The Salishan Conference on

HIGH-SPEED COMPUTING

LANL / LLNL / SNL

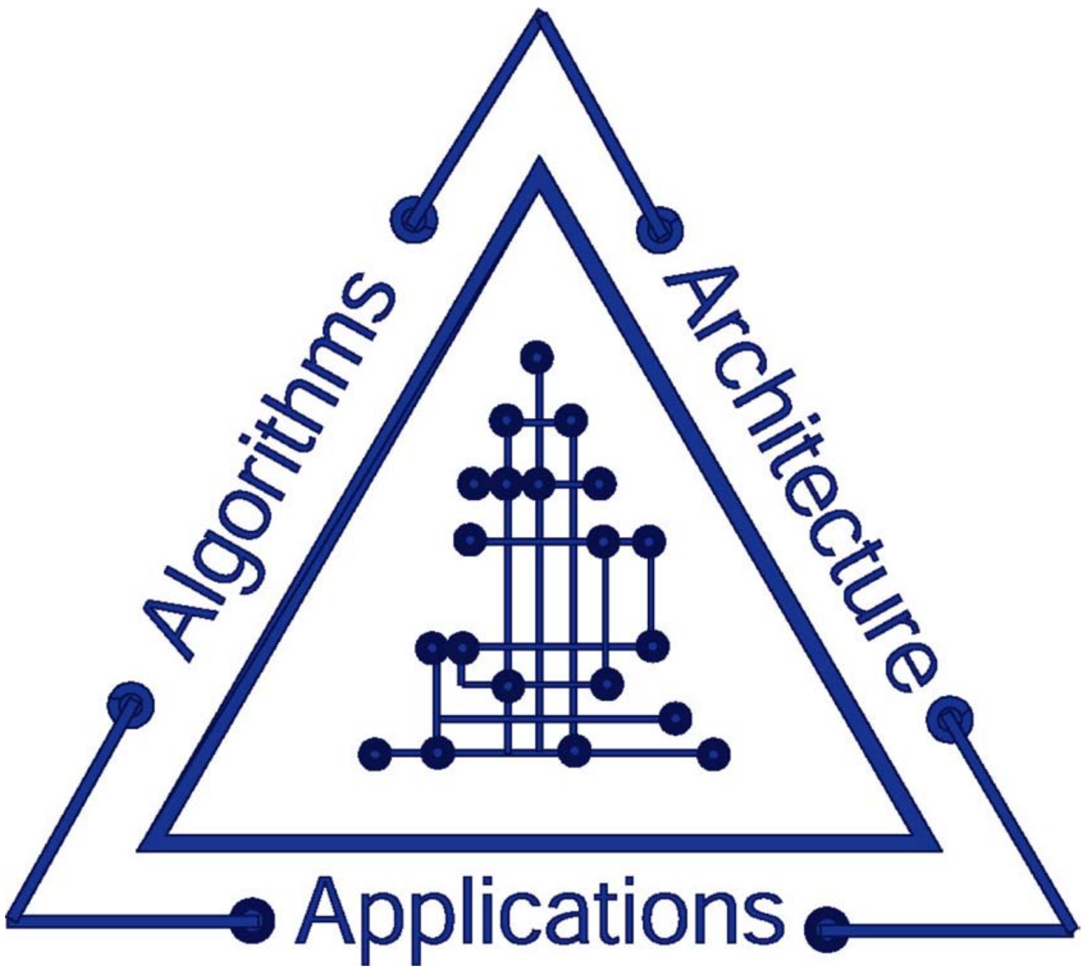

April 27 – 30, 2009

Salishan Lodge
Gleneden Beach, Oregon

The Salishan Conference on High-Speed Computing
at a glance

	Monday	Tuesday	Wednesday	Thursday
8:00 AM		Registration Opens Breakfast	Introduction to Sessions Breakfast	Introduction to Sessions Breakfast
8:30 AM		**Session 1:** *Chair – Alice Koniges* *Energy Security and Climate Change: A New Approach for Global Sustainability in the 21st Century* **Tomás Díaz de la Rubia**, *LLNL* ------------------------- *Energy Modeling and HPC* *John Mitchiner, SNL*	**Session 3:** *Chair – Richard Murphy* *Energy Efficient Computing From Bits to Buildings* *Horst Simon, LBNL* ---------------------- *Major Challenges to Achieve Exascale Performances* *Shekar Borkar, Intel*	**Session 4:** *Chair – Adolfy Hoisie* *Color Me Green! Color Me Fast!* *Wu Feng, VT* ------------------------ *Sequoia Sets New FLOP:Watt Standard* *Mark Seager, LLNL*
9:50 AM		Break	Break	Break
10:10 AM		*Computational Challenges at the Petascale and Beyond for Fusion Energy Sciences* *William Tang,* *Princeton University* ------------------------ *Computational Challenges in Geological Storage of Carbon* *Michael Celia,* *Princeton University*	*Impacts of Energy Efficiency on Supercomputer Programming Models* *Craig Stunkel, IBM* ------------------------ *Of FLITS and FLOPS: Balancing Energy and Interconnect Performance* *K. Scott Hemmert, SNL*	*Energy Smart High Performance Computing* *Moe Khaleel, PNL* -------------------------- *Energy Efficiency in High Performance Computing- An Industry Perspective* *Tahir Cader, HP*
11:30 AM		Panel Discussion	Panel Discussion	Panel Discussion
NOON		Lunch: Council House	**Lunch on Your Own**	Lunch: Council House
1:30 PM		**Session 2:** *Chair – James Ang* *The Outlook for Energy* *John Kuzan, Exxon Mobil* ------------------------ *Petascale Direct Numerical Simulations of Turbulent Combustion* *Jacqueline Chen, SNL*	**No Scheduled Session**	**Session 5:** *Chair – Manuel Vigil* *Green Data Centers* *Stephen Wheat, Intel* ----------------------- *Next Generation Supercomputer Project in Japan, Its Current Status and Applications in Life Science* *Ryutaro Himeno, RIKEN*
2:50 PM		Break		Break
3:10 PM	Registration 4:30-7:00 PM (Salal Room)	*Large-Scale Biomolecular Simulations: Biomedical and Bioenergy Applications* *Kevin Sanbonmatsu, LANL* ------------------------ *Computer Simulations for Nuclear Energy Applications* *Marius Stan, LANL*		*Jaguar: Powering and Cooling the Beast. What Comes Next?* Arthur "Buddy" Bland, ***ORNL*** ----------------------- *Roadrunner Science and Technology* *Andrew White, LANL*
4:30 PM		Panel Discussion		Panel Discussion
6:00 PM	**Welcome/Keynote Address** (Long House) *Potential Roles for High Speed Computing in the Obama Administration* *Vic Reis, DOE*	**Working Dinner/Speaker** (Council House) *Climate Prediction for Decision Support: Intellectual and Computational Challenges* *Philip Duffy, Climate Central*	*Random Access* *(Sign up to speak for 10 minutes)* (Long House)	Informal Discussions Council House
8:00 PM	Informal Discussions Council House	Informal Discussions Cedar Tree Room	Informal Discussions Council House	

