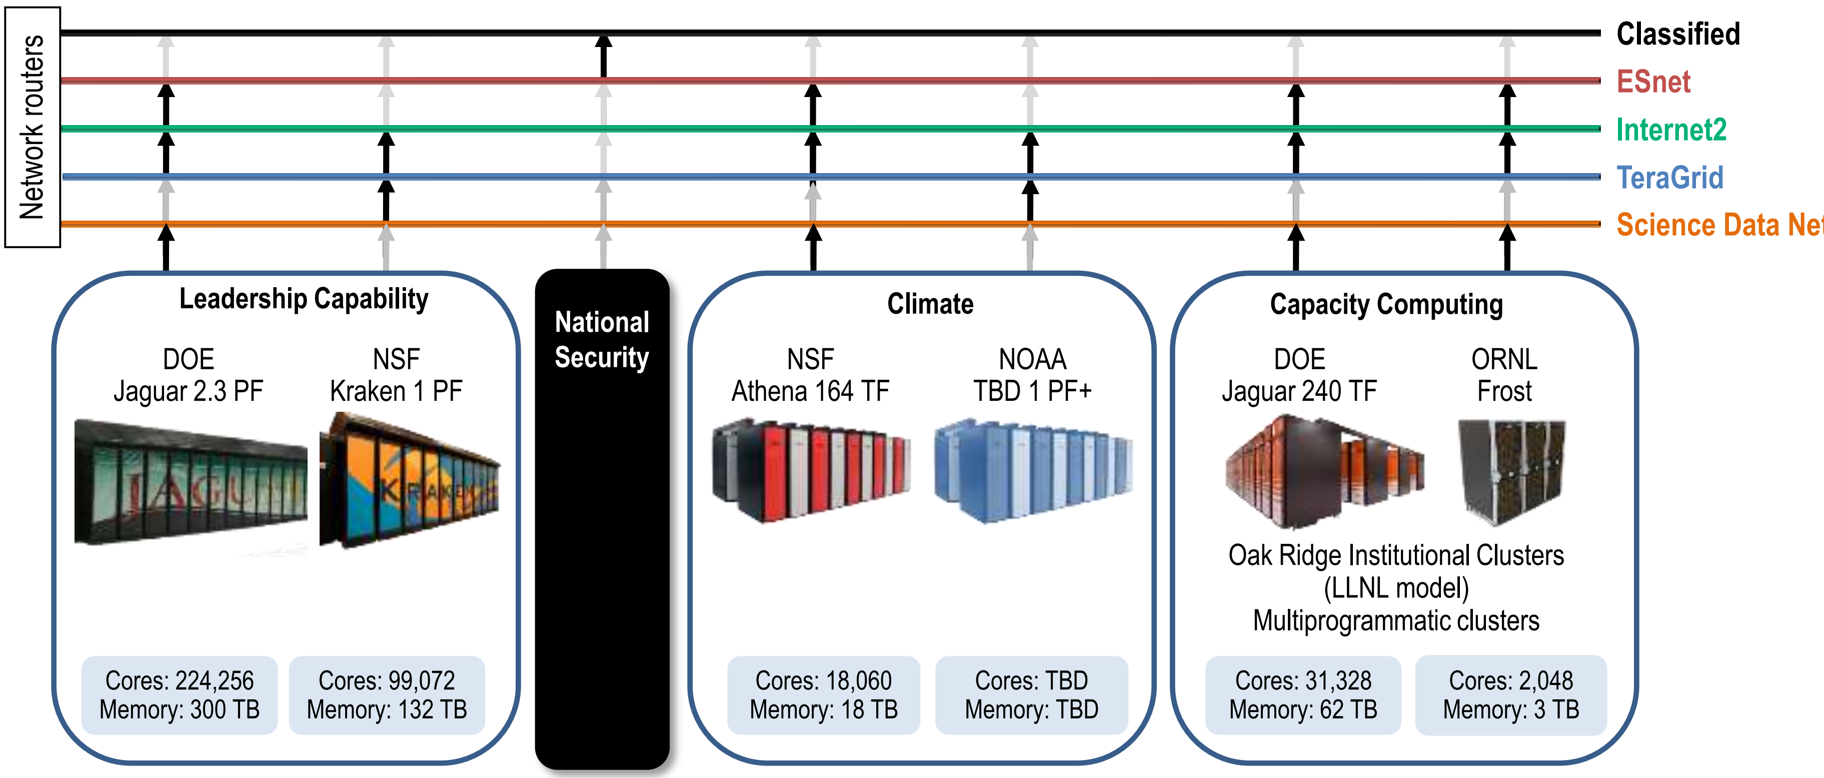

Scientific Visualization Lab

- EVEREST:
  30 ft × 8 ft
- 35 million pixels
- EVEREST cluster

Data analysis

- LENS: 32 nodes
- EWOK: 81 nodes

Experimental systems

- Keeneland, NSF: GPU- based compute cluster
- IBM BG/P
- Power 7
- Cray XMT

Center-wide file system

- Spider
  - 192 data servers
  - 10 PB disks
  - 240 GB/sec

Archival storage

- HPSS
  - Stored data: 8 PB
  - Capacity: 30 PB

December 2009: Summary

- **Supercomputers: 6**
  - **Cores: 376,812**
  - **Memory: 516 TB**
  - **Petaflops: 3.85**
  - **Disks: 14,350 TB**

ORNL’s data center: Designed for efficiency

13,800 volt power into the building saves on transmission losses

480 volt power to computers saves $1M in installation costs and reduce losses

Liquid Cooling is 1,000 times more efficient than air cooling

Result: With a PUE of 1.25, ORNL has one of the world’s most efficient data centers

Variable Speed Chillers save energy

Vapor barriers and positive air pressure keep humidity out of computer center

Flywheel based UPS for highest efficiency

State-of-the-art owned network is directly connected to every major R&E network at multiple lambdas

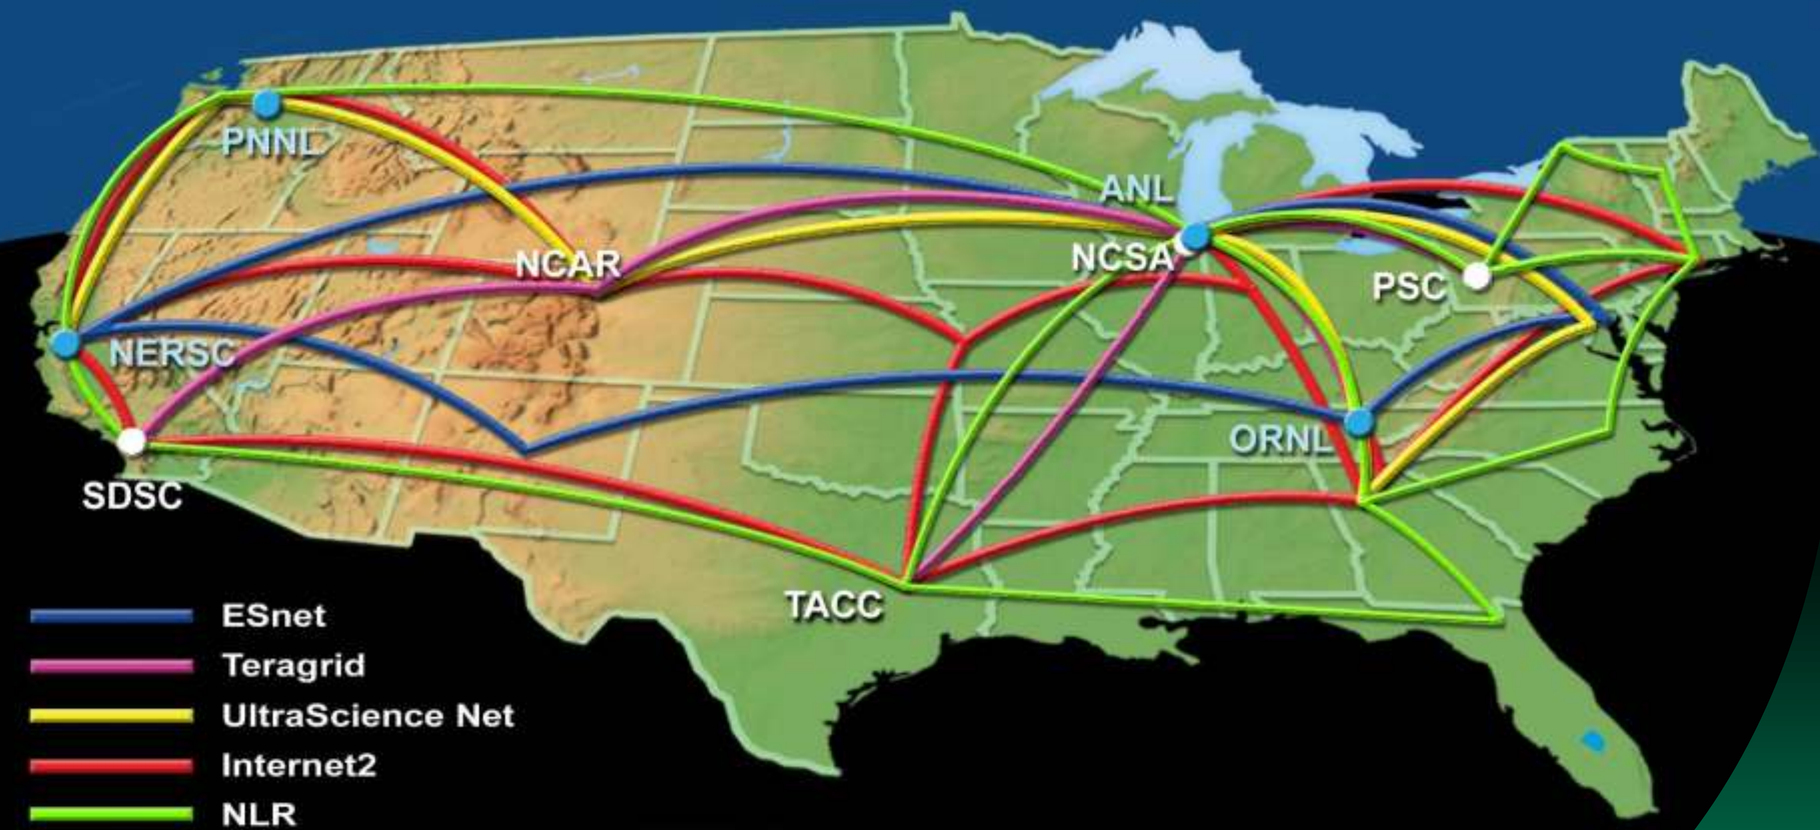

We have increased system performance by 1,000 times since 2004

Hardware scaled from single-core through dual-core to quad-core and dual-socket , 12-core SMP nodes

