

Science Prospects and Benefits with Exascale Computing

December 2007

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



DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via the U.S. Department of Energy (DOE) Information Bridge.

Web site <http://www.osti.gov/bridge>

Reports produced before January 1, 1996, may be purchased by members of the public from the following source.

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone 703-605-6000 (1-800-553-6847)
TDD 703-487-4639
Fax 703-605-6900
E-mail info@ntis.fedworld.gov
Web site <http://www.ntis.gov/support/ordernowabout.htm>

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange (ETDE) representatives, and International Nuclear Information System (INIS) representatives from the following source.

Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831
Telephone 865-576-8401
Fax 865-576-5728
E-mail reports@osti.gov
Web site <http://www.osti.gov/contact.html>

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

National Center for Computational Sciences

**SCIENCE PROSPECTS AND BENEFITS
WITH EXASCALE COMPUTING**

Douglas B. Kothe
Director of Science

Date Published: December 2007

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee 37831-6283
managed by
UT-BATTELLE, LLC
for the
U.S. DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

	Page
LIST OF FIGURES	v
LIST OF TABLES	vii
PREFACE	ix
ACKNOWLEDGEMENTS	ix
EXECUTIVE SUMMARY	xi
1. LEADERSHIP COMPUTING: BUILDING FROM THE LAST REVOLUTION	1
2. MATERIALS SCIENCE	2
2.1 RECENT ACCOMPLISHMENTS WITH LEADERSHIP COMPUTING	2
2.2 NEAR-TERM (PETASCALE) SCIENCE DRIVERS	3
2.3 LONGER-TERM (SUSTAINED PETASCALE AND EXASCALE) SCIENCE DRIVERS	4
3. EARTH SCIENCE	4
3.1 RECENT ACCOMPLISHMENTS WITH LEADERSHIP COMPUTING	4
3.2 NEAR-TERM (PETASCALE) SCIENCE DRIVERS	7
3.3 LONGER-TERM (SUSTAINED PETASCALE AND EXASCALE) SCIENCE DRIVERS	9
4. ENERGY ASSURANCE	10
4.1 RECENT ACCOMPLISHMENTS WITH LEADERSHIP COMPUTING	11
4.2 NEAR-TERM (PETASCALE) SCIENCE DRIVERS	13
4.3 LONGER-TERM (SUSTAINED PETASCALE AND EXASCALE) SCIENCE DRIVERS	13
5. FUNDAMENTAL SCIENCE	18
5.1 RECENT ACCOMPLISHMENTS WITH LEADERSHIP COMPUTING	18
5.2 NEAR-TERM (PETASCALE) SCIENCE DRIVERS	19
5.3 LONGER-TERM (SUSTAINED PETASCALE AND EXASCALE) SCIENCE DRIVERS	20
6. THE ROAD AHEAD	21
7. REFERENCES	23

LIST OF FIGURES

Figure		Page
1	The Jaguar leadership system at ORNL.....	1
2	Crystal structure of the high-temperature cuprate superconductor $\text{YBa}_2\text{Cu}_3\text{O}_7$	2
3	Instantaneous net ecosystem exchange viewed along the terminator	5
4	Getting the weather, climate variability, and climate change impacts right in global climate models over the next decade will require the progressive evolution and coupling of multiple physical models	7
5	Future climate modeling will require much more powerful supercomputers	9
6	U.S. energy flows in 2002	11
7	Direct numerical simulation of flame stabilization in a lifted hydrogen flame	13
8	Recent simulation of a cellulase enzyme attacking a cellulose substrate performed at the LCF	15
9	Many protein and enzyme events of interest to biology and bioenergy research are in the millisecond-to-second timescales, which are orders of magnitude beyond those possible with today's simulations.....	17
10	Rendering of the matter entropy during the nonlinear phase of the standing accretion shock instability in a supernova core	19

LIST OF TABLES

Table		Page
ES.1	The promise of exascale computing	xi
1	Recent materials science accomplishments enabled by leadership computing at ORNL	3
2	Select materials science drivers for leadership computing at the petascale (1 to 3 years)	4
3	Select materials science drivers for leadership computing at the exascale (10 years).....	5
4	Recent Earth science accomplishments enabled by leadership computing	6
5	Select Earth science drivers for leadership computing at the petascale (1 to 3 years).....	8
6	Select Earth science drivers for leadership computing at the exascale (10 years).....	10
7	Recent energy assurance accomplishments enabled by leadership computing	12
8	Select energy assurance drivers for leadership computing at the petascale (1 to 3 years)	14
9	Select energy assurance drivers for leadership computing at the exascale (10 years).....	15
10	Recent fundamental science accomplishments enabled by leadership computing	18
11	Select fundamental science drivers for leadership computing at the petascale (1 to 3 years)	20
12	Select fundamental science drivers for leadership computing at the exascale (10 years)	21
13	Science opportunities and impacts grow with increasingly capable leadership systems.....	22

PREFACE

This discussion represents an initial assessment of the scientific advancements that will become plausible in the era of exascale computing and the profound challenges we face in getting there. We fully expect to refine and update it as we broaden our understanding with the aid of leading researchers and policy makers. Nevertheless, we believe it will contribute to the dialogue that must take place among all those who care about scientific and technological progress in the coming decades.

ACKNOWLEDGEMENTS

A document of this type would be impossible if we were not able to tap into the knowledge and wisdom of researchers from a broad range of institutions and the full gamut of scientific endeavors. At the risk of omitting one or more contributors, we would like to thank the following scientists, who freely shared their time and insights through personal discussion, interviews, and surveys: Pratul Agarwal, Valmor de Almeida, David Dean, John Drake, Tom Evans, Robert Harrison, Anthony Mezzacappa, Thomas Schulthess, Edward Uberbacher, Philip LoCascio, and Patrick Worley of Oak Ridge National Laboratory (ORNL); Jeff Candy of General Atomics; Jacqueline Chen of Sandia National Laboratories; Lei-Quan Lee of Stanford Linear Accelerator Center; Peter Lichtner of Los Alamos National Laboratory; Tommaso Roscilde of Max-Planck Gesellschaft; Benoit Roux of Argonne National Laboratory; and Wei-li Lee of Princeton Plasma Physics Laboratory. Without their insight, contributions, and vision, the case for science at the exascale as outlined in this document would not have had nearly as much impact or meaning.

We also acknowledge Ricky Kendall and the Scientific Computing Group within ORNL's National Center for Computational Sciences. In particular, Bronson Messer and James (Trey) White collected and edited key portions of the information contained in this document. And finally, thanks go to science writer Leo Williams of ORNL, who added a measure of conciseness and clarity of presentation to the final document.

EXECUTIVE SUMMARY

Scientific computation has come into its own as a mature technology in all fields of science. Never before have we been able to accurately anticipate, analyze, and plan for complex events that have not yet occurred—from the operation of a reactor running at 100 million degrees centigrade to the changing climate a century down the road. Combined with the more traditional approaches of theory and experiment, scientific computation provides a profound tool for insight and solution as we look at complex systems containing billions of components. Nevertheless, it cannot yet do all we would like. Much of scientific computation’s potential remains untapped—in areas such as materials science, Earth science, energy assurance, fundamental science, biology and medicine, engineering design, and national security—because the scientific challenges are far too enormous and complex for the computational resources at hand. Many of these challenges are of immediate global importance.

These challenges can be overcome by a revolution in computing that promises real advancement at a greatly accelerated pace. Planned petascale systems (capable of a petaflop, or 10^{15} floating point operations per second) in the next 3 years and exascale systems (capable of an exaflop, or 10^{18} floating point operations per second) in the next decade will provide an unprecedented opportunity to attack these global challenges through modeling and simulation. Exascale computers, with a processing capability similar to that of the human brain, will enable the unraveling of longstanding scientific mysteries and present new opportunities.

Table ES.1 summarizes these scientific opportunities, their key application areas, and the goals and associated benefits that would result from solutions afforded by exascale computing.