MAIN LODGE MAP

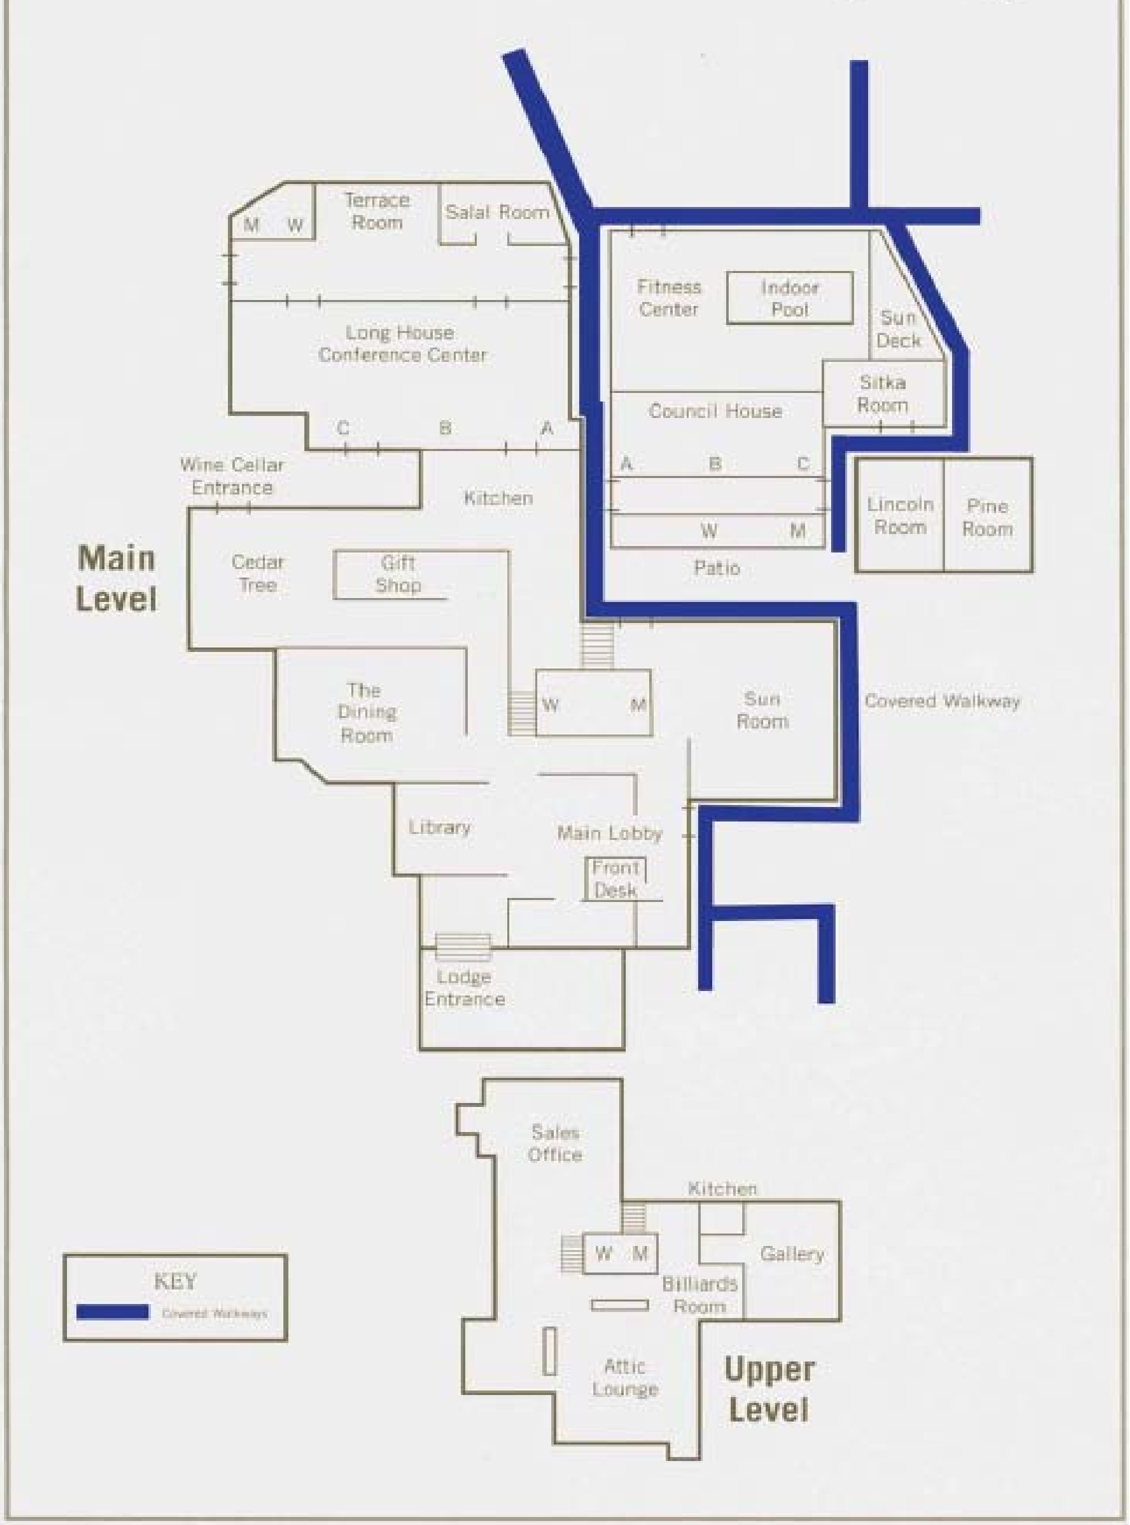

Welcome

Welcome to the Salishan Conference on High-Speed Computing. This conference was founded in 1981 as a means of gathering experts in computer architecture, languages, and algorithms together to improve communication, develop collaborations, solve problems of mutual interest, and provide effective leadership in the field of high-speed computing. Attendance at the conference is by invitation; we limit attendance to about 150 of the world's brightest people. Attendees are from national laboratories, academia, government, and private industry. We keep the conference small to preserve the level of interaction and discussion among the attendees.

The conference agenda and selection of participants has been designed to focus discussion on technical issues of relevance to our conference theme, "Energy-Centric HPC: Systems and Applications." The talks have been selected to give attendees information about the latest technologies and issues facing high-speed computing. The evening sessions are structured to encourage informal discussions and networking among all of the participants.

If you have any comments or suggestions for future topics and/or speakers, we encourage you to speak to any of the conference committee members.

We hope you find this conference stimulating, challenging, and also relaxing – enjoy!

Conference Committee

James Ang & Richard Murphy, *SNL*
Manuel Vigil & Adolfy Hoisie, *LANL*
Alice Koniges, *LBNL (formerly at LLNL)*
Maya Gokhale & John May, *LLNL*

Logistics

Conference sessions and the Random Access session will be held in the Long House. Lunches and the working dinner will be held in the Council House.

For administrative support, please speak to Dee Cadena, Erika Maestas and Gloria Montoya-Rivera located in the registration area (Salal room). If you have specific questions regarding audiovisual equipment or network connectivity, please seek out John Naegle or Dale Land.

Next conference dates:

April 26-29, 2010
April 25-28, 2011

Sponsorship

The Salishan Conference on High-Speed Computing is organized and hosted by Lawrence Livermore, Los Alamos, and Sandia National Laboratories. Additional sponsorship for the evening portions of our program is provided by the corporations listed here.

One of the highlights of the conference is the informal discussions held each evening. These sessions help us to go beyond the formal presentations to exchange ideas, solve problems, and develop friendships.

This year the following companies are helping to sponsor the evening sessions:

Advanced Micro Devices, Inc.

Cray, Inc.

Hewlett-Packard Company

IBM Corporation

Intel Corporation

Microsoft Corporation

NVIDIA Corporation

The Portland Group, Inc.

Silicon Graphics, Inc.

Sun Microsystems, Inc.

We would like to express our thanks to these companies for their generous support.

Table of Contents

THIS PAGE LEFT BLANK
INTENTIONALLY

Conference Theme

Energy-Centric HPC: Systems and Applications