- NNSA and DoD have funded much of the basic system architecture research
  - Cray XT based on Sandia Red Storm
  - IBM BG designed with Livermore
  - Cray X1 designed in collaboration with DoD

Scaling applications and system software is biggest challenge

- DOE SciDAC and NSF PetaApps programs are funding scalable application work, advancing many apps
- DOE-SC and NSF have funded much of the library and applied math as well as tools
- Computational Liaisons key to using deployed systems

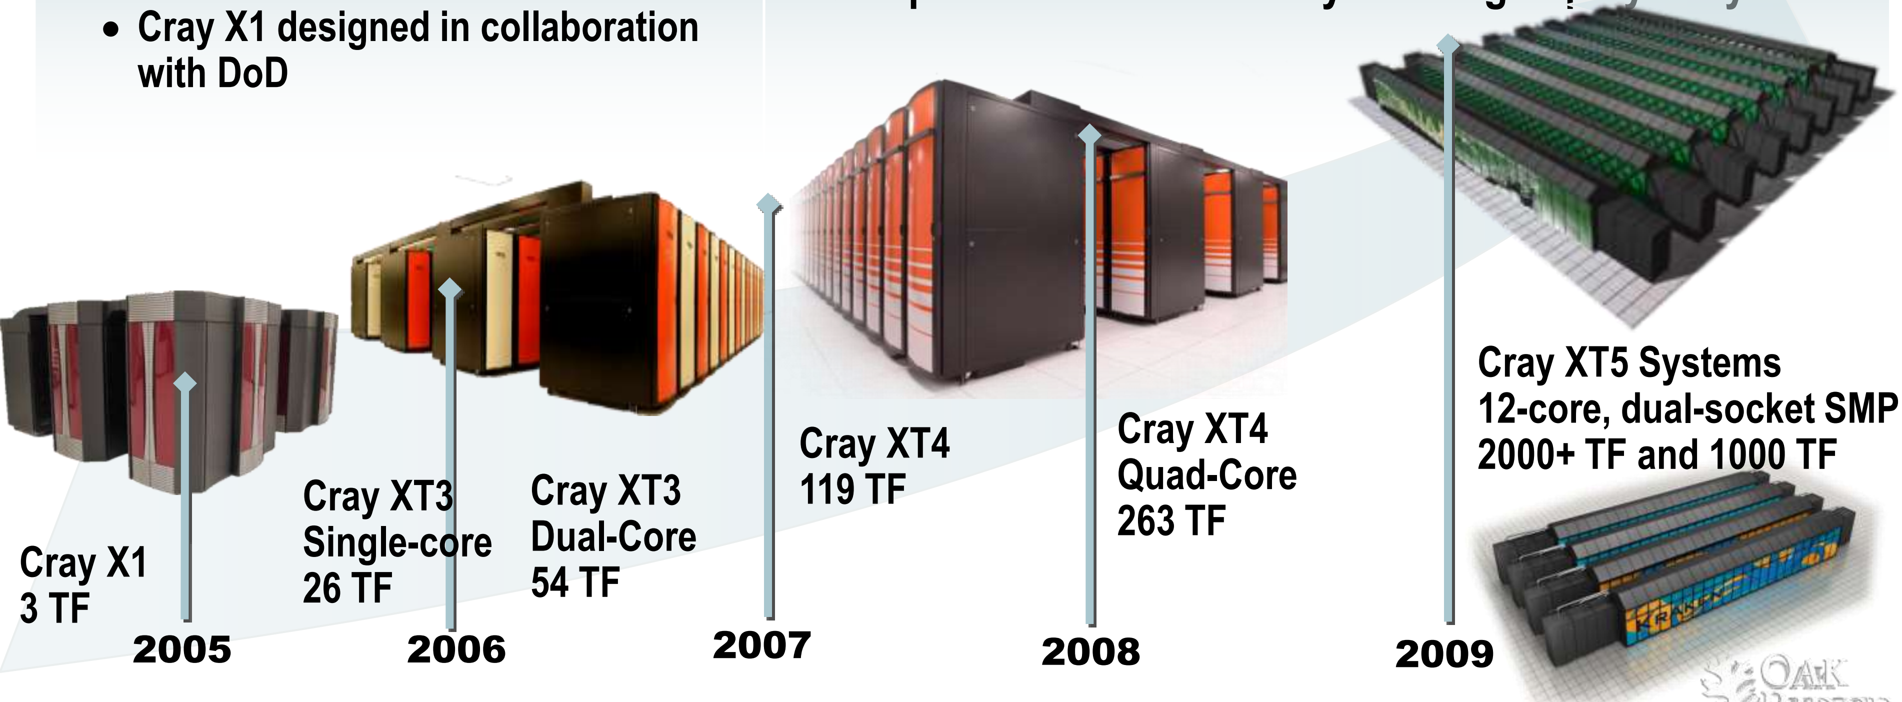

Science requires advanced computational capability 1000x over the next decade

Mission: Deploy and operate the computational resources required to tackle global challenges	Vision: Maximize scientific productivity and progress on the largest scale computational problems
• Deliver transforming discoveries in climate, materials, biology, energy technologies, etc. • Ability to investigate otherwise inaccessible systems, from regional climate impacts to energy grid dynamics	• Providing world-class computational resources and specialized services for the most computationally intensive problems • Providing stable hardware/software path of increasing scale to maximize productive applications development

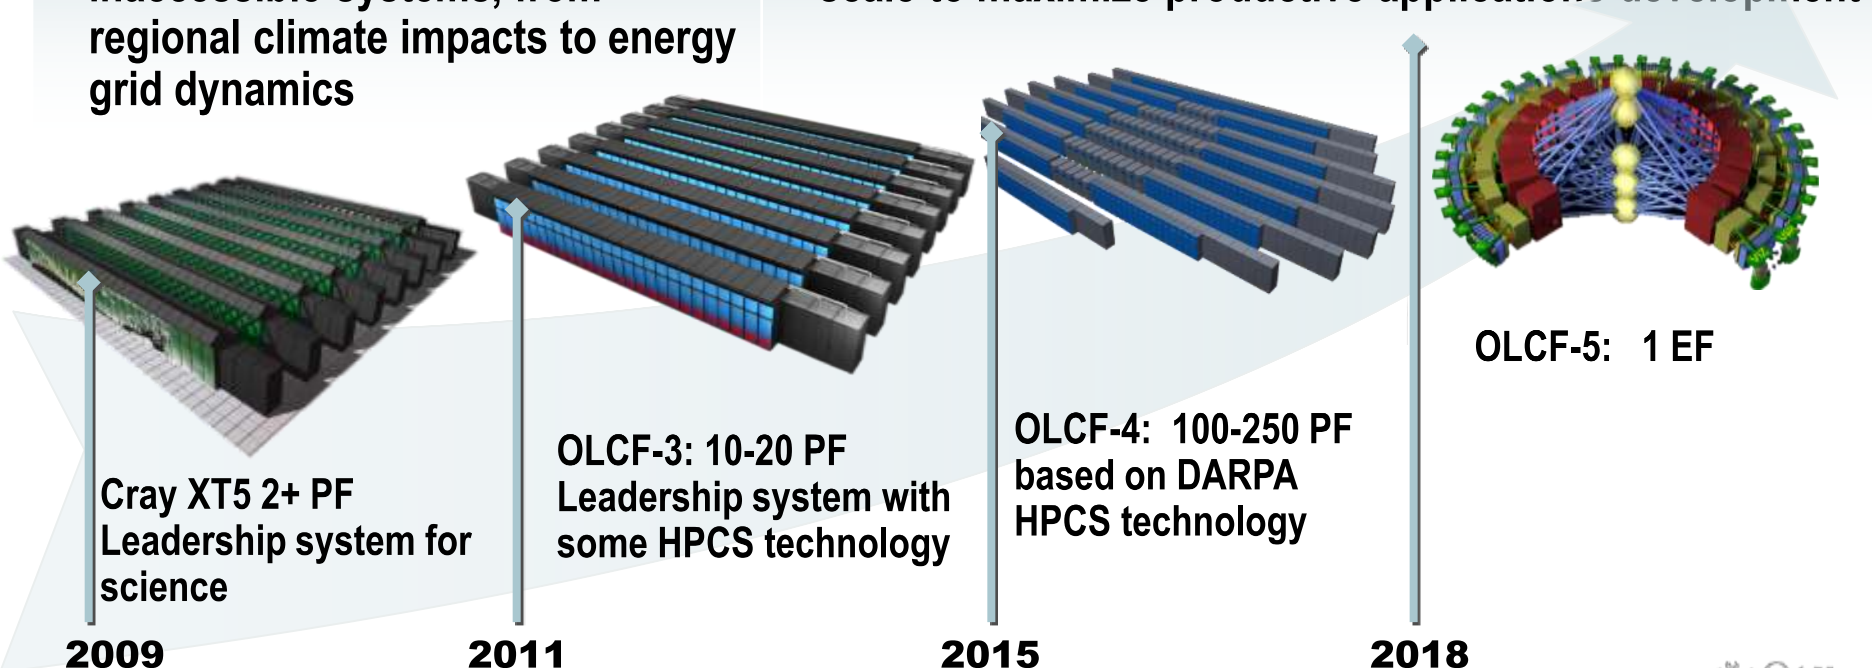

OAK RIDGE National Laboratory

What Will an EF System Look Like?