Table ES.1. The promise of exascale computing

Opportunity	Key application areas	Goal and benefit
Materials science	Nanoscale science; material lifecycles, response, and failure; and manufacturing	Design, characterize, and manufacture materials, down to the nanoscale, tailored and optimized for specific applications
Earth science	Weather, carbon management, climate change mitigation and adaptation, environment	Understand the complex biogeochemical cycles that underpin global ecosystems and control the sustainability of life on Earth
Energy assurance	Fossil, fusion, combustion, nuclear fuel cycle, chemical catalysis, renewables (wind, solar, hydro), bioenergy, energy efficiency, energy storage and transmission, transportation, buildings	Attain, without costly disruption, the energy required by the United States in guaranteed, economically viable, and environmentally benign ways to satisfy residential, commercial, and transportation requirements
Fundamental science	High-energy physics, nuclear physics, astrophysics, accelerator physics	Decipher and comprehend the core laws governing the universe and unravel its origins
Biology and medicine	Proteomics, drug design, systems biology	Understand connections from individual proteins through whole cells into ecosystems and environments
National security	Disaster management, homeland security, defense systems, public policy	Analyze, design, stress-test, and optimize critical systems such as communications, homeland security, and defense systems; understand and uncover human behavioral systems underlying asymmetric operation environments
Engineering design	Industrial and manufacturing processes	Design, deploy, and operate safe and economical structures, machines, processes, and systems with reduced concept-to-deployment time

The challenges presented in this document all require tightly coupled exascale computational platforms. Their solutions must draw from a variety of disciplines and incorporate a range of scales and physical processes. Mathematically they are governed by large sets of intricately linked equations. Distributed computing will not solve these problems, although it is appropriate for a host of other problems (large-scale data analysis, for example). Instead, the critical scientific challenges discussed here require hardware and software structures that can closely integrate many physical processes, each of which is a daunting computational challenge in itself. Such an integration cannot be made across a large, geographically or logically dispersed collection of processors. Real progress on these problems will wholly depend on the descendants of today's terascale (capable of a teraflop, or 10^{12} floating point operations per second) and tomorrow's petascale platforms.

If mathematics is the language of science, computation is its workhorse. Because high-performance computing addresses complexity, accelerates discovery, and enables progress in fields that could not advance without its aid, it has become science's "killer app"—the tool that drives a community to its use once that community views it as indispensable.

The advances outlined in this document are not inevitable. They will require a national commitment of time, talent, and money. Nevertheless, we cannot afford to fail. We must anticipate, mitigate, and adapt to a changing climate. We must provide a range of new energy sources—from fusion to biofuels to hydrogen—that are both cost-effective and plentiful. And we must expand our understanding of the universe, from the idiosyncrasies of subatomic systems to expanding galaxies. Each of these goals requires computing power orders of magnitude beyond what we have available, but each is well within our reach if we commit to achieving it.

Following a brief retrospective on leadership computing (Chap. 1), this document discusses the prospects for each of these scientific opportunities—in materials science (Chap. 2), Earth science (Chap. 3), energy assurance (Chap. 4), and fundamental science (Chap. 5)—drawing on data from and interviews with leading researchers [1, 2]. This discussion focuses on the following areas:

- science accomplishments to date on leadership computing systems,
- near-term (1- to 3-year) petascale science drivers, and
- longer-term (5- to 10-year) exascale science drivers.

After looking at what will be feasible in the next 5 to 10 years and the major challenges remaining, this document concludes that important, high-impact global challenges have a high likelihood of being tackled and solved with the aid and guidance of exascale computing resources (Chap. 6).

1. LEADERSHIP COMPUTING: BUILDING FROM THE LAST REVOLUTION

Researchers at the turn of the millennium realized they would need far more powerful computing tools to meet the challenges of an increasingly complex world [3–6]. In response to their demands, the Leadership Computing Facility (LCF) was established to provide resources 100 times more powerful than the most advanced systems available at the time. The LCF effort is exemplified by Oak Ridge National Laboratory’s (ORNL’s) Jaguar supercomputer (Fig. 1).

The 2004 LCF proposal summarized the research outlook of more than half a dozen science domains [7]. Nearly two dozen leading researchers and key collaborators united in calling for a national computing center, one that nurtures the kind of research that changes the world and the way we understand it. Some of their comments included the following:

- “The [LCF] system proposed here will allow the community to model the global climate, including a dynamic carbon cycle and improved cloud physics while increasing the atmospheric resolution to a 30 km grid and including eddy resolving ocean models.”
- “Capability computing is required in fusion energy studies to enable researchers to describe the new physics regimes needed to model advanced experiments such as ITER [the multinational ITER reactor]—higher spatial resolutions, higher temperature plasma, longer simulation times and higher model dimensionality.”
- “On the [LCF system], it would be computationally feasible to...achieve an understanding of the essential physics of high- T_c [high-temperature] superconductors.”



Fig. 1. The Jaguar leadership system at ORNL.

The LCF met these demands in two ways. First, it brought in state-of-the art supercomputers, systems that offered previously unseen power for open scientific research. Second, it took a revolutionary approach to the job of running a supercomputing center. The LCF has a staff of scientists dedicated to ensuring that each project takes the fullest advantage of these unique, next-generation systems, and the supercomputers themselves are reserved for projects so ambitious they cannot be accomplished anywhere else.

The promise of the LCF has been realized: Multicentury climate simulations (unheard of just a few years ago) are shedding new insight on climate change and its causes and impacts; unique three-dimensional simulations of tokamak plasma are driving ITER design choices; and researchers have discovered that a purely electronic model successfully describes high-temperature superconductors (see Fig. 2). The LCF looks back and celebrates the scientific achievements made possible through terascale computing. It also looks forward to the next generation of systems that will usher in a new era of science exploration and application.

These and other science achievements and opportunities are described in the following chapters.

2. MATERIALS SCIENCE

2.1 RECENT ACCOMPLISHMENTS WITH LEADERSHIP COMPUTING

Leadership computing has enabled a broad range of discoveries and remarkable advances in nanoscience and materials science (for example, see Fig. 2). From spintronics with applications to ultrahigh-density magnetic storage, thermoelectric converters to the theory of high-temperature superconductivity, supercomputers have accelerated the pace of advances in both basic science and application. Computational predictions [8] supported by experimental confirmation [9] brought us the MgO-based magnetic tunnel junctions used in today's high-density hard drives, meaning computers, digital cameras, and MP3 players can now hold many gigabytes of data. Table 1 presents a brief list of similar accomplishments, with many more available in Yang and Chen 2007 [10]. Computational materials scientists clearly understand the impact and essential role of leadership computing in materials science and technology, as indicated by the following comments:




Fig. 2. Crystal structure of the high-temperature cuprate superconductor $\text{YBa}_2\text{Cu}_3\text{O}_7$. The superconducting copper-oxide planes are modeled with a two-dimensional Hubbard model [6]

- “Quantum mechanical ab initio calculations are usually done with 200 to 300 atoms. We’re doing calculations with a unit cell of more than 1,000 atoms. People would not be able to dream of doing these calculations without a large computing facility.” (Jihui Yang, General Motors)
- “Theory and computation can be of tremendous value in helping to understand experimental results but also can provide new insight into physical and chemical processes that would otherwise be extremely difficult to obtain.” (Bobby Sumpter, ORNL)
- “Without LCF and the INCITE [Innovative and Novel Computational Impact on Theory and Experiment] award, we simply cannot run our job. ... Due to the larger computer, we can do new science, not just higher-quality science.” (Lin-Wang Wang, Lawrence Berkeley National Laboratory)

Leadership computing allows for deeper penetration—higher-quality and higher-productivity science—into any investigation. This holds for a broad range of challenges in materials science and technology, including these examples:

- Nanomaterials theory
 - Emergent behavior/strongly correlated systems: spin fluctuations and pairing mechanism in cuprates (high-temperature superconductivity)
 - Spin and charge transport in nanostructures
 - Nanobiotechnology and statistical physics of nanoscale systems
 - Nanomagnetism (particles, wires, multilayers)

Table 1. Recent materials science accomplishments enabled by leadership computing at ORNL

Accomplishment	Description	Impact	Leadership computing role
Turning vehicle exhaust into electricity	Use of first-principles calculations to study a thermoelectric material [10]	Ability to convert a portion of waste heat generated by motor vehicle engines into usable electricity, making vehicles more fuel-efficient	Simulation of various properties of the material in a supercell of more than 1,000 atoms, rather than 200 to 300 atoms as in previous simulations
Resolving molecular attachment	Resolution of the molecular structure of successive layers of methane as they attach to magnesium oxide [11]	Improve our understanding of adsorption, a process with enormous importance in areas as diverse as fuel storage, manufacturing, and airport security	Enabling of calculations needed to understand this system
Casting light on high-temperature superconductivity	Solution of two-dimensional Hubbard model, showing a purely electronic model can describe superconductivity in high-temperature superconductors [12]	Pave the way for advances that may revolutionize power generation as well as electric transmission and reduce the cost of energy by adding significantly to the understanding of this phenomenon	Demonstration that Hubbard Model can be reliably used to study the physics of high-temperature superconducting cuprates
Boosting solar-cell efficiency	Use of Linear-Scaling Density-Functional Theory to study the electronic structure of zinc-tellurium alloy [13]	Use alloy in solar-energy applications; boost the efficiency of solar cells by as much as 50 percent by adding oxygen to the alloy	Study of very large systems (thousands of atoms) at the quantum mechanical level, feasible only on leadership-class computers using an aggregate of several thousand CPUs

- Structural materials
 - Micromechanics and nanomechanics of bulk and nanostructured materials
 - Nuclear fuel performance/structural reactor materials
 - Lightweight materials for transportation
 - Designing and modeling materials processing: phase stability, kinetics, etc.