The focus of the 2009 Conference on High-Speed Computing is energy. Two aspects of energy make this a timely theme. As interest grows in developing a national energy strategy, a wide spectrum of scientific and engineering high performance computing applications can be used to analyze and develop conventional, nuclear, and alternative energy technologies. Additionally, our conference participants, leaders in high performance computing, are acutely aware of the issues surrounding supercomputing in an energy-constrained environment.

While Moore's Law has provided steady, predictable increases in processor performance, power concerns have caused a shift in processor architecture toward multi-core designs. Unfortunately, the power efficiency of computation has declined with increases in computational performance. We have now reached the point where the cost of facilities and electricity will soon exceed the cost of a supercomputer. If we do not pursue initiatives in power-aware supercomputing, this imbalance will only get worse. Power-aware HPC will also increase the importance of resilience and reliability in the overall energy to solution metric, as the power cost of re-computing failed computations from a checkpoint may be unbearable compared with the cost of improving hardware resilience. Clearly, these new power constraints will fundamentally change the way in which future supercomputer architectures are designed, applications are developed, and systems are managed.

The conference program will cover a broad spectrum of energy-related topics, from systems to simulations. In the simulations arena, we'll cover timely topics related to development and analysis of conventional and new energy technologies. In the systems area, we'll address energy efficiency from device to data center level.

To this end, the 5 sessions of this year's Salishan conference are as follows:

Session 1: Energy Applications I
Session 2: Energy Applications II
Session 3: Energy Efficient Computing I
Session 4: Energy Efficient Computing II
Session 5: HPC Platforms Update

Conference Program

Energy-Centric HPC: Systems and Applications

Monday, April 27, 2009

4:30 -7:00 PM **Registration (Salal Room)**

6:00 PM **Welcome/Keynote Address**

Title: Potential Roles for High Speed Computing in the Obama Administration

Speaker: Vic Reis, *U.S. Department of Energy*

8:00 PM **Reception and Informal Discussions (Council House)**

Tuesday, April 28, 2009

8:00 AM	**Registration Opens (Salal Room)** Breakfast available (Terrace)
8:30 AM	**Session 1: Energy Applications I**

Title: **Energy Security and Climate Change: A New Approach for Global Sustainability in the 21st Century**
Speaker: Tomás Díaz de la Rubia, *Lawrence Livermore National Laboratory*

Title: **Energy Modeling and HPC**
Speaker: John Mitchiner, *Sandia National Laboratories*

9:50 AM **Break**
Refreshments available (Terrace)

10:10 AM **Session 1: Energy Applications I**

Title: **Computational Challenges at the Petascale and Beyond for Fusion Energy Sciences**
Speaker: William Tang, *Princeton University*

Title: **Computational Challenges in Geological Storage of Carbon Dioxide**
Speaker: Michael Celia, *Princeton University*

11:30 AM **Panel Discussion**

Tuesday, April 28, 2009 (cont.)