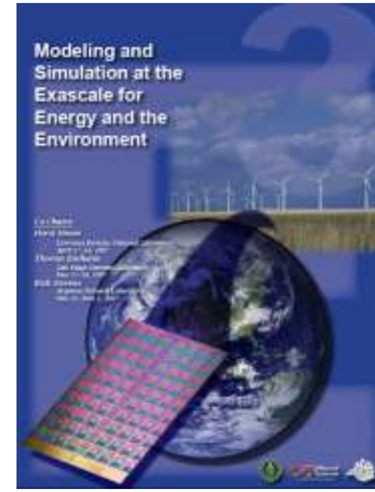

- All projections are daunting
  - Based on projections of existing technology both with and without “disruptive technologies”
  - Assumed to arrive in 2016-2020 timeframe
- Example 1
  - 400 cabinets, 115K nodes @ 10 TF per node, 50-100 PB, optical interconnect, 150-200 GB/s injection B/W per node, 50 MW
- Examples 2-4 (DOE “Townhall” report*)

Example system	Ops/ cycle	Freq [GHz]	Cores/ socket	Peak/ socket [TF/s]	Sockets	Total cores	Peak/ system [EF/s]	Power [MW]
A	4	3.0	64	0.768	1300k	85M	1.0	130
B	8	16.0	128	16.0	120k	15M	2.0	60 - 80
C	8	1.5	512	6.1	200k	100M	1.8	20 - 40

*www.er.doe.gov/ASCR/ProgramDocuments/TownHall.pdf

We have always had inflection points where technology changed

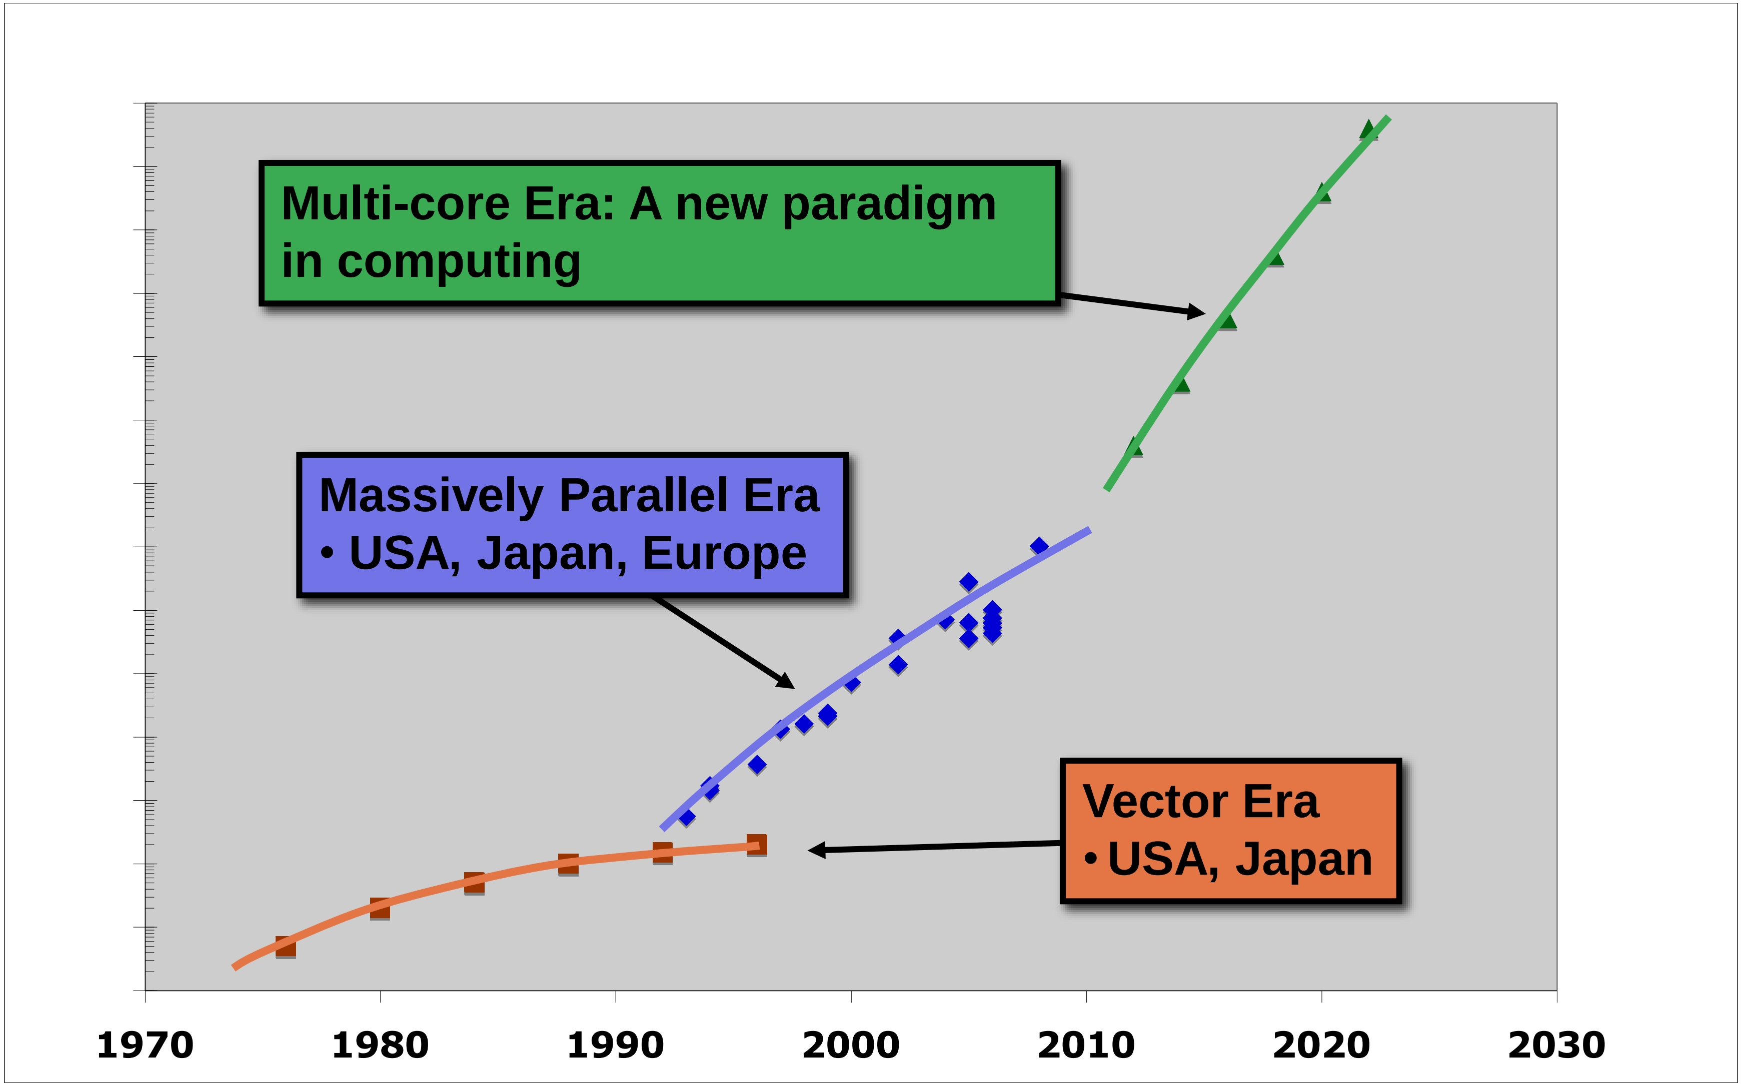

We have a unique opportunity for advancing math and computer science critical to mission success through multi-agency partnership

- Two national centers of excellence in HPC architecture and software established in 2008
  - Funded by DOE and DOD
  - Major breakthrough in recognition of our capabilities

Institute for Advanced Architectures and Algorithms

- **Jointly funded by NNSA and SC in 2008 ~$7.4M**
- **ORNL-Sandia partnership**

Extreme Scale System Center

- **$7M in 2008**
- **Aligned with DOE-SC interests**

IAA is the medium through which architectures and applications can be co-designed in order to create synergy in their respective evolutions.

Computing and Computational Sciences Directorate

Associate Laboratory Director
Jeff Nichols
Debbie McCoy, Assistant to the ALD
Linda Malone, Executive Secretary

Strategic Programs
Gil Weigand

Directorate Operations
Becky Verastegui, Director

Industrial Partnerships
Suzy Tichenor

Office of the Chief Information Officer
Mike Bartell
Chief Information Officer

Computational Sciences and Engineering Division
Brian Worley
Director

Matrix Support
Kyle Spence
Business Manager
Nancy Wright, Haifa Abbott
HR Manager
Wendy Peper
Recruiting Manager
Willy Besancenez
Procurement Manager