2.2 NEAR-TERM (PETASCALE) SCIENCE DRIVERS

Several key materials science drivers (motivators), objectives (tangible deliverables), and impacts have been identified as leading candidates for leadership computing accomplishments within the next 3 years (Table 2). These anticipated achievements are based on detailed interactions with leading members of the computational materials science community [1, 14].

Table 2. Select materials science drivers for leadership computing at the petascale (1 to 3 years)

Application area	Science driver	Science objective	Impact
Nanoscale science	Material-specific understanding of high-temperature superconductivity theory	Understand the quantitative differences in the transition temperatures of high-temperature superconductors	Macroscopic quantum effect at elevated temperatures (>150K) New materials for power transmission and oxide electronics
	Thermodynamics of nanostructures	Understand and improve colossally magneto-resistive oxides and magnetic semiconductors	Magnetic data storage Economically viable ethanol production
		Develop new switching mechanism in magnetic nanoparticles for ultrahigh-density storage Simulate and design molecular-scale electronics devices	Energy storage via structural transitions in nanoparticles
	Evolution of an understanding of biological system behavior	Elucidate the physical-chemical factors and mechanisms that control damage to DNA	Medicine, biomimetics, sequence dependencies, and inhibiting agents of hazardous bioprocesses
Material response	Elucidation of the causes leading to eventual brittle or ductile fragmentation and failure of a solid	Understand macro-cracking due to coalescence of subscale cracks, local deformation due to void coalescence, and dynamic propagation of cracks or shear bands	Reduction of engineering margins to within required safe operating envelop

2.3 LONGER-TERM (SUSTAINED PETASCALE AND EXASCALE) SCIENCE DRIVERS

Materials science drivers, objectives, and impacts have been identified for leadership computing accomplishments considered possible on an exascale leadership computing platform deployed within the next decade (Table 3). These more speculative achievements are based on recent workshops [2] and personal communication with select members of the computational materials science community.

3. EARTH SCIENCE

3.1 RECENT ACCOMPLISHMENTS WITH LEADERSHIP COMPUTING

The LCF at ORNL provided more than a third of the U.S. contribution of computational resources to the February 2007 report of the Intergovernmental Panel on Climate Change (IPCC). High-performance computing guided the studies and conclusions that went into the report, leading to the 2007 Nobel Peace Prize for the IPCC in recognition of its work.

Earth science simulations at the LCF continue to play a key role in U.S. climate change research, bringing ORNL's leadership computing resources to studies of weather, carbon management, climate change mitigation and adaptation, and the environment, to name a few. For example, LCF systems provide much of the computing power for the Community Climate System Model (CCSM), a fully coupled, global climate model that provides state-of-the-art computer simulations of the Earth's past, present, and future climates (Fig. 3). In fact, the last few months of 2007 have seen many high-impact

Earth science accomplishments, some of which are outlined in Table 4. Recent efforts include the evolution, development, and application of CCSM; high-resolution global ocean studies of thermohaline circulation and deep water formation; carbon-land model intercomparisons; and the impact of snow emissivity on regional climate models (using the WRF code).

Table 3. Select materials science drivers for leadership computing at the exascale (10 years)

Application area	Science driver	Science objective	Impact
Nanoscale science	First-principles design of increasingly complex materials with specific, targeted properties	Understand and use isolated nanostructures to design materials made out of nano building blocks	Smart materials for nanoelectronics, photovoltaics, information technology, and medicine
	Predictive description of microscopic behavior of water to understand systems in aqueous environments	Perform molecular dynamics with forces found with Quantum Monte Carlo computations	Detailed understanding of the structure of water—fundamental understanding of biological systems
	Understanding of synthesis of alloy nanoparticles with potential impact for design of new catalysts	Define the thermodynamics of compositions of alloy nanoparticles	Magnetic data storage Economically viable ethanol production Energy storage via structural transitions in nanoparticles
	Physics of strongly correlated electron materials	Explain the fundamental mechanism of high-temperature superconductivity, including materials specificity and inhomogeneities	New materials for practical applications in oxide electronics and next-generation power transmission




Fig. 3. Instantaneous net ecosystem exchange viewed along the terminator. The eastern half of this image is in sunlight and the terrestrial ecosystems are taking up carbon (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 (shown in red).

Table 4. Recent Earth science accomplishments enabled by leadership computing

Accomplishment	Description	Impact	Leadership computing role
Eulerian and Lagrangian studies of turbulent transport in the global ocean	Simulation of the ocean at high resolution (0.1° longitude) for 100 years and use of a suite of passive tracers and Lagrangian particles to understand the role of eddies and water transport in the global ocean	Create most realistic global ocean circulation simulation to date, helping to correct deficiencies with earlier simulations, which could result in different conclusions being drawn in future climate scenarios using eddying ocean models	Significantly improved fidelity compared to earlier runs, so will be used as the initial condition for a fully coupled climate simulation of unprecedented scale Ability to simulate in 3 months on Jaguar what took 10 times as long on previous-generation machines
Carbon-land model intercomparison project	Development and critical analysis of two terrestrial biogeochemistry models	Develop self-consistent carbon cycle and land models required for predictive climate change simulation capability	Model development aided and accelerated by required intercomparison simulations
CCSM	Development, validation, and application of next-generation model for Earth's past, present, and future climate	Develop and apply key tool in the U.S. for prediction and policy decisions on climate change mitigation and adaptation scenarios	Principal developmental and application platform for IPCC simulations
Reactive flows in porous media	Modeling of uranium migration at Hanford 300 area and exploration of CO ₂ sequestration in reservoirs	Understand role of fingering in CO ₂ sequestration; develop Hanford 300 conceptual model	Ability to model three-dimensional systems at sufficient spatial and temporal resolution to capture relevant physical processes

Looking ahead, four abrupt climate events will soon be tested against proxy records to search for the best route for melting water flux from retreating glaciers. A study of ocean eddies and their role in ocean transport will tell us if an eddying ocean model changes projections sufficiently to justify the much higher computational cost. And work on CCSM will move forward with continued development of the CCSM carbon, nitrogen, and biogeochemistry packages.

LCF assistance to the climate science community goes beyond providing the world's most powerful supercomputers for open scientific research. LCF scientific computing staff also provide key software and algorithm contributions to the climate teams; these include the insertion of new and more efficient linear solver preconditioners, the introduction of new data-analysis and visualization tools for carbon dioxide transport, and aid in identifying and fixing bugs in the global ocean model. As indicated by the following testimonies from key climate scientists using the LCF systems, these contributions are well appreciated:

- “[On Jaguar,] we got 100-year runs in 3 days. This was a significant upgrade of how we do science with this model. Forty years per day was out of our dreams.” (Peter Gent, National Center for Atmospheric Research [NCAR])
- “The most impressive new result in 10 years.” (Peter Gent, NCAR, referring to the El Niño/Southern Oscillation)
- “[The LCF has] been instrumental in being a bridge between our science and development teams and the computational resources.” (Lawrence Buja, NCAR)

- “[The Scientific Computing Liaisons] really helped manage the code optimization.” (John Drake, ORNL)
- “Applying leadership class computing to the climate problem allows us to include more realistic processes in the model; run at higher resolutions, which results in better transient solutions; carry out more ensemble runs, which reduce the uncertainty of the results; and run long historical simulations, allowing us to demonstrate that models reproduce the important features of the real climate system.” (Warren Washington, NCAR, principal investigator, Climate Science Computational End Station Development and Grand Challenge Team)
- “We now have models that are highly scalable and are limited only by available hardware. DOE’s [the Department of Energy’s] proposal to make available the Leadership Computing Facility for this project provides an unprecedented match between capability computing resources and a state-of-the-art model.” (Venkatramani Balaji, National Oceanic and Atmospheric Administration’s Geophysical Fluid Dynamics Laboratory at Princeton University)

The concept and application of computational end stations as originally proposed in the LCF at ORNL has indeed come to fruition and is best typified by the successful Climate Science Computational End Station.

3.2 NEAR-TERM (PETASCALE) SCIENCE DRIVERS

As indicated in Fig. 4, in the next 3 to 5 years climate computational researchers will be focusing their efforts on improving understanding and predictability in three principal areas: weather (land, surface, atmosphere), climate variability (ocean, sea ice, biogeochemistry), and climate change impacts (carbon cycle, water cycle, economics). Several key Earth science drivers, objectives, and impacts have been identified as leading candidates for leadership computing accomplishments within the next 3 years (Table 5). These anticipated achievements are based on interactions with leading members of the computational Earth science community [1, 14].