Noon **Lunch (Council House)**

1:30 PM **Session 2: Energy Applications II**

Title: **The Outlook for Energy**
Speaker: John Kuzan, *Exxon Mobil Corporation*

Title: **Petascale Direct Numerical Simulations of Turbulent Combustion**
Speaker: Jacqueline Chen, *Sandia National Laboratories*

2:50 PM **Break**
Refreshments available (Terrace)

3:10 PM **Session 2: Energy Applications II**

Title: **Large-Scale Biomolecular Simulations: Biomedical and Bioenergy Applications**
Speaker: Kevin Sanbonmatsu, *Los Alamos National Laboratory*

Title: **Computer Simulations for Nuclear Energy Applications**
Speaker: Marius Stan, *Los Alamos National Laboratory*

4:30 PM **Panel Discussion**

Tuesday, April 28, 2009 (cont.)

6:00 PM **Working Dinner/Speaker (Council House)**

Title: Climate Prediction for Decision Support: Intellectual and Computational Challenges

Speaker: Philip Duffy, *Climate Central, Inc. and University of California Institute for Research on Climate Change and its Societal Impacts*

8:00 PM **Reception and Informal Discussions** (Cedar Tree Room)

Wednesday, April 29, 2009

8:00 AM **Introduction to Sessions**
Breakfast available (Terrace)

8:30 AM **Session 3: Energy Efficient Computing I**

Title: **Energy Efficient Computing – From Bits to Buildings**
Speaker: Horst Simon, *Lawrence Berkeley National Laboratory*

Title: **Major Challenges to Achieve Exascale Performance**
Speaker: Shekhar Borkar, *Intel Corp.*

9:50 AM **Break**
Refreshments available (Terrace)

10:10 AM **Session 3: Energy Efficient Computing I**

Title: **Impacts of Energy Efficiency on Supercomputer Programming Models**
Speaker: Craig Stunkel, *IBM T.J. Watson Research Center*

Title: **Of FLITS and FLOPS: Balancing Energy and Interconnect Performance**
Speaker: K. Scott Hemmert, *Sandia National Laboratories*

11:30 AM **Panel Discussion**

Wednesday, April 29, 2009 (cont.)

Noon **Lunch on Your Own**

1:30 PM **No Scheduled Session**

6:00 PM **Random Access (Long House)**

The Random Access session consists of timely communications from participants on areas of interest to the Conference. Presentations are strictly limited to 10 minutes. A sign-up board is provided in the registration area.

8:00 PM **Reception and Informal Discussions (Council House)**

This conference sponsors a student poster session on Wednesday evening. Conference participants are encouraged to stop by the poster sessions and visit with the students.

Thursday, April 30, 2009

8:00 AM **Introduction to Sessions**
Breakfast available (Terrace)

8:30 AM **Session 4: Energy Efficient Computing II**

Title: **Color Me Green! Color Me Fast!**
Speaker: Wu Feng, *Virginia Tech*

Title: **Sequoia Sets New FLOP:Watt Standard**
Speaker: Mark Seager, *Lawrence Livermore National Laboratory*

9:50 AM **Break**
Refreshments available (Terrace)

10:10 AM **Session 4: Energy Efficient Computing II**

Title: **Energy Smart High Performance Computing**
Speaker: Moe Khaleel, *Pacific Northwest National Laboratory*

Title: **Energy Efficiency in High Performance Computing - An Industry Perspective**
Speaker: Tahir Cader, *Power & Cooling Strategist*, *HP*

11:30 AM **Panel Discussion**

Thursday, April 30, 2009 (cont.)

Noon **Lunch (Council House)**

1:30 PM **Session 5: HPC Platforms Update**

Title: **Green Data Centers**
Speaker: Stephen Wheat, *Intel Corp.*

Title: **Next Generation Supercomputer Project in Japan, Its Current Status and Applications in Life Science**
Speaker: Ryutaro Himeno, *Next-Generation Supercomputer R&D Center, RIKEN*

2:50 PM **Break**
Refreshments available (Terrace)

3:10 PM **Session 5: HPC Platforms Update**

Title: **Jaguar: Powering and Cooling the Beast. What Comes Next?**
Speaker: **Arthur "Buddy" Bland**, *Oak Ridge National Laboratory*

Title: **Roadrunner Science and Technology**
Speaker: Andrew White, *Los Alamos National Laboratory*

4:30 PM **Panel Discussion**

6:00 PM **Wrap-up, Reception and Informal Discussions (Council House)**

Abstracts

Keynote Address

Potential Roles for High Performance Computing in the Obama Administration

Vic Reis, *U.S. Department of Energy*

After a brief review of the history of high performance computing, we will discuss lessons learned from the DOE's Accelerated Strategic Computing Initiative (ASCI). Based upon input from various experts (e.g. speakers at this conference) we will suggest possible applications that might lead the future development of High Performance Computing, and list potential barriers. Input will be taken - indeed encouraged - from the audience, and displayed in near-real time. This will be compared with some of the specific goals set out by the Obama administration.

With your active participation, future generations will be astonished at the collective wisdom of the 2009 Salishan Conference.