Computer Science and Mathematics Division
Barney Maccabe
Director

Information Technology Services Division
Becky Verastegui
Director

National Center for Computational Sciences
Jim Hack
Director

Leadership Computing Facility
Buddy Bland
Project Director

National Institute for Computational Sciences
Phil Andrews
Project Director

Joint Institute for Computational Sciences
Jeff Nichols
Director

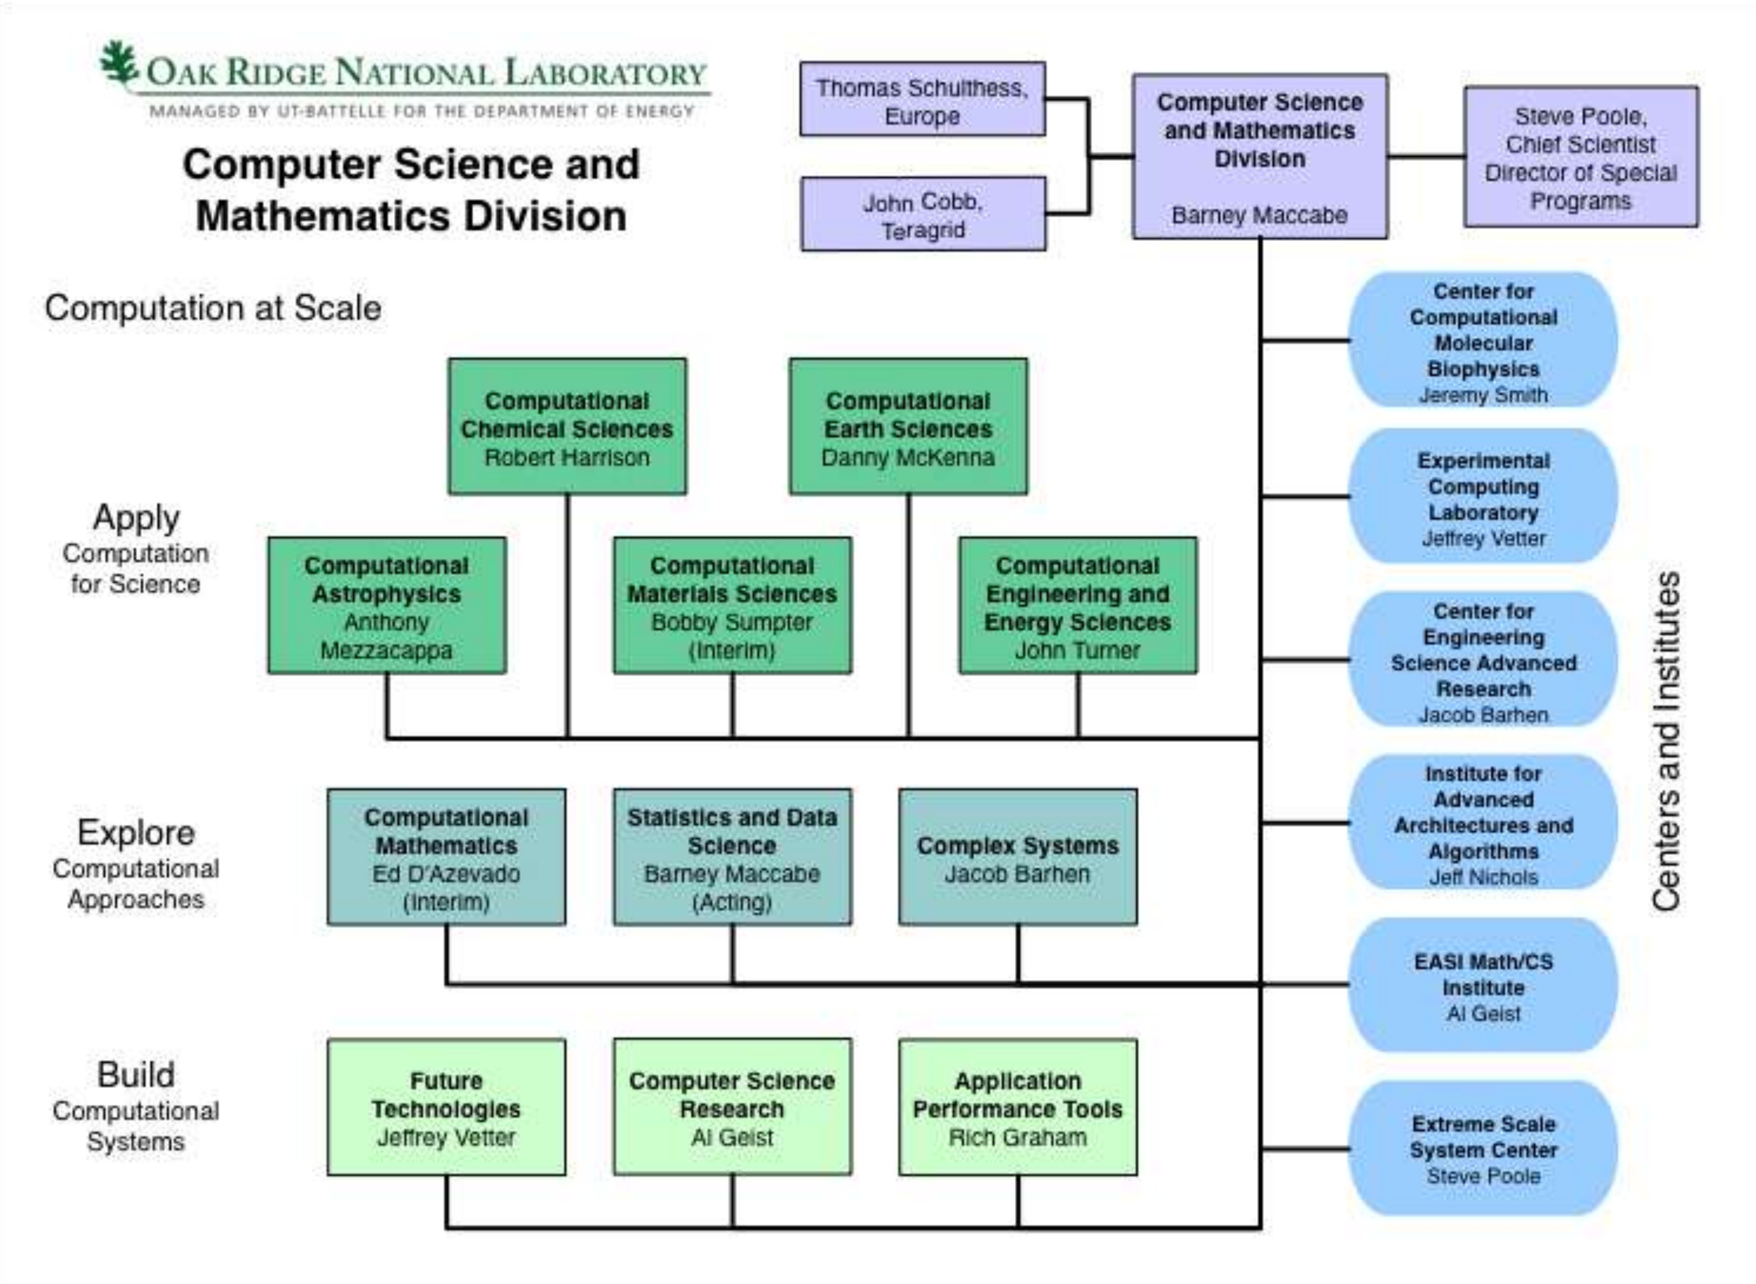

OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY
Computer Science and Mathematics Division
Thomas Schulthess, Europe
John Cobb, Teragrid
Computer Science and Mathematics Division
Barney Maccabe
Steve Poole, Chief Scientist Director of Special Programs
Computation at Scale
Apply
Computation for Science
Computational Chemical Sciences
Robert Harrison
Computational Earth Sciences
Danny McKenna
Computational Astrophysics
Anthony Mezzacappa
Computational Materials Sciences
Bobby Sumpter (Interim)
Computational Engineering and Energy Sciences
John Turner
Explore
Computational Approaches
Computational Mathematics
Ed D'Azevado (Interim)
Statistics and Data Science
Barney Maccabe (Acting)
Complex Systems
Jacob Barhen
Build
Computational Systems
Future Technologies
Jeffrey Vetter
Computer Science Research
Al Geist
Application Performance Tools
Rich Graham
Center for Computational Molecular Biophysics
Jeremy Smith
Experimental Computing Laboratory
Jeffrey Vetter
Center for Engineering Science Advanced Research
Jacob Barhen
Institute for Advanced Architectures and Algorithms
Jeff Nichols
EASI Math/CS Institute
Al Geist
Extreme Scale System Center
Steve Poole
Centers and Institutes

Computer Science Challenges and Example Projects

- Architectures and Interconnection Networks
- Programming Models
- Performance Modeling, Measurement, and Engineering
- System Software and **Runtime Environments**
- Translation and Performance Tuning
- **Software Development**
- I/O and Storage
- Fault Tolerance
- **Connections to Commodity**

Vancouver: A Software Stack for Productive Heterogeneous Exascale Computing

Jeffrey Vetter, ORNL
Wen-Mei Hwu, UIUC
Allen Malony, University of Oregon
Rich Vuduc, Georgia Tech

Objectives

- Enhance programmer productivity for the exascale
  - Increase code development ROI by enhancing code portability
  - Decrease barriers to entry with new programming models
- Create next-generation tools to understand the performance behavior of an exascale machine

Approach

- Programming tools
  - GAS programming model
  - Analysis, inspection, transformation
- Software libraries: autotuning
- Runtime systems: scheduling
- Performance tools
- Impact on DOE Applications

The proposed Maestro runtime simplifies programming heterogeneous systems by unifying OpenCL task queues into a single high-level queue.

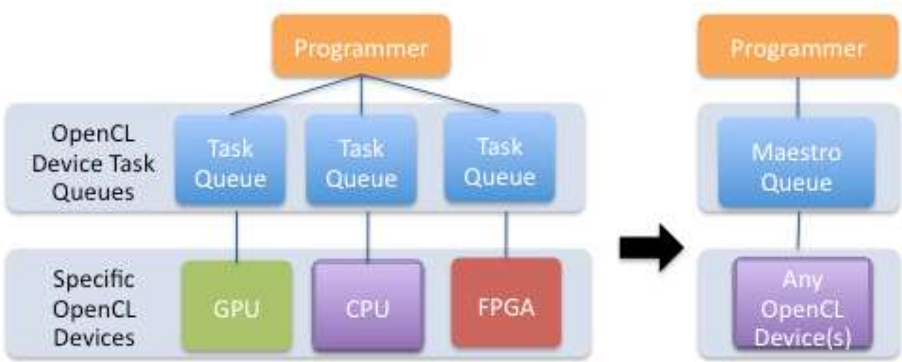

Impact

- Reduced application development time
- Ease of porting applications to heterogeneous systems
- Increased utilization of hardware resources and code portability

 CSMD Georgia Tech College of Computing Computational Science and Engineering ILLINOIS UNIVERSITY OF OREGON OAK RIDGE National Laboratory

Vancouver: A Software Stack for Productive Heterogeneous Exascale Computing

Jeffrey Vetter, ORNL
Wen-Mei Hwu, UIUC
Allen Malony, University of Oregon
Rich Vuduc, Georgia Tech

Challenges

- The coming generation of exascale systems will require massive parallelism, leading to challenges in performance tools and runtime system development
- These systems are also likely to include heterogeneity, for which development tools, programming models, systems software, and scientific support libraries are immature
- Proposed software must be flexible to accommodate changes in rapidly evolving hardware

Formal Problem Statement

- The new trajectory of heterogeneous exascale systems comes with multiple challenges, including poor programmer productivity, application portability, lack of standard tools and libraries, and very sensitive runtime efficiency
- Few languages can span the wide range of concurrency and granularity. This is also true for performance tools, debuggers, resource managers, and libraries

Research Products/Artifacts

- Develop libraries and infrastructure for benchmarking, auto-tuning, data movement, and task scheduling
- Design new performance tools for code analysis, inspection, and transformation
- Research new programming models that support globally addressable data and hierarchical parallelism

Example results from the SHOC benchmark suite show comparative performance across two scientific kernels for multi-core CPUs and GPUs.