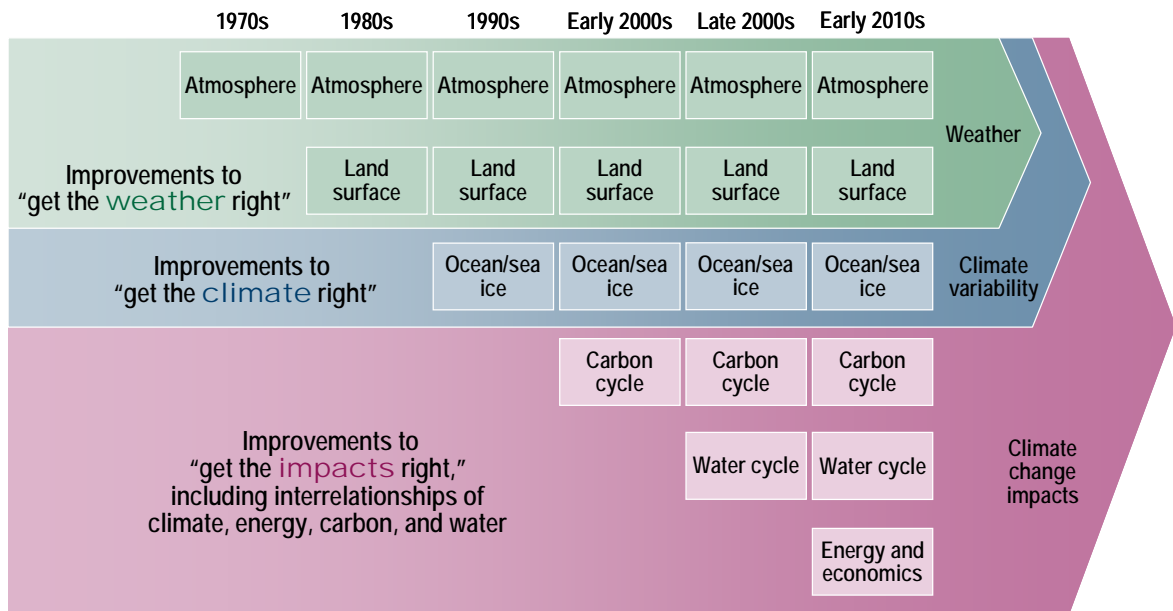


Fig. 4. Getting the weather, climate variability, and climate change impacts right in global climate models over the next decade will require the progressive evolution and coupling of multiple physical models.

Table 5. Select Earth science drivers for leadership computing at the petascale (1 to 3 years)

Application area	Science driver	Science objective	Impact
Climate	Simulation of dynamic ecological and chemical evolution of the climate system	Predict future climates based on scenarios of anthropogenic emissions and other changes from options in energy policies	Information for IPCC and policy decisions on climate change
	Development, delivery, and support for CCSM	Integrate and couple models for ocean, sea ice, land, atmosphere, cycles for carbon/water, and biogeochemistry	Tool for climate change prediction for input into policy making
	Accurate representation of ocean circulation	Develop fully coupled eddy-resolving ocean and sea ice model to reduce the coupled model biases where ice and deep-water parameters are governed by the accurate representation of current systems	Reduction of current uncertainties in coupled ocean-sea ice system model
Environment	Performance of multiscale, multiphase, multicomponent modeling of a three-dimensional field CO ₂ injection scenario	Include oil phase and four-phase liquid-gas-aqueous-oil system to describe dissipation of the supercritical CO ₂ phase and escape of CO ₂ to the surface	Demonstration of viability of and potential for a predictive groundwater transport model

Climate science drivers at the petascale reflect the ability to incorporate substantially more fidelity in the models. For example, CCSM at the petascale will be able to include the following:

- models for tropospheric chemistry (100 species),
- dynamic vegetation,
- terrestrial carbon and nitrogen cycles,
- ocean ecosystems,
- land ice sheets,
- stratospheric chemistry,
- full sulfur cycle,
- increase in ensemble size for climate change studies,
- coupled-ocean eddy-resolving simulations,
- cloud microphysics and interaction of aerosols with water,
- realistic land-use patterns, and
- tropical event simulation on climate timescales.

Biogeochemical processes, for example, will be modeled with an order-of-magnitude more variables than under the current approach [1]. Systematic errors and biases in climate models can also be reduced at the petascale, thereby increasing their utility as a predictive tool. Such reductions are likely to occur through improvements in existing physical parameterizations and a more accurate incorporation of phenomena. The internal dynamics of a system, for example, are more accurately represented at higher resolution.

3.3 LONGER-TERM (SUSTAINED PETASCALE AND EXASCALE) SCIENCE DRIVERS

Earth science and climate change research will be focused on two principal activities in the decade ahead:

- *mitigation*: evaluating strategies and informing policy decisions for climate stabilization and
- *adaptation*: preparing for committed climate change with decadal forecasts and regional impacts.

Simulations of 100 to 1,000 years will be typical for mitigation activities, while shorter simulations of 10 to 100 years will be used for adaptation. Each set of simulations must be predictive and quantifiable to reliably inform policy makers. Requirements for leadership computing can be tied to these activities and goals, as shown in Table 5. For example, estimates call for compute factors 10^{10} to 10^{12} greater than those of today to meet goals for spatial resolution, model completeness, simulation times, and breadth and depth of ensembles and scenarios (Fig. 5).

Climate models are currently more reliable at short timescales and long, asymptotic scales. Model predictions in the 20- to 50-year range are currently less accurate, resulting in more uncertain forecasts. Table 6 gives some select Earth science drivers, objectives, and impacts for leadership computing accomplishments considered possible on an exascale platform deployed within the next decade. These are based on recent workshops [2] and personal communication with select members of the computational Earth science community. This community has also done well in articulating what it believes to be attainable biogeochemical objectives over the next decade [2], including the following:

- integrated models and measurements of biogeochemical cycles,
- development of next-generation ecological models, and
- better theory for and quantification of uncertainty.

To reiterate, climate science opportunities at the exascale are abundant [15] and include advancement in the understanding of the following:

- decadal prediction on regional scales (accuracy in global models),
- climate extremes (heat waves, drought, floods, synoptic events, etc.),
- climate variability (low-frequency variability),
- water cycle (particularly in the tropics),
- human-induced impacts on the carbon cycle,
- sea-level rise (melting of the Greenland and Antarctic ice sheets), and
- abrupt climate change.

The rate limiters above are decadal prediction, abrupt climate change, and climate variability.

Issue	Motivation	Compute factor
Spatial resolution	Provide regional details	10^3 - 10^5
Model completeness	Add "new" science	10^2
New parameterizations	Upgrade to "better" science	10^2
Run length	Long-term implications	10^2
Ensembles and scenarios	Range of model variability	10^1
Total compute factor		10^{10} - 10^{12}

Fig. 5. Future climate modeling will require much more powerful supercomputers.

Table 6. Select Earth science drivers for leadership computing at the exascale (10 years)

Application area	Science driver	Science objective	Impact
Climate	Characterization and bounding of the coupled Earth system	Maintain tolerable time-integration rates while increasing model resolution and complexity; integrate models and observations; model biogeochemical cycles and coupled physical and biogeochemical systems at the process level	Understanding and prediction of stability and sustainability of rain forests, polar ice and ice sheets, agricultural ecosystems, precipitation, and methane hydrates; understanding of extreme weather; quantification of mitigation strategies
	Dynamical linking of socioeconomic and climate responses	Couple infrastructure, climate, demographic, informational, and energy economic models to predict adaptation as communities react to stresses on infrastructure systems and propose potential policies	Identification of future energy-infrastructure needs
	Decadal climate prediction	Incorporate cloud-resolving (1 to 5 kilometers) atmosphere, longer time integration (100 to 300 years, 1,000-year spin-ups), and larger ensembles (five to 20)	Understanding of and preparation for committed climate change
Environment	Radioactive waste management and environmental stewardship	Simulate radionuclide plumes at high resolution using stochastic parameter and property fields with meter and centimeter vertical resolution	Improvements to cleanup of DOE Complex
	Carbon sequestration in geologic formations	Model dissipation of supercritical CO ₂ injected in the subsurface accounting for fingering phenomena in kilometer-scale basin simulations	Management of carbon through active capture and storage

4. ENERGY ASSURANCE

The production, use, and flow of energy in the United States result in a complex system, as shown in Fig. 6 for the year 2002. This system of energy flows is tightly balanced between supply and demand, and it does not have a “silver bullet,” namely one particular energy type whose net change can significantly alter energy assurance on its own. Energy assurance—encompassing production, distribution, and consumption—is the ability to obtain, without costly disruption, the energy required by the United States in assured, economically-viable, and environmentally-benign ways to satisfy residential, commercial, and transportation needs. Energy assurance aims to avoid costly disruptions. While the causes of some such disruptions are beyond our control (weather, accidents, etc.), many others can be proactively controlled (reliability, redundancy, diversity, etc.). In those areas over which we do have control, small changes result in big impacts in energy supply, giving us many opportunities to further decrease the number of costly disruptions.