Session 1: Energy Applications I

Energy Security and Climate Change: A New Approach for Global Sustainability in the 21st Century

Tomás Díaz de la Rubia, *Lawrence Livermore National Laboratory*

Projections by the Department of Energy's Energy Information Administration and most other international studies show that worldwide electric power demand will increase from the current level of about 2 Terawatts (TW) to 5 TW by 2050 and likely to as much as 10 TW by 2100. A recent IEA 2008 Energy Technologies Perspectives report shows that for the next 30 to 50 years burning fossil fuels will continue to provide most of the world's electricity. In fact, in these baseline scenarios CO2 emissions will be almost two and a half times the current level by 2050. In addition, the most recent report from the Intergovernmental Panel on Climate Change has placed a 90% likelihood that human sources of carbon dioxide emissions are significantly affecting the global climate. Clearly, this increasing demand is placing enormous pressure on natural resources, the global ecosystem, and international political stability. Alternative sources of energy are required in order to meet increased energy demand, stabilize the increase of atmospheric carbon dioxide, and mitigate the concomitant climate change. In response, governments are urgently trying to develop new economical, sustainable, and environmentally friendly energy technologies.

In this talk, I will discuss an approach to generating carbon-free, economically competitive power from nuclear energy that mitigates proliferation concerns, minimizes nuclear waste and eliminates concerns related to reactor core meltdown accident scenarios. The approach, Laser Inertial Confinement Fusion-Fission Energy (LIFE), combines a modest, neutron-rich fusion source with a subcritical fission blanket into an engine capable of generating several thousand Megawatts. A LIFE engine can utilize a variety of fertile and fissile fuels, eliminates the need for uranium enrichment and for Spent Nuclear Fuel reprocessing, and minimizes the production of long-lived actinides in nuclear waste to below DOE attractiveness level E (the lowest in the safeguards tables). LIFE thus represents a once-through, closed fuel cycle that can extend the capacity of current underground nuclear waste repository designs by factors of 20 to 100. Moreover, LIFE engines can burn the existing inventories of SNF and excess plutonium thereby drastically shrinking the nation's—and the world's—stockpiles of these special nuclear materials. Because LIFE is safe and minimizes proliferation concerns associated with the nuclear fuel cycle, we envision this technology as capable of providing a global solution to carbon-free energy generation in the 21st century. I will describe progress at LLNL's National Ignition Facility towards achieving fusion ignition and burn – the *sine qua non* condition for LIFE – and will discuss the specifics of the LIFE engine design and the basic and applied research challenges associated with making this vision a reality. I will close by discussing how success with LIFE could help meet the carbon-free energy demand gap for the planet and help mitigate potential climate change in the second half of the 21st century.

Energy Modeling and HPC

John Mitchiner, *Sandia National Laboratories*

Energy modeling, today, encompasses a broad array of spatial and temporal scales and physical phenomena including multi-physics, socio-economic context and even cognitive modeling of human behavior. This presentation will use energy modeling programs at Sandia National Laboratories to illustrate the breadth of current work. The discussion will include highlights from:

- Nuclear Energy Advanced Modeling and Simulation Program
- Nuclear Regulatory Commission Safety Analysis
- Burner Reactor Integrated Safety Code
- National Infrastructure Simulation and Analysis Center
- Photovoltaics Analysis
- Geologic Storage of Waste and CO_2
- Concepts for Designing and Controlling the Energy Grid dominated by Intermittent, Non-Dispatchable Sources

The talk will conclude with a discussion of future challenges in energy modeling and the potential for HPC to contribute.

Computational Challenges at the Petascale and Beyond for Fusion Energy Sciences

William Tang, *Princeton University*

Major progress in magnetic fusion research has led to ITER – a multi-billion dollar burning plasma experiment supported by seven governments (EU, Japan, US, China, Korea, Russia, and India) representing over half of the world's population. Currently under construction in Cadarache, France, it is designed to produce 500 million Watts of heat from fusion reactions for over 400 seconds with gain exceeding 10 - thereby demonstrating the scientific and technical feasibility of magnetic fusion energy. Strong research and development programs are needed to harvest the scientific information from ITER to help design a future demonstration power plant with a gain of 25. Advanced computations at the petascale and beyond in tandem with experiment and theory are essential for acquiring the scientific understanding needed to develop whole device integrated predictive models with high physics fidelity. Since ITER and leadership class computing are prominent missions of the DOE today, producing such a world-leading predictive capability for fusion represents a key exascale-relevant strategic project for the future.

Session 1: Energy Applications I

Computational Challenges in Geological Storage of Carbon Dioxide

Michael A. Celia, *Princeton University*

Carbon capture and storage (CCS) is an emerging technology for carbon mitigation. It involves injection of massive quantities of captured CO_2 into deep geological formations, where the CO_2 should remain for at least several hundred to several thousand years. Injection of CO_2 results in a multi-component, multi-phase flow system governed by highly nonlinear, coupled partial differential equations. Broad questions involving the fate of the injected CO_2, including possible leakage of CO_2 out of the injection formation and the fate of displaced fluids like resident brines, lead to very challenging modeling and analysis problems. Because important leakage pathways can be very localized, and their properties can be highly uncertain, an overall analysis of the system requires resolution of multiple length scales in the context of a probabilistic approach. These requirements render standard numerical simulators ineffective due to excessive computational demands. Development of simulation tools to answer important practical questions requires new computational paradigms that might include new solution algorithms as well as application of a set of simplifying assumptions to reduce the overall problem complexity. One way to develop simulators of appropriate complexity is through hybrid formulations that combine different numerical or analytical solutions across large domains within a multi-scale framework. Such an approach requires advanced new computational tools that can ultimate provide simulation capabilities for realistic, large-scale systems.

Session 2: Energy Applications II

The Outlook for Energy

John Kuzan, *Exxon Mobil Corporation*

The presentation will highlight Exxon Mobil's global energy outlook through 2030. The projections indicate that, at that time, the world's population will be ~8 billion, roughly 25% higher than today. Along with this population rise will be continuing economic growth. This combination of population and economic growth will increase energy demand by over 50% versus 2000. As demand rises, energy efficiency will become increasingly important. The pace of improvement is likely to accelerate, reflecting the development and deployment of new technologies for personal transportation and power generation. Ensuring access to energy for people around the world has been and will remain a great challenge. Effective technology solutions to the energy challenges before us will naturally rely on modeling complicated processes and that in turn will lead to a strong need for high-performance computing.