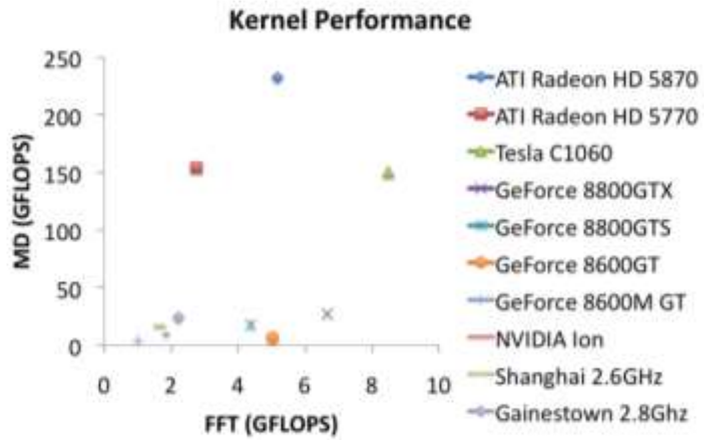

CSMD Georgia Tech College of Computing Computational Science and Engineering ILLINOIS UNIVERSITY OF OREGON OAK RIDGE National Laboratory

COMPOSE-HPC: Software Composition for Extreme Scale Computational Science and Engineering

David E. Bernholdt, ORNL
Tom Epperly, LLNL
Manoj Krishnan, PNNL
Matt Sottile, Galois
Rob Armstrong, SNL

Objectives

- Develop a flexible, extensible toolkit to help software developers address various kinds of software composition challenges
- Provide examples of how the toolkit can be applied to specific composition problems
- Develop new approaches to facilitate composition of parallelism (threads and processes)

A schematic of the KNOT tool chain

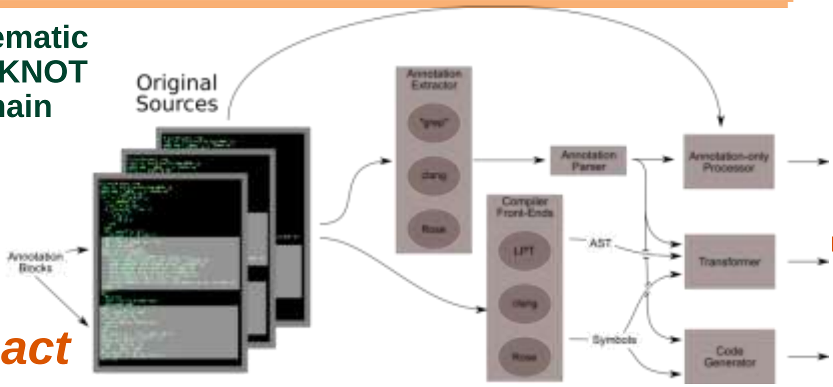

Impact

- Practical tools to help software developers produce higher quality code that works effectively in modern HPC environments
- Bring the capabilities of code transformation and code generation to bear on the challenges of software composition

Approach

- Develop the Knot Nimble Orchestration Toolkit (KNOT), consisting of three components:
  - An annotation parsing facility (PAUL) to interpret guiding annotations embedded as comments in user source code
  - A transformation facility (ROTE) to apply source-to-source transformations to the code, based on annotations and other inputs
  - A code generation facility (BRAID) capable of manipulating compiler-like intermediate representations, transforming, optimizing, and generating source code based on the results
- Develop examples of using KNOT to address different composition problems: language interoperability, contract enforcement, automatic performance instrumentation, data marshalling for GPUs
- Build on this infrastructure to facilitate the composition of parallelism
  - Allow codes to express their preferred threaded execution model, and call modules with different threading models
  - Simplifying the expression and exploitation of MPMD parallelism

CSMD

OAK RIDGE
National Laboratory

COMPOSE-HPC: Software Composition for Extreme Scale Computational Science and Engineering

David E. Bernholdt, ORNL
Tom Epperly, LLNL
Manoj Krishnan, PNNL
Matt Sottile, Galois
Rob Armstrong, SNL

Challenges

- Software composition takes many forms, so tools addressing such issues must carefully balance between flexibility and usability
- Software composition also depends on the non-scientific details of the HPC environment, tools, and system software

Research Products/Artifacts

- A flexible tool chain ("KNOT") allowing user-defined transformations and code generation, providing developers the tools to automate some of their software composition challenges
- Demonstrations of the use of KNOT, including language interoperability (a more flexible "Babel 2.0"), contract enforcement, automatic performance instrumentation, and data marshalling for GPUs
- Tools to express the desired threaded execution models of software modules and compose modules with disparate models
- Tools to express and manage concurrency through an MPMD execution model

Formal Problem Statement

- Composition of software plays a central role in the development of modern applications for computational science and engineering
- Many composition issues can be addressed with minimal intrusion on the user's code through a combination of transformations and code generation
  - Different types of composition requires *different* transformations and code generation
  - We are developing flexible tools to support the creation and application of customized composition capabilities by users
  - We are developing examples of how the tools can be used for different types of composition
- Leveraging these tools, we will also investigate methods to facilitate the composition of parallelism at both the thread and process levels

Evaluating the Role of Cloud Computing for Scientific Discovery

Rob Gillen, ORNL
Sudharshan Vazhkudai, ORNL

Objectives

- Expose Petascale Scientific Data to the Cloud
- Develop Parallel Extensions to the Cloud (APIs)
- Develop Additional Scientific Services for the Cloud
- Establish cloud computing test bed
- Establish a level of expertise in cloud computing at ORNL

Approach

- Partner with commercial cloud vendors to understand issues surrounding (and work at remediation) data movement.
- Deploy representative codes both in native and cloud "ports" to discover areas of benefit as well as future work
- Develop experiment workflows that integrate cloud and traditional fixed assets

Significant Improvements in Transfer Due to Concurrency and Adaptive Compression

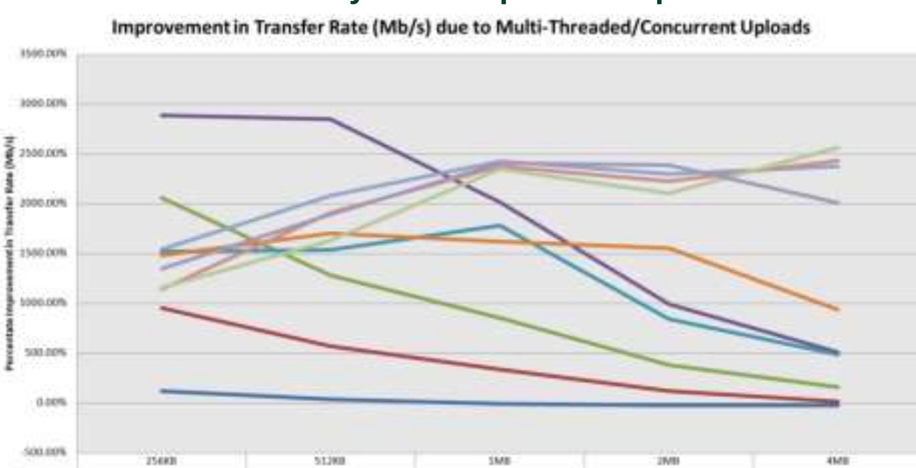

Impact

- While cloud computing is generally focused on horizontally-scaling of Internet properties, the inherent characteristics make it interesting as an on-demand, massively scalable scientific compute platform – particularly for mid-range computing.
- As the primary design target of these platforms is non-technical, work needs to be done to adapt APIs and tools to support scientific workloads

CSMD

OAK RIDGE National Laboratory

Evaluating the Role of Cloud Computing for Scientific Discovery

Rob Gillen, ORNL
Sudharshan Vazhkudai, ORNL

Challenges

- The variability of the public Internet exacerbates the issues of data movement found in traditional datacenters
- Most scientific codes were written assuming high-speed or specialized interconnects which are unavailable in the cloud
- Appropriate use of cloud computing environments requires non-incremental thinking to code design

Formal Problem Statement

- The non-availability of specialized networks in the cloud requires novel approaches to problems (master/worker, queues, map/reduce, etc.).
- The verbosity of "cloud-friendly" data protocols is at odds with the scale of scientific data yet domain-specific formats are at odds with the openness of the cloud.

Research Products/Artifacts

- Libraries demonstrating approaches for improved data movement
- Sample cloud post-processing/data-distribution applications
- Sample ports of scientific codes – foundational to further usages
- Cloud computing test bed for ORNL researchers
- Experiment codes demonstrating mixed-use (cloud + fixed) applications

Cloud-Processed Climate Data Visualization (Post Processing)

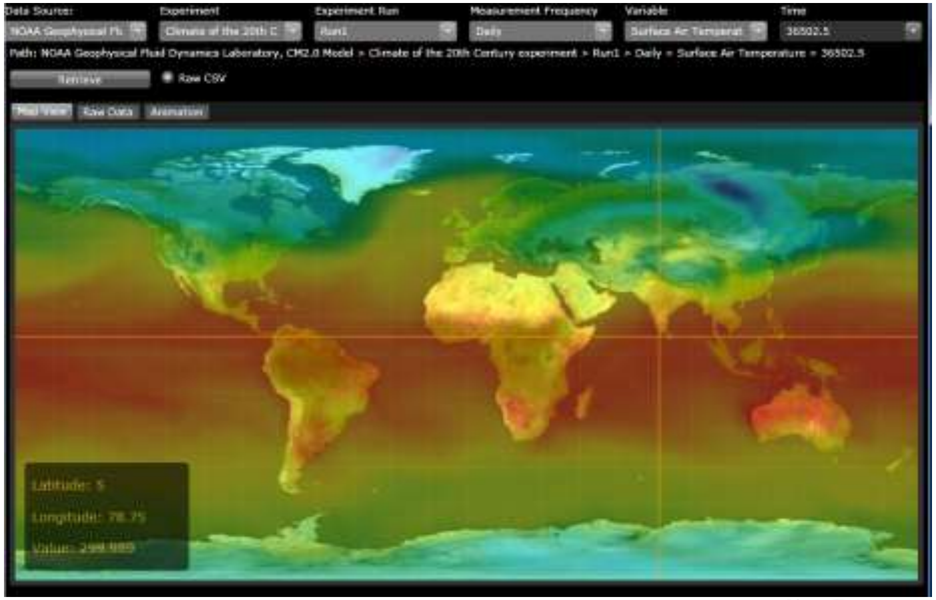

CSMD

OAK RIDGE National Laboratory

Mathematics Challenges and Example Projects

- Libraries and Frameworks
- Fault Oblivious Algorithms
- **Algorithmic models and Structures**
- Asynchronous Algorithms
- Verification and Validation of
- **Uncertainty Quantification**
- Data Analysis

Advanced Dynamically Adaptive Algorithms for Stochastic Simulations on Extreme Scales

Rick Archibald, ORNL
Dongbin Xiu, Purdue University

Objectives

- Extend the existing deterministic adaptive mesh refinement (AMR) algorithms to high dimensional random space to facilitate adaptive stochastic simulations.
- Develop and extend the existing edge-detection methods to stochastic spaces to identity and locate smooth sub-domains.
- Develop advanced, dynamically adaptive, stochastic collocation methods for extreme-scale computing.