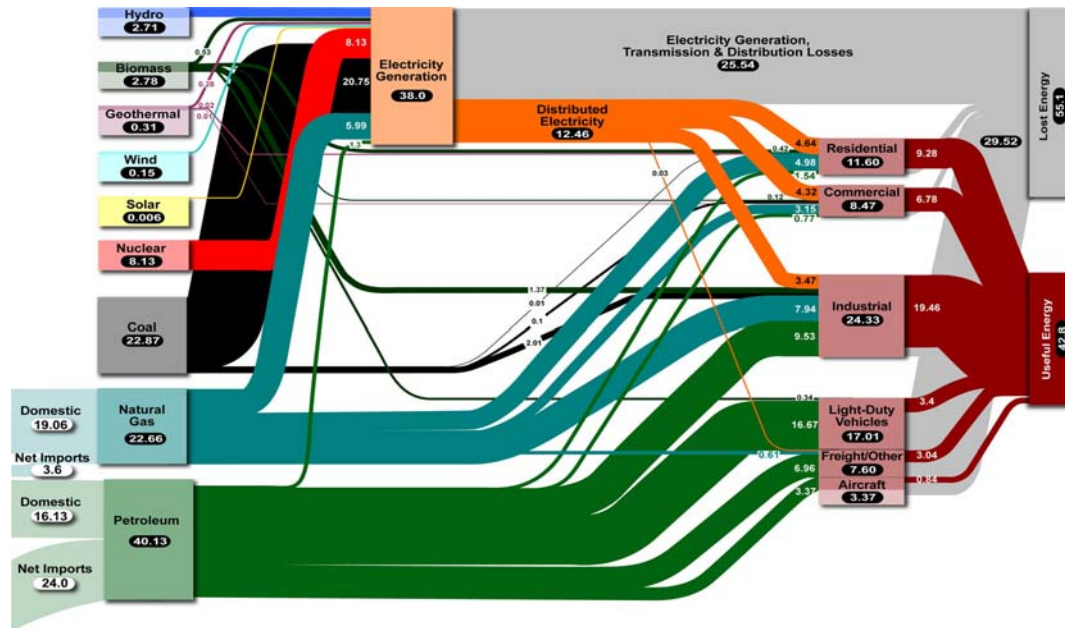


Fig. 6. U.S. energy flows in 2002. (Lawrence Livermore National Laboratory report, UCRL-TR-129990-02, 2002) [16]

These changes must incorporate improvements to our energy legacy (oil, gas, coal, nuclear, conservation, efficiency, etc.), as well as development of new approaches (wind, solar, fusion, biomass, hydrogen, carbon sequestration, new distributions, etc.). It is here that leadership computing is needed; its value will accelerate to the exascale and beyond for systems-level modeling [17].

Energy assurance has been and will continue to be a major driver for leadership computing resources, as illustrated by the accomplishments achieved to date and those that will be achieved by future petascale and exascale systems. In analyzing the science drivers, challenges, and likely outcomes, we are confident that exascale leadership computing will deliver an immeasurable return on investment in the form of increased U.S. energy assurance.

4.1 RECENT ACCOMPLISHMENTS WITH LEADERSHIP COMPUTING

Energy production research—both fundamental and applied—has benefited greatly from the use of leadership computing. Recent simulations at the LCF have provided invaluable insight into core turbulence and radio-wave heating of ITER-like tokamak plasmas for magnetic fusion, turbulent combustion processes in diesel and gas turbine engines, chemical catalysis, coal gasification, biomolecular dynamics and function, and fast nuclear reactor core flow and fuel performance. Table 7 illustrates a few of these accomplishments.

Testimonies from scientists on these teams illustrate the role of leadership computing and LCF staff in helping to deliver these accomplishments:

- “This allows us to look at the waves in ITER. We need much more resolution with ITER because it’s so big and the wavelength is so small.” (Fred Jaeger, ORNL)
- “If low-temperature compression ignition systems employing lean, dilute fuel mixtures make their way into next-generation autos, fuel efficiency could increase by as much as 25 to 50 percent.” (Jacqueline Chen, Sandia National Laboratories)
- “[Our LCF liaison] has been a crucial person in this effort, especially for code optimization.” (Jeff Candy, General Atomics)
- “NCCS [ORNL’s National Center for Computational Sciences] has enabled the vital

breakthrough. The allocation has helped investigating the multiple enzymes from multiple species in a short duration.” (Pratul Agarwal, ORNL)

Table 7. Recent energy assurance accomplishments enabled by leadership computing

Accomplishment	Description	Impact	Leadership computing role
ITER plasma core turbulence characterization	Validation of fusion plasma microturbulence	Identify and understand core plasma turbulence modes to be controlled in the ITER fusion reactor	More validation studies and increased physics fidelity
Wave-plasma interaction and extended magneto-hydrodynamics in fusion systems	Understanding of the radiofrequency (RF) heating mechanisms and the interaction of RF with magnetohydrodynamics (MHD) and other plasma processes in the ITER fusion reactor	Use at least 40 megawatts of RF power to heat the plasma in the \$13 billion ITER reactor, making it ten times hotter than the center of the sun	Higher fidelity and more predictive RF simulation more than triple the resolution of earlier simulations; coupled models of RF with MHD
Turbulent combustion characterization	Stabilization studies of lifted turbulent H ₂ /air jet flames in ignitive coflow	Develop fundamental understanding of “turbulence chemistry” interactions in combustion: flame stabilization in an ignitive environment relevant to diesel engines and gas turbines	Enabling of combined Reynolds number and chemical complexity consistent with laboratory-scale flames and autoignition experiments
Coal gasification pilot plant studies	Development and application of a gasifier model to help design gasifier systems for DOE’s FutureGen project	Design advanced, zero-emission fossil fuel plant	Performance scaling and optimization of multiphase flow model through new parallel algorithm insertion
Rational design of chemical catalysts	Dynamics of alkalis on catalytic metal surfaces and their promotion in Fischer-Tropsch processes	More efficiently produce synthetic petroleum substitute from coal or natural gas	Fast turnaround time for rigorous parametric studies
Biomolecular structure, dynamics, and function	Evolutionary conservation of protein structure based on reaction promoting vibrations	Enable low-cost bioethanol production	Fast turnaround time for reaction path simulations; reactive systems simulation instead of selective states

LCF staff members were indispensable to these achievements; their contributions included the importing of a new sparse linear solver, implementation of asynchronous input/output and automated end-to-end workflows, and achievement of more than 70% of peak performance in the compute-intensive portion of a fusion plasma application. Figure 7 is a direct numerical simulation of turbulent combustion recently performed on the Jaguar LCF system, namely the first fully-resolved simulation of a three-dimensional lifted flame in heated coflow with detailed chemistry.




Fig. 7. Direct numerical simulation of flame stabilization in a lifted hydrogen flame. Image shows the mixture fraction iso-surface (tan) and volume rendering of the OH radical concentration. [18]

4.2 NEAR-TERM (PETASCALE) SCIENCE DRIVERS

Petascale leadership systems, only a few years down the road, give us increased hope for tackling the nation's energy assurance challenges, as illustrated by the science objectives and impacts presented in Table 8. The following are among these possibilities:

- understanding how some microbial enzymes in nature (cellulase, etc.) break down cellulose into sugars;
- coupling wall, edge, and core physics into one integrated ITER simulation tool;
- deciphering the subtle coupling between turbulent transport and chemical kinetics controlling the stabilization of lifted fuel jets in lean, low-temperature combustion;
- exploiting insight provided by the first coupled, geometrically faithful, and physics-inclusive simulations of an entire nuclear reactor core; and
- understanding physicochemical design coprocesses—from atomic to continuum to geological time and length scales—involved in the reprocessing of spent nuclear fuel for producing recycled fuel and safe, stable waste forms.

In short, petascale leadership computing platforms will help to move energy assurance research from simplified, single-physics studies to more realistic systems that include the relevant phenomena. It will be an important next step toward predictability.

Figure 8 provides just one example of the types of bioenergy simulations that can be performed today.