Petascale Direct Numerical Simulations of Turbulent Combustion

Jacqueline Chen, *Sandia National Laboratories*

Transportation is the second largest consumer of energy in the United States, responsible for 60% of our nation’s use of petroleum, an amount equivalent to all of the oil imported into the U.S. As the nation moves away from petroleum in the coming decades new alternative fuel sources will emerge. The overarching grand challenge is to develop a validated predictive modeling capability to tailor the design of diverse alternative fuel sources with next-generation engines for transportation. These new engines are expected to be characterized by higher pressures, lower temperatures, and higher levels of dilution to reduce energy consumption, pollutants and green-house gas emissions. They will operate in new combustion regimes that are more sensitive to fuel chemistry than conventional engines. Recent petascale Direct Numerical Simulation (DNS) have been performed to investigate underlying complex turbulence-chemical interactions that may be encountered in engines. In particular, DNS is used to understand how flames are stabilized in autoignitive environments and how mixture stratification affects flame propagation and structure. An end-to-end parallel data analysis and visualization pipeline is being constructed to glean insight from 100’s of terabytes of simulated field and particle data.

Session 2: Energy Applications II

Large-Scale Biomolecular Simulations: Biomedical and Bioenergy

Kevin Sanbonmatsu, *Los Alamos National Laboratory*

Biomolecular simulation applications easily exhaust petascale machines and present an excellent application area that pushes the state-of-the-art in supercomputing. We have performed the largest biomolecular simulation published to date, a record that has stood since 2005. Even this simulation is 8 orders of magnitude away from simulations of realistic time scales (~ seconds of physiological time). We will present results describing the molecular mechanism of the ribosome, a molecule central to all living systems and the target for 50% of the antibiotics used today. Our results are generated from all-atom explicit solvent molecular dynamics simulations and from recent coarse-grained studies which achieve longer time scales. In addition, we will present preliminary results of simulations of the cellulosome, a molecule that has solved the problem of cellulose processing (a key bottleneck for bioenergy production).

Computer Simulations for Nuclear Energy Applications

Marius Stan, *Los Alamos National Laboratory*

The increasing demand for energy supply is coupled with the need to explore the optimization of energy production, storage, and transportation. In the U.S. A., the nuclear energy community (both fission and fusion) is in the process of evaluating several paths forward that are strongly influenced by economic drivers and the political landscape. In all but one scenario, model development and computer simulations can assist in areas such as nuclear fuel fabrication, fuel performance, reactor design, hydrodynamics, and studies of reactor safety that include accident scenarios. In this context, the presentation reviews the state of the art and opportunities for using computer simulations, in synergy with theory and experiment, as the main investigation means and predictive capabilities for nuclear energy applications. Special attention is paid to requirements for multi-scale simulations, from atomistic to continuum, and code design elements, such as modularity and data transfer. A brief review of the major drivers and techniques for uncertainty quantification in nuclear energy demonstrates that computer simulations have the potential to save millions of dollars in nuclear plant licensing and fuel qualification. To turn that into reality, conservatism must be replaced by a realistic treatment of uncertainty evaluation involving changes in thinking, rhetoric, and the actual methodology. The talk ends with a discussion of experimental validation and international collaborations in this area.

Dinner Speaker

Climate Prediction for Decision Support: Intellectual and Computational Challenges

Philip Duffy, *Climate Central, Inc. and University of California Institute for Research on Climate Changes and its Societal Impacts*

The recent widespread acceptance of the human role in climate change, and of the likely harmful impacts of climate change, has fundamentally altered the field of climate modeling and prediction. Whereas until recently climate prediction was a research activity with few or no practical applications, planners now want to incorporate climate change into investment and other decisions made in the public and private sectors. Investments in long-lived physical infrastructure are particularly likely to be affected by climate change. The use of climate predictions in decision support requires that these predictions have fine resolution in space and time and rigorously quantified uncertainties. These requirements present difficult intellectual challenges, and lead to dramatically increased computational demands in a field that was already among the most demanding computationally. I will discuss these issues in general terms and illustrate with examples form recent work at DOE and elsewhere.

Session 3: Energy Efficient Computing I

Energy Efficient Computing – From Bits to Buildings

Horst Simon
Lawrence Berkeley National Laboratory

In a recent survey by IDC, facilities managers named power and cooling by an overwhelming majority to be the most pressing issues of concern to them. A study of Exaflops computing came to the conclusion that by projecting todays technology, an Exaflops computer might require 120 MW of power, if it can be built at all. A different study commissioned by the EPA estimates that power consumption by servers doubled in the period from 2000 to 2005 worldwide, and that total amount of electricity consumed by servers worldwide now costs about $7.2B. This is already today the same order of magnitude as the the investment in HPC technology ($9.2B). We have thus reached a critical threshold that should give us cause to consider the question of power consumption as a potentially limiting factor to the future growth in HPC. I will try to address this very question: what are the power limitations of current technology, and how can we change the equation to assure the future rapid growth of HPC performance without contributing even more to carbon emissions and global warming. In particular, I will discuss several research projects that we have started in Berkeley to address the issue of reducing power consumption in HPC, both at the systems and at the building level

Major Challenges to Achieve Exascale Performance

Shekhar Borkar, *Intel Corp.*

Compute performance increased by orders of magnitude in the last few decades—Mega to Giga took about 15 years, Giga to Tera in 12 years, Tera to Peta in 11 years. This was made possible by continued technology scaling, improving transistor performance to increase frequency, increasing integration capacity to realize novel architectures, and reducing energy consumed per logic operation to keep power dissipation within limit. The technology treadmill will continue to fulfill the performance demand, and one would expect to reach Exascale level performance in about 10 years. However, it's the same Physics that helped you in the past will now pose some barriers—"Business as usual" will not be an option.