Approach

- Minimize cost by ***leading*** computational simulations and running only the essential computational simulations scenarios.
- Automate ***adaptive mesh refinement*** (AMR) strategies in stochastic space of arbitrary dimension
- Build ***advanced dynamically adaptive stochastic collocation*** approaches

Strong (left) and weak (right) scaling data for our block-structured AMR system AMROC.

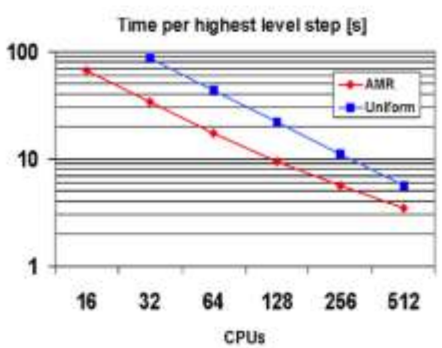

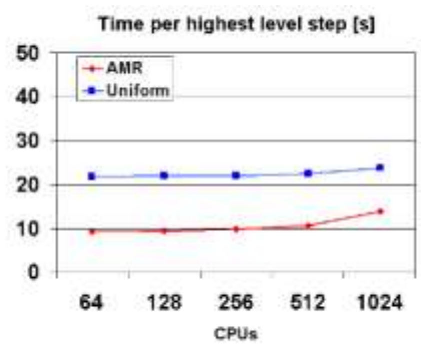

Impact

- This set of advanced numerical algorithms and computational tools for modeling uncertainty in complex stochastic systems will enable application scientists in the high-performance computing (HPC) environment to provide predictive modeling and simulation on the next-generation computer architecture on extremely large scales.
- Essential for predictive modeling and simulation-based decision-making.

Rick Archibald, ORNL
Dongbin Xiu, Purdue University

Advanced Dynamically Adaptive Algorithms for Stochastic Simulations on Extreme Scales

Challenges

- A fundamental challenge in uncertainty quantification is how to construct and evaluate approximations of non-smooth, high dimensional functions with a limited number of simulations one can afford. This is the context of capacity-based simulation methods we focus here.
- We have a strategy for minimizing the number of unknowns needed to approximate uncertainty in computational simulation.

Formal Problem Statement

- In a stochastic simulation, one is typically interested in the relationship between the variables that drive the system (inputs) and the system response (outputs). Complete characterization of this relationship is the ultimate goal.
- Non-intrusive methods that couple to existing and future computational models have a natural parallelism that can take full advantage of extreme scale computing.

Research Products/Artifacts

- The project seek to develop mature algorithms and software suitable for extremely large scale stochastic simulations and make the algorithms and software accessible to the simulation scientists in both national laboratories and academia.
- Develop mathematical methods to provide the framework needed to quantity uncertainty, construct surrogate models, and evaluate high-dimensional integrals

Left: Efficient sampling of stochastic space with unknown discontinuities. Right: Optimal scaling as the stochastic dimension increases.

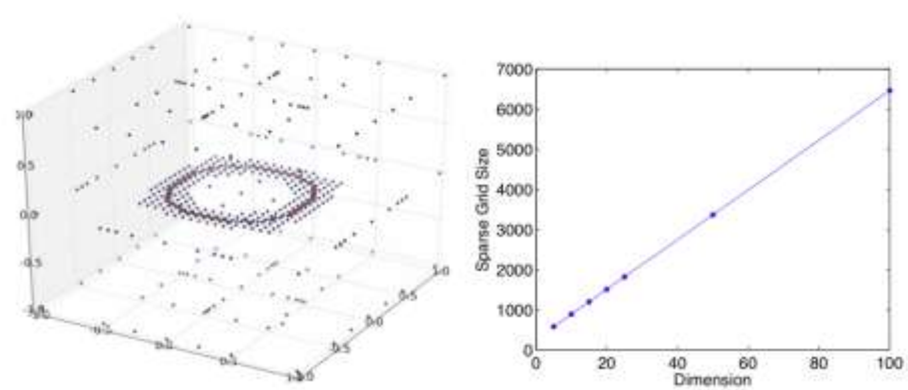

Scalable algorithms for extended magnetohydrodynamics

L. Chacon, D. del-Castillo-Negrete, ORNL
J. Shadid, SNL

Objectives

- Develop next-generation simulation capabilities for extended magnetohydrodynamics via novel, efficient, massively parallel algorithms.
- Enable simulation of realistic conditions in space and fusion plasmas via novel mathematical/algorithmic formulations.

Approach

- Develop novel preconditioning strategies for stiff-wave partial differential equations.
- Develop novel mathematical formulations that enable the accurate numerical treatment of stiff, strongly anisotropic physical processes.

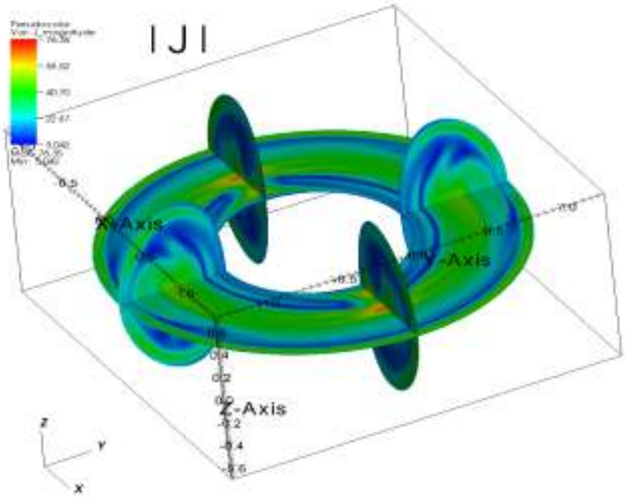

Tokamak simulation in realistic geometry

Impact

- Extended MHD modeling is a cornerstone for many areas of research of interest to DOE, such as magnetic fusion, high-energy-density physics, and the study of solar and magnetospheric phenomena.
- An efficient, scalable extended MHD simulation capability will enable the study of plasmas in previously unaccessible physical regimes.

Scalable algorithms for extended magnetohydrodynamics

L. Chacon, D. del-Castillo-Negrete, ORNL
J. Shadid, SNL

Challenges

- Extended MHD is a very stiff hyperbolic system.
- It also features extreme transport anisotropy due to the presence of the magnetic field.
- These two features challenge the state of the art, both in accuracy (e.g., numerical pollution of anisotropy) and efficiency.

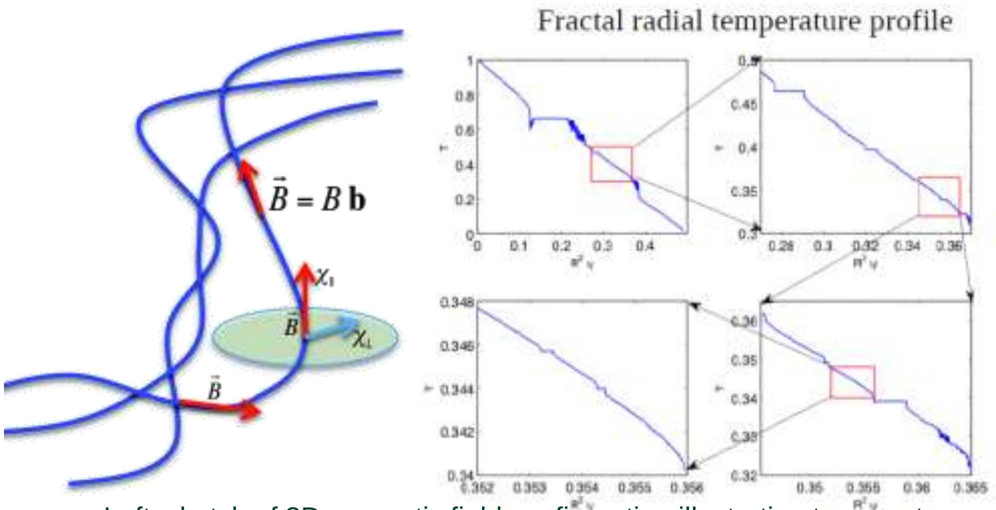

Left: sketch of 3D magnetic field configuration illustrating transport anisotropy. Right: fractal temperature profile in stochastic 3D magnetic field configuration obtained with our algorithms.

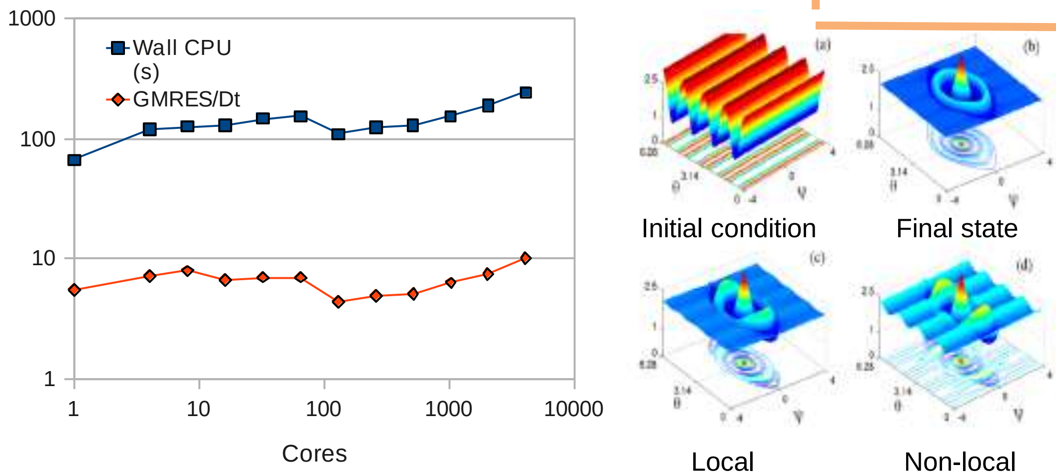

Left: parallel scalability of resistive MHD solver. Right: anisotropic heat transport in presence of magnetic island.