4.3 LONGER-TERM (SUSTAINED PETASCALE AND EXASCALE) SCIENCE DRIVERS

Energy assurance drivers for exascale computing, as shown in Table 9, bring even more promise for new and innovative solutions to the nation's energy challenges. Highlights include the following:

- breaking down the natural resistance of plant cell walls by microbial and enzymatic deconstruction, collectively known as biomass recalcitrance [19], to enable economically viable biofuels;

Table 8. Select energy assurance drivers for leadership computing at the petascale (1 to 3 years)

Application area	Science driver	Science objective	Impact
Bioenergy	Exploration of the multiscale structure, dynamics, and function of enzyme complexes	Evolve multimillion-atom systems over 0.1 to 1.0 millisecond for multiple (200 to 2,000) trajectories	Identification of efficient means of converting biomass (cellulose) to ethanol
Fusion	Development of a quantitative, predictive tool for wave heating of ITER plasmas	Complete simulation of mode conversion heating in ITER with a realistic antenna geometry and non-Maxwellian alpha particles	Design and operation of ITER RF antenna
	Understanding and quantification of physics and properties of ITER scaling and H-mode confinement	Conduct strongly coupled and consistent wall-to-edge-to-core modeling of ITER plasmas; attain a realistic assessment of ignition margins	Design and operation of ITER
Chemistry	Computational catalysis	Accurately describe large systems with modern hybrid and meta density functional theory functionals	Generation of quantitative catalytic reaction rates and guidance of small-system calibration
	Heavy-element chemistry for advanced fuel cycles and environmental restoration	Simulate select liquid–liquid and gas–gas interfaces with accurate thermochemistry and spectroscopy	Replacement of many expensive experiments and shortening of timescales from decades to years for implementation of new nuclear fuel reprocessing
Combustion	Development of a predictive engineering simulation tool for new engine design	Understand flame stabilization in lifted autoigniting diesel fuel jets relevant to low-temperature combustion for engine design at realistic operating conditions	Potential for 50-percent increase in efficiency and 20-percent savings in petroleum consumption with lower-emissions, leaner-burning engines
Nuclear energy	Design, safety analysis, and licensing support of fast reactor core physics	Create coupled model of core thermal hydraulics, structural mechanics, neutronics, and fuel performance with acceptable resolution for geometry and neutron energy and direction	Reduction of reactor concept-to-license cycle time and cost; reduction of operating margins while maintaining adequate safety margins
	Converged reactor neutron flux and spectrum distribution	Analyze 10^{12} spatial elements and 30,000 discrete energy points; one calculation per temperature-feedback iteration per quasi-static time step for a total of 10^{21} unknowns performed 1,000 times	Accurate reactor core energy distribution and output for a given fuel type and age
	Design, safety analysis, and licensing support of separations reprocessing plants	Perform coupled, multiscale simulation of all major unit operations in a separations plant with acceptable resolution of chemical species in all streams	Reduction of plant concept-to-license cycle time and cost; production of economic products for fuel fabrication and stable waste forms for safe disposition




Fig. 8. Recent simulation of a cellulase enzyme attacking a cellulose substrate performed at the LCF [20].

Table 9. Select energy assurance drivers for leadership computing at the exascale (10 years)

Application area	Science driver	Science objective	Impact
Bioenergy	Biomass recalcitrance	Understand the complexity of plant cell wall structure and its relationship to recalcitrance through large-scale (microbial and plant cell wall structure and cellulosome, etc.) simulations of 10 to 100 million atoms over millisecond timeframes	The most important current barrier to the emergence of a cellulosic biofuels industry
Nuclear energy	Simulator for facilities within an operating closed fuel cycle	Develop an integrated set of models and simulations of the complete set of physical processes and facilities within an operating fuel cycle	A decision-making tool to help predict the outcome of changes made to the system as it operates
	Aqueous fuel reprocessing	Perform a first-principles simulation of interfacial transport in solvent extraction cycles for aqueous-based reprocessing	Design novel flowsheets with reliable scale-up to commercial reprocessing plants

Table 9. (continued)

Application area	Science driver	Science objective	Impact
Fusion	Reliable, whole-device modeling of ITER	Couple auxiliary heating, MHD dynamics, and plasma core and edge codes	Design and operation of ITER
	Predictive and self-consistent simulation of ITER plasma profile evolution	Understand anomalous transport at the plasma edge and MHD coupling to turbulence	Design and operation of ITER
Combustion	Understanding of “flameless” combustion of diverse fuels at high pressure in a turbulent environment relevant to advanced fuel-efficient, low-emissions engine concepts	Perform direct simulation of nonconventional, mixed-mode, turbulent combustion of biofuels under compression-ignition, aero-thermo-chemical regimes accounting for emission using statistical moments and models for particulate matter	New combustion systems designed to use alternative fuels with high efficiency while meeting stringent requirements on emissions
National security	Urban modeling for homeland security	Enable highly accurate three-dimensional radiation transport with up to ~1 million unknowns per grid cell using high numbers of moments, energy groups, and angles per cell	Threat analysis and prediction of improvised nuclear explosive effects
Chemistry	Systematic, large-scale exploration of optimal materials for catalysis or nuclear material separation agents	Combine density functional theory with evolutionary search for complex materials or an accurate combinatorial approach to screen the best separation material out of $O(10^3)$ compounds	Virtual design of catalysts and separating agents

- building and using a “virtual simulator” for advanced nuclear fuel cycles to guide the design and construction of new major nuclear facilities (thermal and recycling reactors, reprocessing facilities, and waste form treatment);
- demonstrating the scientific and technical feasibility of fusion power with the successful design and operation of the ITER facility [21];
- guiding and optimizing biofuels-consuming combustion engine design and operation by understanding fuel effects on fundamental combustion processes; and
- selecting and designing chemical catalysts quickly, cheaply, and efficiently via “virtual selection,” or with a minimum of experiments.

The bioenergy driver in Table 8 is just one example of the opportunities and insights afforded by exascale systems. The need to routinely model biomolecular complexes on their natural timescale, namely milliseconds (see Fig. 9), is currently impossible because today’s molecular dynamics simulation tools possess algorithms that must integrate forward in time with extremely small (10^{-12} second) time steps to maintain the requisite accuracy and stability. Exascale systems will put this requirement in the past, for the first time allowing the kinds of turnkey simulations of enzyme function and protein dynamics needed for real progress in understanding structure and function.

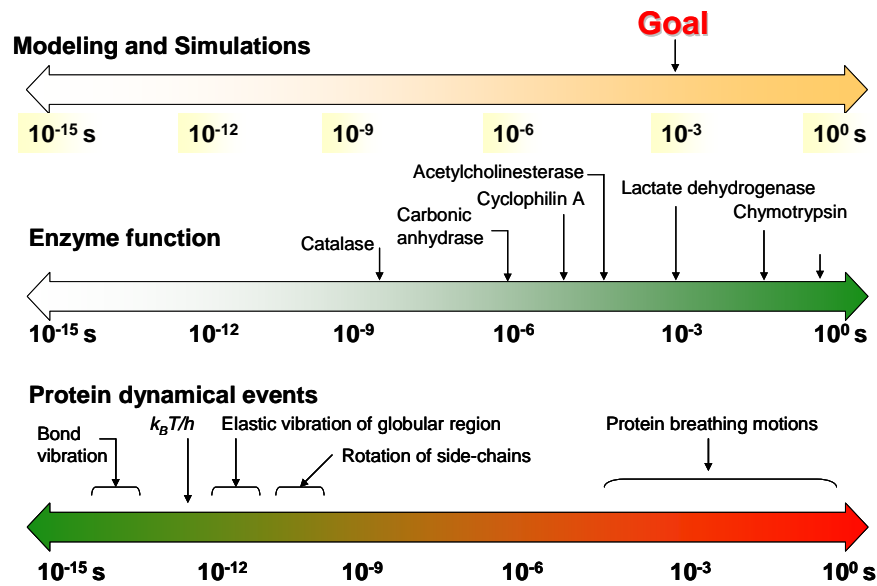


Fig. 9. Many protein and enzyme events of interest to biology and bioenergy research are in the millisecond-to-second timescales [22, 23], which are orders of magnitude beyond those possible with today’s simulations. Exascale leadership systems will change this landscape, making these simulations routine.

Recognition of the ITER facility’s dependence on petascale and exascale computing has led to the recently chartered Fusion Simulation Project (FSP), whose primary objective is to “create high-performance software to carry out comprehensive predictive integrated modeling simulations, with high physics fidelity, relevant to ITER” [24]. ITER is DOE’s highest-priority scientific facility [25]. The commitment to ITER and the urgent need for the FSP represent major energy assurance drivers for exascale computing. The FSP charter calls for development and application of a fully verified and validated comprehensive modeling capability to meet the following requirements [24]:

- Because each ITER discharge will be very costly (more than \$1 million), predictive, whole-device simulation is needed to optimize discharge scenarios.
- Simulation-guided controls must be developed to suppress large-scale instabilities adversely affecting confinement *prior* to completion of ITER construction.
- Accurate and reliable predictions are needed for edge transport barrier (enhancing core plasma confinement) and edge instabilities (causing fluctuations in divertor and first wall power).