First, the energy and power will pose as a major challenge. A straight forward extrapolation shows that an Exascale machine would consume in excess of a Giga-watt. Second, memory & communication bandwidth to feed such a machine with conventional technology would be prohibitive. Third, orders of magnitude increased parallelism, let alone explosion of parallelism created by energy saving techniques, would increase unreliability. And programming system will be posed with even sever challenge of harnessing the performance with concurrency.

This talk will discuss potential solutions in all disciplines, such as circuit design, architecture, system design, programming system, and resiliency to pave the road towards Exascale performance.

The author would like to acknowledge the Exascale study workgroup sponsored by Dr. Bill Harrod (DARPA, IPTO) and his colleagues for inspiring much of this work.

Impacts of Energy Efficiency on Supercomputer Programming Models

Craig Stunkel, *IBM T.J. Watson Research Center*

Driven by the insatiable performance requirements of current and future HPC applications, system designers are already considering how to build systems that can scale well beyond today's Petascale systems. Supercomputing systems must become drastically more power-efficient and reliable to make Exascale computing practical. The high-end application challenges will drive technology and architecture changes. Simpler, more power-efficient cores, perhaps running at lower frequency, may result in systems with ten million cores or more. SIMD units may play a larger role in computation. Special-purpose cores and/or accelerators have already been introduced to increase efficiency, resulting in heterogeneous systems. These and other innovations that increase energy efficiency will stress today's systems software stack. Likewise, they will impact applications and the programming models and languages that they rely on. Given these challenges, we will examine recent directions in programming models for HPC systems. We will also discuss potential migration paths for application designers that seek to scale to Exascale performance.

Session 3: Energy Efficient Computing I

Of FLITS and FLOPS: Balancing Energy and Interconnect Performance

K. Scott Hemmert, *Sandia National Laboratories*

As the high performance computing community pushes peak supercomputer performance faster than Moore's law, part counts are increasing, driving the power requirements to higher and higher levels. As this happens, energy costs become an ever-increasing percentage of the total cost of procuring and operating large machines. Given this, it is important to consider energy-to-solution when architecting and procuring systems for a given set of applications. This talk explores this theme from the point of view of the interconnection network, including the importance of well-balanced systems and finding more efficient ways to utilize available interconnect performance.

Session 4: Energy Efficient Computing II

Color Me Green! Color Me Fast

Wu Feng, *Virginia Tech*

The "Supercomputing in Small Spaces" project, established in 2001, sought to dramatically reduce energy consumption while maintaining high performance in order to improve the overall efficiency, reliability, and availability of supercomputers. In short, it sought to be (temperature) cool at a time when being cool wasn't cool. Since that time, the SSS project has continued to evolve in a multitude of green directions. For traditional clusters, the project has explored low-power systems architectures as well as commodity systems architectures that are enhanced with power-aware software. More recently, the project has leveraged emerging chip multiprocessing architectures, such as the Cell, GPGPU, and reconfigurable multicore, to deliver energy-delay products that are upwards of three orders of magnitude better than today's traditional chip multiprocessing architectures. So, where previously, supercomputers fought for elbow room at the red end of the spectrum, the move towards the green end of the spectrum continues to gain momentum, as evidenced by rapid advances by supercomputers on the Green500.

Sequoia Sets New FLOP:Watt Standard

Mark Seager, *Lawrence Livermore National Laboratory*

The introduction of the BlueGene line on the TOP500 list in June of 2004 marked the start of a new "low power" approach to High Performance Computing. The contemporary reigning TOP500 champion at that time was the "Earth Simulator" which had an astronomical FLOP:Watt 40TF/s:12.8MW = 3.1TF/MW was the pinnacle of the performance at any cost bubble. By focusing on system on chip design principles using embedded microcontrollers with custom HPC components, the first generation BlueGene/L achieved a breakthrough 367TF/s:1.3MW = 282.3 TF/MW or an amazing 91x improvement over the "Earth Simulator" and ushered in a new era of low power, high performance computing.

Since then the BlueGene team, jointly funded by DOE/SC and NNSA/ASC programs, has developed a second generation and is currently developing the third generation systems based on BlueGene technology that will be deployed to the ASC Program at the Lawrence Livermore National Laboratory as part of the ASC Sequoia build contract. In this talk we discuss system improvements in multiple areas with each generation including dramatic improvement in FLOP:Watt and usability. This talk will also cover the ASC program plan for improving the predictive simulation capability of applications by leveraging these systems for the Stockpile Steward Program.

Session 4: Energy Efficient Computing II

Energy Smart High Performance Computing

Moe A. Khaleel, Andres Marquez, Steve Elbert, Tahir Cader,
Pacific Northwest National Laboratory

Power and cooling will play an ever increasing role in high performance computing. For example, new extreme-scale supercomputers can use as much energy as about 20,000 homes.

By integrating software, hardware, sensors, energy flow, and heat management technologies, the Pacific Northwest National Laboratory's Energy Smart Data Center (ESDC) is driving supercomputer energy efficiency to unprecedented levels. Through specialized instrumentation, ESDC provides instant feedback on energy use, including the monitoring of air temperature, fluid temperature, and flow rate. The information collected can then be used by PNNL, along with its customers and partners, to design the most energy-efficient approaches to power and cooling.

The ESDC employs advanced supercomputer technology, characteristic in size and energy consumption of supercomputers used in typical data centers, making it an ideal testbed for independent evaluation and validation of new technologies and equipment. The testbed is highly instrumented with power distribution units, chillers, cooling towers, and other utilities. ESDC houses NW-ICE, a dedicated research cluster being used in the development of tools and metrics for data center productivity and data center energy efficiency in collaboration with our partners.