- Demonstrated scalable solver (algorithmic and in parallel) for resistive and extended MHD. Demonstrated scalability up to 4K processors.
- Developed novel mathematical formulation for anisotropic transport in the limit of . Approach has enabled previously out-of-reach simulations of temperature transport in 3D stochastic magnetic fields.

CSMD

Computational Science Challenges and Example Projects

- Application Scalability

- **Astrophysics**
- Chemistry
- Materials Science
- **Earth Science**
- **Engineering and Energy Sciences**

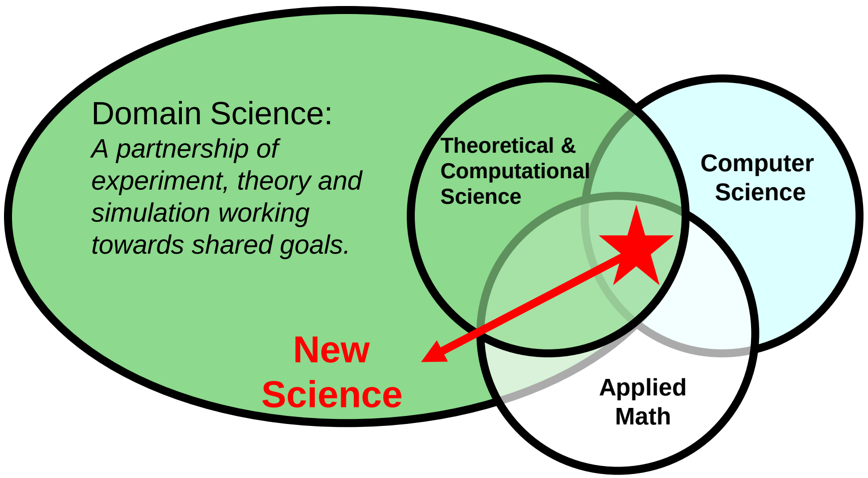

Exascale Supernova Initiative

Anthony Mezzacappa, ORNL
Bronson Messer, ORNL

Objectives

- Ascertain the core collapse supernova explosion mechanism.
- Make detailed predicts of supernova element synthesis, gravitational wave production, and neutrino signatures.
- Use core collapse supernovae as laboratories for fundamental neutrino and nuclear physics.
- Develop simulation capabilities for general relativistic astrophysical phenomena.
- Develop exa-scalable algorithms and programming models for general relativistic radiation magnetohydrodynamics with AMR.

Approach

- Perform a series of staged simulations, staged in spatial dimensionality (e.g., 2D vs. 3D) and physics (e.g., multifrequency neutrino transport vs. multifrequency and multiangle neutrino transport), over the next decade.
- Develop exa-scalable Newton-Krylov approaches for transport solution.
- Continue with space-filling-curve approach to AMR while looking at task-based approaches.
- Continue with MPI+OpenMP while looking at CAF in the intermediate term and Chapel in the longer term.

Convection developing below the supernova shock wave in the first "4D" simulation performed to date, with three spatial dimensions and multifrequency neutrino transport.

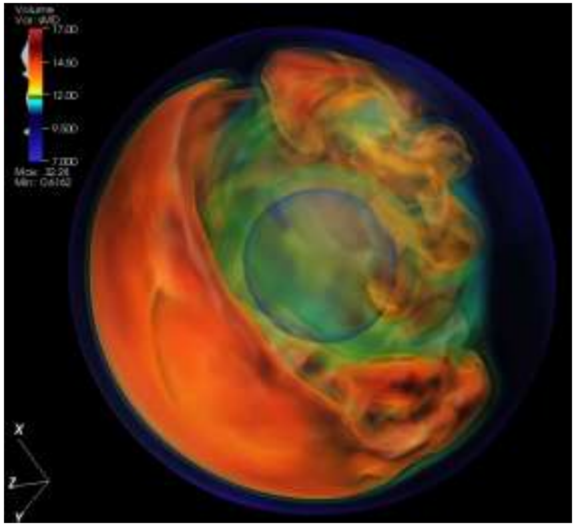

Impact

- Core collapse supernovae are the dominant source of elements in the Universe necessary for life. They are also laboratories for fundamental physics not accessible in terrestrial experiments.
- First series of 2D models with multifrequency neutrino transport to obtain explosions. Explosions obtained over the broadest range of stellar masses to date.
- First complete gravitational wave signatures computed in the context of 2D models.
- First realistic 3D multiphysics simulation initiated.
- 3D MHD studies demonstrated that the SASI may be the origin of neutron star magnetic fields.

Exascale Supernova Initiative

Anthony Mezzacappa, ORNL
Bronson Messer, ORNL

Challenges

- Core collapse supernovae are multiphysics phenomena involving general relativistic radiation magnetohydrodynamics (GRMHD) with multifrequency and multiangle neutrino transport. Thus, they are multidimensional (7D), multiphysics phenomena.
- Simulations of core collapse supernovae require exa-scalable approaches to adaptive mesh refinement and the solution of the general relativistic radiation magnetohydrodynamics equations.
- Analysis of multidimensional, multivariate supernova simulation data requires effective strategies for data management and new approaches to scientific visualization.

Formal Problem Statement

- Progress toward ascertaining the core collapse supernova mechanism will entail the development of a GRMHD capability with, at a minimum, multifrequency neutrino transport and, ultimately, multifrequency and multiangle Boltzmann neutrino transport and quantum kinetics.
- Simulations will be enabled through the development of scalable Newton-Krylov solution algorithms for the linear systems underpinning the solution of the neutrino transport equations.
- Successful simulations of core collapse supernovae will be founded on successful approaches (e.g., task-based) to AMR.

Research Products/Artifacts

- In our effort to solve the supernova problem and develop a predictive capability for key supernova observables we are developing two multiphysics codes: CHIMERA and GenASiS. While these codes will be used primarily on the supernova problem, they can be applied to a broader class of general relativistic astrophysical problems.
- Algorithms for the solution of radiation transport problems on massively parallel architectures.
- Effective approaches to scientific workflows, from data management strategies to scientific visualization, are being developed. These approaches will benefit other application areas.

Snapshot of entropy (equivalently temperature) below the supernova shock wave (surface) during an explosion.

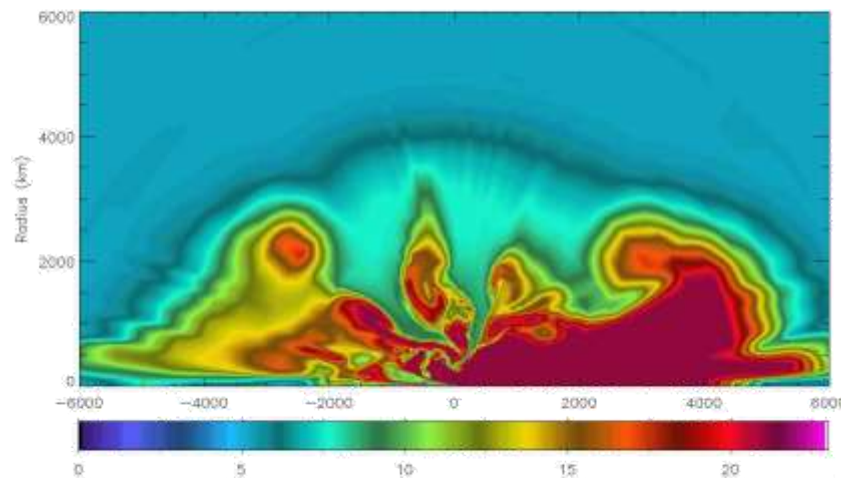

Ultra High Resolution Global Climate Simulation to Explore and Quantify Predictive Skill

Richard Archibald, Kate Evans,
Michael Ham, Pat Worley
Jim Hack (PI) - ORNL

Objectives

- Test the hypothesis that a higher resolution CCSM is necessary
- Include explicit simulation of non-linear phenomena and interactions on the small scale that have feedbacks on large scale climate features
- Provide accurate and explicit simulations of local to regional scale phenomena, including low-probability, high-impact hydrological events

1/3 degree resolution global spectral atmospheric model with the CCSM

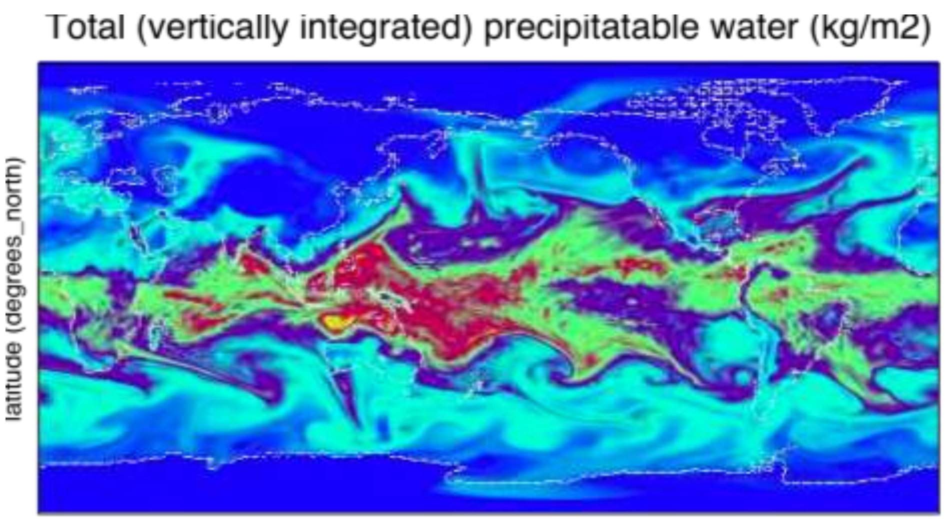

Approach

- For some regions there has been only very limited study of key aspects of regional climate change, particularly with regard to extreme events
- Atmosphere-Ocean General Circulation Models show no consistency in simulated regional precipitation change in some key regions
- In regions where fine spatial climate regimes are affected by topography, there is insufficient information about climate change

Impact

- High resolution global model capability to perform for a number of climate studies
- Better representations of mesoscale storm structure and intensity
- Higher resolution atmospheric models needed to correctly simulate the climatology of monsoon season rainfall events over Asia.