ITER is dependent upon exascale computing for both optimal design (prior to 2017) and operational success (beyond 2017). This is true not only for the reasons cited above, but also because ITER is far larger than current experimental fusion reactors, with a toroidal magnetic field, minor radius, and electron temperature one to two orders of magnitude greater than those of current operational devices (CDX-U, DIII-D, etc.). As a result, a whole-device simulation of ITER will require 10^{12} more space-time points than current simulations of experimental devices. These 12 orders of magnitude cannot be achieved without both key software (algorithm) advancements and exascale computing [24].

U.S. energy assurance cannot be realized without science, technology, and innovation. A successful strategy will require prolific science and technology activity targeting all energy sectors to drive innovation and technology-based solutions. Such a strategy must use exascale leadership computing to meet an aggressive schedule—two decades—for achieving U.S. energy assurance.

5. FUNDAMENTAL SCIENCE

5.1 RECENT ACCOMPLISHMENTS WITH LEADERSHIP COMPUTING

Fundamental science—defined here as fields such as high-energy physics, accelerator physics, nuclear physics, and astrophysics—has traditionally relied heavily on computer and computational science as an enabling tool for discoveries, breakthroughs, and new insight. Leadership computing has arguably been an even more vital tool because the breadth of fundamental science exploration often increases with the size—in memory and speed—of the computational resource employed. A few recent accomplishments in fundamental science are illustrated in Table 10, with noteworthy advances that include the following:

- Astrophysics simulations at the LCF have decisively changed the standard model of core-collapse supernova evolution and explosion (Fig. 10). The (wholly computational) discovery of the standing accretion shock instability (SASI) is a fundamentally new piece of supernova theory, and the SASI has become the framing concept for most current supernova simulations. The first generation of multidimensional radiation-hydrodynamic simulations of core-collapse supernovas with sophisticated spectral neutrino transport and detailed nuclear burning are beginning to incorporate the effects of the SASI.
- The first-ever nuclear physics coupled-cluster simulations of ^{40}Ca have led to the determination of its ground state properties, including realistic three-body interactions.

Table 10. Recent fundamental science accomplishments enabled by leadership computing

Accomplishment	Description	Impact	Leadership computing role
Astrophysics core-collapse supernova pulsar mechanism explanation	Discovery of fundamental instability of supernova shocks through simulation	Determine that the SASI not only directly impacts the explosion mechanism in core-collapse supernovas, but also provides an explanation for pulsar spin-up	Ability to explore long (greater than half a second) physical timescales, essential to the discovery of the SASI
Nuclear physics ground state determination of ^{40}Ca	First-ever ^{40}Ca calculation using modern coupled-cluster theory and interactions, including triples correlations	Support nuclear structure research with exotic beams and enable reliable predictions of unknown nuclei crucial for astrophysical modeling	Required computation with several thousand processors
Astrophysics core-collapse supernova explosion mechanism explanation	Simulations of core-collapse supernovas with realistic neutrino transport and nuclear burning coupled to multidimensional hydrodynamics leading to explosions	Solve core-collapse supernova explosion-mechanism puzzle after more than 40 years	Ability to model neutrino transport, nuclear burning, and multidimensional hydrodynamics




Fig. 10. Rendering of the matter entropy during the nonlinear phase of the standing accretion shock instability in a supernova core [26].

Testimonies from scientists on these teams illustrate the role of leadership computing and LCF staff in helping to deliver these accomplishments:

- “We cannot do the medium-mass nuclei in coupled-cluster theory without the resources. The same holds for diagonalization techniques; mass tables from DFT [density functional theory] all require both turnaround and large runs.” (David Dean, ORNL)
- “With the machine being more stable and larger, it has been easier to perform production science runs. [The LCF liaison] has been instrumental in our progress this quarter. This is the very best aspect of the NCCS–LCF structure and relationship.” (Anthony Mezzacappa, ORNL)

5.2 NEAR-TERM (PETASCALE) SCIENCE DRIVERS

It is no surprise that the fundamental science community is prepared to take advantage of petascale computing with an abundance of fundamental science drivers and very specific objectives. This community, like the climate community, is very articulate and quantitative about its science needs and objectives and about what it expects to gain from petascale computing. In astrophysics, for example, petascale-level core-collapse supernova simulations will attempt to definitively determine the explosion mechanism and make quantitative statements regarding many of the most important observables (nucleosynthesis, gravitational waves, neutrino signatures, etc.). This is an exciting statement indeed, as the explosive mechanism has been both elusive and controversial for decades. This and other select fundamental science challenges at the petascale are highlighted in Table 11.

Table 11. Select fundamental science drivers for leadership computing at the petascale (1 to 3 years)

Application area	Science driver	Science objective	Impact
Astrophysics	Understanding of the core-collapse supernova mechanism for a range of progenitor star masses	Perform core-collapse simulations with sophisticated spectral neutrino transport, detailed nuclear burning, and general relativistic gravity	Understanding of the origin of many elements in the periodic table and the creation of neutron stars and black holes
	Meeting of need for ~1-percent accuracy in ability to make Type Ia supernova standard candles	Discriminate among Type Ia supernova detonation mechanisms and predict correlations for observers to use in supernova calibration	Better understanding of dark energy properties; guidance of Joint Dark Energy Mission instrumentation, scientific observation strategy, and data analysis
Lattice quantum chromodynamics	Computation of strong interactions between particles so precise that theoretical uncertainties no longer limit understanding of their interactions	Calculate weak interaction matrix elements of strongly interacting particles to the accuracy needed to make precise tests of the Standard Model; determine the properties of strongly interacting matter at high temperatures and densities, such as those that existed immediately after the Big Bang; calculate the masses of strongly interacting particles	Unification of the forces and particles of nature and understanding of the cosmos and destiny of the universe
Nuclear physics	Ab initio understanding of nuclear properties and nuclear reaction mechanisms	Compute from first principles the properties of medium-mass nuclei (mass 40 to 100) with two- and three-body nuclear forces	Predictive capability for nuclear properties and scattering cross sections relevant to DOE programs

5.3 LONGER-TERM (SUSTAINED PETASCALE AND EXASCALE) SCIENCE DRIVERS

Exascale simulation will usher in an era of precision predictive power for numerical experiments in fundamental science, examples of which are given in Table 12 for astrophysics, biology, and nuclear physics.

The impact of delivering on these objectives at the exascale cannot be overstated and include the following:

- understanding the chemical evolution of the galaxy;
- simulating human interaction with microbes (bacteria, viruses, etc.);
- designing drugs virtually;
- designing the International Linear Collider (ILC) accelerator;
- understanding the evolutionary state of a nucleus as it undergoes fission and fusion reactions;
- reconstructing protein structure and function from raw genomic data; and
- improving understanding of proteomics, protein folding, and docking.

Delivering on these and other fundamental science drivers with exascale leadership computing would, by any measure, represent a quantum leap forward in science advancement and discovery.

Table 12. Select fundamental science drivers for leadership computing at the exascale (10 years)

Application area	Science driver	Science objective	Impact
Astrophysics	Detailed simulations of core-collapse supernovas, including nucleosynthesis, gravitational waves, and neutrino signatures	Perform core-collapse simulations with Boltzmann neutrino transport and nuclear kinetics capable of isotopic evolution for a wide range of stars	Understanding of the chemical evolution of the galaxy and the place of supernovas
Biology	Protein and ligand interactions	Predict protein structure and classification, interacting protein partners, protein-protein complexes, and structure-function changes	Understanding of human interactions with microbes (bacteria and viruses)
	Conformational changes in the activation of a voltage-gated potassium channel that occurs on the surface of nerve, muscle, and secretory cells	Simulate biological systems (whole organisms such as cells), including multiple proteins undergoing assembly and folding	Regulation of cell-membrane excitability; repetitive, low-frequency firing in some neurons; and recovery of the nerve-fiber membrane at the end of the action potential
	Accurate, statistically-relevant, real-time biomolecular complex analysis	Perform classical atomistic simulations of 100,000-atom systems to 10-million-atom systems with multiple trajectories in parallel (ensemble-based approach) and ab initio quantum-mechanical/molecular-mechanical methods for quantum treatment of biological systems	Timely and cost-efficient drug design
Accelerator physics	Optimization and design of future particle accelerators for better efficiency at lower costs and development of advanced accelerator concepts	Include electromagnetic, thermal and mechanical effects for the ILC RF unit to determine optimal linear accelerator design	Increased return on investment of large DOE accelerator facilities
Nuclear physics	Deciphering of the evolution in time of fission and fusion processes	Use time-dependent coupled-cluster theory to investigate the time evolution of “below-the-barrier” events to deduce fragment mass and energy distribution	Reducing uncertainty in nuclear fusion and fission reactions

6. THE ROAD AHEAD

While conventional wisdom holds that scientific breakthroughs cannot be planned, we can be confident they will take place. The urgency of our scientific needs, combined with the development of

exascale computing resources and a driven, talented scientific community, will virtually ensure that the pace of scientific breakthrough continues to accelerate in the coming decade. In fact, with the deployment of increasingly capable leadership systems, the conventional wisdom will give way to a new reality: the planning of scientific discovery, simulation by simulation. Such planning has already been demonstrated for applications in climate, chemistry, combustion, astrophysics, fusion, and communication systems (see Table 13).