To design computing centers that are more energy efficient, we are first determining how much energy is used for powering the banks of processors, cooling the computers, and operating the facility. A valuable tool in obtaining these metrics is FRED (Fundamental research in energy Efficient Data Centers), developed by our researchers to measure parameters such as temperatures and power usage. The instrumentation and database capabilities provided by FRED enable the calculation of a variety of metrics that assess energy usage and efficiency of the ESDC.

Most of the commonly advocated metrics do not account for the absolute performance or useful work that is performed by a data center. DOE's National Energy Research Scientific Computing Center, for example, performed measurements on the Cray XT4 and concluded that the HPL code is a suitable workload for power/performance measurements, but they only ran it for three minutes—not long enough to see all the power phases. The ESDC has conducted a detailed study of the Weather Research and Forecasting Model and CP2K (Car-Parrinello atomistic and molecular simulations) to measure power as a function of useful work and concluded that HPL is not representative. Well-tuned codes such as HPL have distinctive power signatures that correlate well with performance whereas typical applicant codes do not, i.e., they can waste power while being computationally inefficient.

We advocate a holistic approach to metrics that includes the consideration of the energy used by the computer, the energy used in power delivery and cooling, partial performance metrics and absolute performance (e.g. time-to-solution). Using the ESDC testbed to the fullest, we recently conducted an experiment to demonstrate the calculation of Data Center Energy Productivity (DCeP)—a metric comparing the amount of useful work performed by the data center to the amount of energy consumed producing the work—and concluded that DCeP can be used to distinguish between different operational states in the data center.

Energy Efficiency in High Performance Computing – An Industry Perspective

Tahir Cader, *Power & Cooling Strategist, HP*

In its 2007 "Report to Congress on Server and Data Center Energy Efficiency", the EPA presented data that states that in 2006, data centers used 1.5% (60 billion kW-hr/year) of all the electricity produced in the US that year. In addition, the report claims that if nothing significant is done about the situation, this consumption will rise to 2.9% (>120 billion kW-hr/year) by 2011. While the EPA report covered all data center types, there is an identical, if not more pressing, crisis looming in the High Performance Computing (HPC) space. There are a number of major forces converging on HPC that are rapidly driving the cost of powering and cooling supercomputers to unsustainable levels: (1) rack power levels approaching or exceeding 40 kW/rack, (2) data center power densities approaching and exceeding 500 W/ft2, (3)rapidly inflating costs of energy brought on by diminishing oil supplies, political instabilities, and pending climate change legislation. In order to survive, HPC will have to adapt to these forces in addition to rapidly evolving technologies. What will HPC look like 5 years from now, even 10 years from now? The talk will focus on a number of areas that are likely to have a significant impact on HPC in the 5 – 10 year time frame. In particular the following questions will be posed and discussed:

Session 4: Energy Efficient Computing II

1. Will HPC be forced to consolidate the number of clusters here in the US?
2. What role will containerized data centers play in HPC?
3. Does it make sense to co-locate our data centers with power plants, or maybe bring the power plant to the data center?
4. Will highly instrumented/monitored and controlled HPC sites provide major relief?
5. Can data centers sited within microgrids help stem the rise in utility rates?
6. Air-cooled or liquid-cooled clusters?
7. What does a Sustainable HPC site look like? Will pending Cap and Trade legislation force HPC in this direction?

Session 5: HPC Platforms Update

Green Data Centers

Stephen Wheat, *Intel*

The deployed performance of the Top500 systems (www.top500.org) has recently been more than doubling every year, which is twice the rate of Moore's Law. As the insatiable need for performance is expected to continue, the systems themselves continue to become "larger", increasing the challenges of physical constraints such as floor space, power, and cooling. These factors are interrelated, as are means to innovate in them. Platform density will almost certainly continue to increase. Therefore, in addition to the larger total power and cooling loads, the localized factors at the cabinet level become critical. Nevertheless, localized approaches will need to comprehend the global challenge, particularly when the data center needs to be "Green" as well. This presentation will review a range of activities, including research and development projects, at Intel and in the industry, focusing on promising results for making HPC more efficient (green) while making data centers more capable. The presentation will also show that Green and higher performance can be complimentary terms as we move through the PetaScale era into the ExaScale era.

Next Generation Supercomputer Project in Japan, Its Current Status and Applications in Life Science

Ryutaro Himeno, *Next-Generation Supercomputer R&D Center*, *RIKEN*

We started the Next Generation Supercomputer Development Project in Japan, 2006. It will start operation from April in 2011 but it will be enhanced during 2011FY. Full completion will be in March, 2012. The system is a heterogeneous one composed of both vector and scalar units, whose performance goal is 10PetaFLOPS. We are currently in the final design stage of the system. We estimate the total electric consumption including all peripheral system as well as cooling system will be 30MW. A gas cogeneration system will provide electric power (not all) and chilled water for cooling.

We, RIKEN is currently focusing on developing wide variations of application software in the Life Science.

Session 5: HPC Platforms Update

"Jaguar: Powering and Cooling the Beast. What Comes Next?"

Arthur "Buddy" Bland, *Oak Ridge National Laboratory*

Computer designs from Cray have always been known for their outstanding performance as well as innovative power and cooling systems. The Cray XT5 system is no exception. In this talk, I will give an overview of the cooling systems that have been used on many systems over the years. I will then cover the features in Jaguar and ORNL's computing facility that promote energy conservation while delivering exceptional performance on applications, and finally, I will discuss some of the implications of this for future data centers.

Roadrunner Science and Technology

Andrew White, *Los Alamos National Laboratory*

Roadrunner is a resource for petascale science and engineering as well as an advanced architecture. It achieved the first sustained petaflop/s on LINPACK, June 2008, in Poughkeepsie, NY. The system is now in operation in Los Alamos. We will provide a brief up-to-date status of hardware, software and file systems. We will discuss several of the initial applications in detail and how they have been targeted for this platform. Finally, we will describe several newly observed effects of environment scale and complexity and their resolution.