Ultra High Resolution Global Climate Simulation to Explore and Quantify Predictive Skill

Richard Archibald, Kate Evans,
Michael Ham, Pat Worley
Jim Hack (PI) - ORNL

Challenges

- We still need to treat subgrid scale processes (see panel below right)
- Creation of high accuracy high resolution remapping and mask file creation for accurate interpolation between components
- Communication and I/O fragility and overhead
- Scalability across all components
- Keeping up with CCSM4 Development activities

Formal Problem Statement

- Simulation biases in earlier investigations point to need to address several important details about the high-resolution configuration
- Experimental protocol to be performed within CSMD includes (1) an experimental suite designed to test hypotheses, (2) exploration of 3 atmospheric dynamical cores and (3) a systematic plan for simulation evaluation

Research Products/Artifacts

- New configurations within 3 dynamical cores at a variety of resolutions within the CCSM4/5
- Knowledge of when resolution refinement is necessary and perhaps when it is not
- Climate simulation data covering. preindustrial. present day and near term predictions

Illustration of the scale of clouds versus traditional global model resolution (T42), courtesy NASA Goddard Space Flight Center Visualization Studio

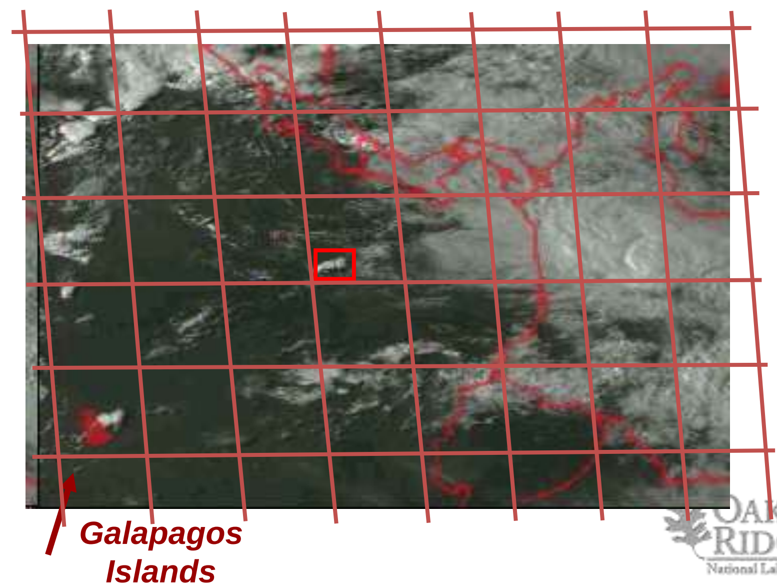

Mesoscale Simulations of Microstructure Evolution in Nuclear Fuels

B. Radhakrishnan,
S.B. Gorti (CSMD)

Objectives

- Develop predictive models for the evolution of grains, pores and bubbles in nuclear fuels in the reactor environment
- Develop strategies for incorporating lower length scale physics at the mesoscale as well as for upscaling of mesoscale calculations to provide input for larger length scale simulations of fuel swelling and fission gas release

Approach

- Utilize Potts model with inputs from lower length scale models or experimental results to simulate microstructural evolution
- Utilize high performance computing to simulate large three dimensional domains with high spatial resolution

Integrated Model for Fuel Performance

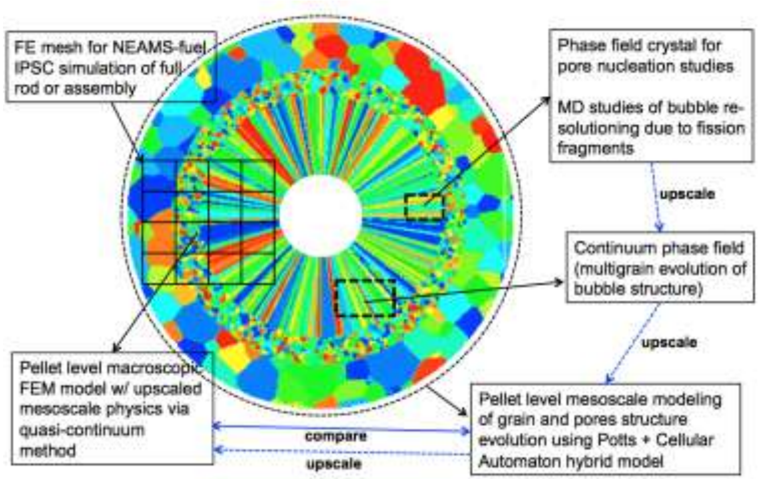

Impact

- Simulations will provide a predictive, physics based capability to model fuel swelling and fission gas release
- Mesoscale simulations will form an integral part of the next generation fuel-performance codes to replace currently used empirical models

Mesoscale Simulations of Microstructure Evolution in Nuclear Fuels

B. Radhakrishnan,
S.B. Gorti (CSMD)

Challenges

- Ability to capture the material-specific lower length scale physics at the mesoscale
- Ability to incorporate irradiation effects on the source and sink terms for bubbles at the mesoscale
- Ability to develop upscaling methods that can be linked efficiently with continuum models

Formal Problem Statement

- Utilize a multi-scale approach to develop predictive, material-specific, mesoscale simulations of grain, pore and bubble structure evolution in oxide fuels
- Develop functional forms for pore, bubble and grain size distribution to be incorporated into continuum models of fuel swelling and fission gas release in oxide fuels

Research Products/Artifacts

- A scalable, parallel Potts model code that is capable of evolving fuel microstructure in three-dimensions with high spatial resolution
- Investigation of pore migration and pore-grain boundary interaction in a temperature gradient

Simulations of pore-grain boundary interaction in a temperature gradient

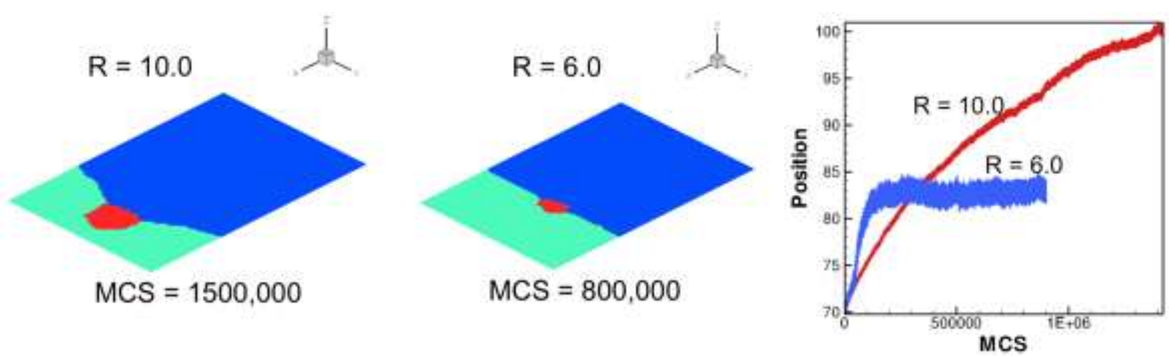

Mesoscale Simulations of Microstructure Evolution in Nuclear Fuels

B. Radhakrishnan,
S.B. Gorti (CSMD)

Progress

- Developed a parallel Potts model code for simulating coupled evolution of grain, pore and bubble structure in the presence of a temperature gradient
- Investigated the migration velocity of a pore in the presence of a temperature gradient for migration by surface or volume diffusion mechanism
- Simulations showed a linear variation of pore velocity with 1/R, where R is the bubble radius, in agreement with analytical results
- Simulations realistically capture the interaction of pores with grain boundaries

Accomplishments

- Mesoscale simulation capability for fuel microstructure evolution
- Validation of pore migration kinetics using analytical results
- Capturing the formation of the central hole in a cylindrical fuel pin

Migration of pore by surface diffusion in the presence of a temperature gradient

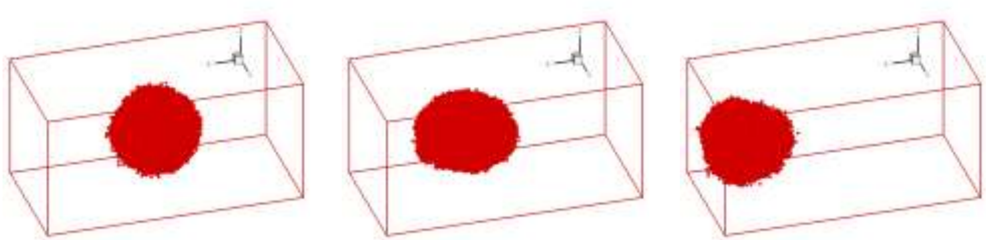

MCS = 10 **MCS = 100,000** **MCS = 200,000**

Pore migration leading to formation of central hole in a cylindrical fuel pin

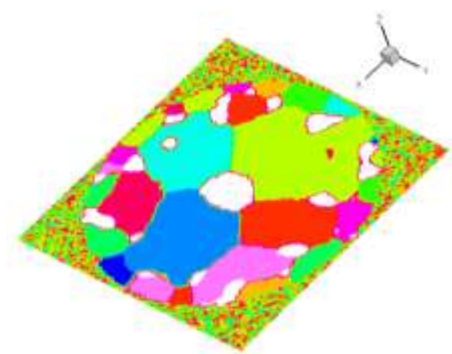

Our strengths in key areas will ensure success

SCIENCE

Exaflops

Applications

- Multidisciplinary application development teams
- Partnerships to drive application performance
- Science base and thought leadership

People

- Exceptional expertise and experience
- In-depth applications expertise in house
- Strong partners
- Proven management team

Software

- Broad system software development partnerships
- Experienced performance/optimization tool development teams
- Partnerships with vendors and agencies to lead the way
- Leverage DOD and NSF investments

Petaflops

Systems

- Driving architectural innovation needed for exascale
- Superior global and injection bandwidths
- Purpose built for scientific computing
- Leverages DARPA HPCS technologies

Facilities

- Power (reliability, availability, cost)
- Space (current and growth path)
- Global network access capable of 96 X 100 Gb/s

Oak Ridge National Laboratory

Oak Ridge National Laboratory:
Meeting the challenges of the 21st century

www.ornl.gov