Table 13. Science opportunities and impacts grow with increasingly capable leadership systems

Application area	Science at 1 petaflop	Science at a sustained petaflop	Science at 1 exaflop
Climate	Mitigation: full chemistry; cycles of carbon, nitrogen, and sulfur; ice-sheet model; multiple ensembles Adaptation: high-resolution (1/4 degree longitude) atmosphere, land, sea ice, and ocean	Mitigation: increased resolution, longer simulations, more ensembles for reliable projections; coupling with socioeconomic and biodiversity models Adaptation: limited cloud-resolving simulations, large-scale data assimilation	Mitigation: multicentury ensemble projections for detailed comparisons of mitigation strategies Adaptation: full cloud-resolving simulations, decadal forecasts of regional impacts, and extreme-event statistics
Chemistry	Dynamics of few-electron systems such as interaction with intense radiation	Treatment of absorption problem with larger unit cells to avoid error sources	Extension of interaction with intense radiation to more realistic systems containing more electrons
Combustion	Stabilization of autoigniting diesel jets Predictive models accounting for turbulent transport coupling with multistage ignition and cool-flame chemistry in diesel combustion	Liquid jet combustion with spray dynamics, evaporation, and chemical reaction at phase boundaries Active control of spray shape and injection parameters for clean combustion in flexible-fuel vehicles	Emission in direct injection combustion using multiple statistical moments and stochastic models for particulate matter Reduction of soot and emissions from combustion of diesel and alternative fuels such as biobutanol
Astrophysics	Determination of the nature of the core-collapse supernova explosion mechanism Fully integrated, three-dimensional neutrino radiation hydrodynamics simulations with nuclear burning	Detailed nucleosynthesis (element production) from core-collapse supernovas Large nuclear network capable of isotopic prediction (along with energy production)	Precision prediction of complete observable set from core-collapse supernovas: nucleosynthesis, gravitational waves, neutrino signatures, light output Full three-dimensional Boltzmann neutrino transport, three-dimensional MHD and relativistic hydrodynamics, nuclear burning

Table 13. (continued)

Application area	Science at 1 petaflop	Science at a sustained petaflop	Science at 1 exaflop
Fusion	Study of mesoscale dynamics, structure formation, heat and particle transport in collisionless trapped electron mode turbulence Comparison of fluctuation spectra and zonal flow patterns with basic energy sciences data	Initial predictive capability for the effective coupling of the plasma core and the plasma edge region and for MHD dynamics and auxiliary heating of the plasma through RF waves	ITER predictive capability: ion-electron temperature gradient coupling with realistic mass ratios and core-edge coupling Coupling with edge code and looking into longer timescale phenomena
Communication systems (civilian and defense)	Study of interplay of weather, signal propagation, and related phenomena on system performance Study of security and performance of extant protocols on Internet-scale configurations	Increase fidelity of communications network physical-layer effects Incorporation of detailed communications-network protocols	Optimization of simulation-based design process for robustness and on-demand tuning for very rapid deployment

Planning for scientific discovery will become increasingly important in the coming years, but it must be preceded by planning for the tools of discovery—next-generation computing hardware, algorithms, and applications. We have every reason to assume that the nation’s research scientists, armed with these tools, will readily meet the challenges at hand.

7. REFERENCES

1. National Center for Computational Sciences. 2007. *Computational Science Requirements for Leadership Computing*, <http://nccs.gov/media-center/nccs-reports>, Oak Ridge, Tenn.
2. Department of Energy. 2007. *Modeling and Simulation at the Exascale for Energy and the Environment*, Office of Science, <http://www.mcs.anl.gov/~insley/E3/E3-draft-2007-08-09.pdf>, Washington, D.C.
3. Computing Research Associates. 2003. *The Roadmap for the Revitalization of High-End Computing*, http://www.nitrd.gov/subcommittee/hec/hecrtf-outreach/20040112_cra_hecrtf_report.pdf.
4. Department of Energy. 2003. *A Science-Based Case for Large-Scale Simulation*, Office of Science, available at <http://www.pnl.gov/scales/>.
5. Department of Energy. 2000. *Scientific Discovery through Advanced Computing*, Office of Science, <http://www.scidac.gov/SciDAC.pdf>.
6. Department of Energy. 2003. *The Challenge and Promise of Scientific Computing*, Office of Science, http://www.er.doe.gov/sub/Occasional_Papers/1-Occ-Scientific-Computation.PDF.
7. Oak Ridge National Laboratory. 2004. *National Leadership Computing Facility: A Partnership in Computational Sciences*, <http://www.ccs.ornl.gov/nlcf/>.

8. Butler, W. H., et al., “Spin-dependent tunneling conductance of Fe|MgO|Fe sandwiches,” *Physical Review B* **63**, 054416 (2001).
9. Parkin, S.S.P., et al., “Giant tunnelling magnetoresistance at room temperature with MgO (100) tunnel barriers,” *Nature Materials* **3**, 862–867 (2004).
10. Yang, J., and C. Chen. 2007, unpublished work.
11. Drummond, M., et al., “An Electronic Structure Investigation of Surface-Adsorbate and Adsorbate-Adsorbate Interactions in Multilayers of CH₄ on MgO(100),” *Journal of Physical Chemistry C* **111**, 966–976 (2007).
12. Schulthess, T., “Solving a Macroscopic Model for High-Temperature Superconductivity,” *Scientific Computing* (2006), <http://www.scientificcomputing.com/>.
13. Wang, L. W. 2007, unpublished work.
14. Department of Energy. 2006. *Petascale Applications Workshop*, Office of Science Advanced Scientific Computing Research Program Office Workshop, Washington, D.C.
15. Hack, J., and E. Bierly, *Computational and Informational Technology Rate Limiters to the Advancement of Climate Change Science*, presentation given to the DOE Advanced Scientific Computing Research Advisory Committee, November 6–7, 2007, <http://www.sc.doe.gov/ascr/ASCAC/presentationpage1107.html>.
16. Kaiper, G. V., *U.S. Energy Flow Trends—2002*, Lawrence Livermore National Laboratory, June 2004, UCRL-TR-129990-02, <https://eed.llnl.gov/flow/pdf/USEnFlow02-quads.pdf>.
17. Weigand, G., *Energy Assurance and High Performance Computing*, presentation given at the ORNL booth at Supercomputing 2007, November 2007.
18. Yoo, C. S., J. H. Chen, and R. Sankaran, *Direct numerical simulation of a turbulent lifted hydrogen/air jet flame in an autoignitive heated coflow*, presentation at the 2nd ECCOMAS Thematic Conference on Computational Combustion, 2007.
19. Mann, R., *BioEnergy Science Center: A DOE Bioenergy Research Center*, presentation at the 2007 Fall Creek Falls Workshop, Nashville, Tennessee, September 2007.
20. Himmel, M., J., et al. unpublished work (2007).
21. Information available at <http://www.iter.org>.
22. Agarwal, P. K., “Enzymes: An integrated view of structure, dynamics and function,” *Microbial Cell Factories* **5**:2 (2006), <http://www.microbialcellfactories.com/content/5/1/2>.
23. Agarwal, P. K., and S. R. Alam, “Biomolecular Simulations on Petascale: Promises and Challenges,” *Journal of Physics Conference Series* **46**, 327–333 (2006).
24. Keyes, D., A. Kritz, and W. Tang, *Fusion Simulation Project (FSP): Workshop Report*, presentation at the DOE Advanced Scientific Computing Research Advisory Committee (ASCAC) meeting, November 2007.
25. Department of Energy. 2007. *Four Years Later: An Interim Report on Facilities for the Future of Science: A Twenty-Year Outlook*, Office of Science, 2007, [http://www.sc.doe.gov/Scientific_User_Facilities/History/20-Year-Outlook-Interim%20Report\(August\).pdf](http://www.sc.doe.gov/Scientific_User_Facilities/History/20-Year-Outlook-Interim%20Report(August).pdf).
26. Blondin, J., and A. Mezzacappa, “Pulsar spins from an instability in the accretion shock of supernovae,” *Nature* **445**, 58–60 (2007).