The Salishan Conference on

HIGH SPEED COMPUTING

Algorithms
Architecture
Applications

April 23 – 26, 2018

Salishan Lodge
Gleneden Beach, Oregon

Welcome

The Association for High Speed Computing welcomes you to the Salishan Conference on High Speed Computing. This conference was founded in 1981 gathering experts in computer architectures, languages, and algorithms 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 only; we limit attendance to about 150 of the world's brightest people. Participants 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: Maximizing Return on Investment for HPC in a Changing Computing Landscape. The speakers have been selected to address our theme and give attendees information about the latest technologies and issues facing high performance computing (HPC). The evening sessions are structured to encourage informal discussions and networking among all 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 and/or complete the electronic survey at the end of the conference **(http://salishan.ahsc-nm.org/2018Survey.html)**.

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

Conference Committee: Katie Lewis and Robin Goldstone, *LLNL*
Jim Ang and Ron Brightwell, *SNL*
Carolyn Connor and Christoph Junghans, *LANL*

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, Gloria Montoya-Rivera, or Jan Susco, located in the registration area (Salal Room). If you have specific questions regarding audiovisual equipment or network connectivity, please seek out administrative support.

Visit our website at: **http://salishan.ahsc-nm.org**

Next Conference Dates: April 22-25, 2019 April 27-30, 2020 April 26-29, 2021

Main Lodge Map

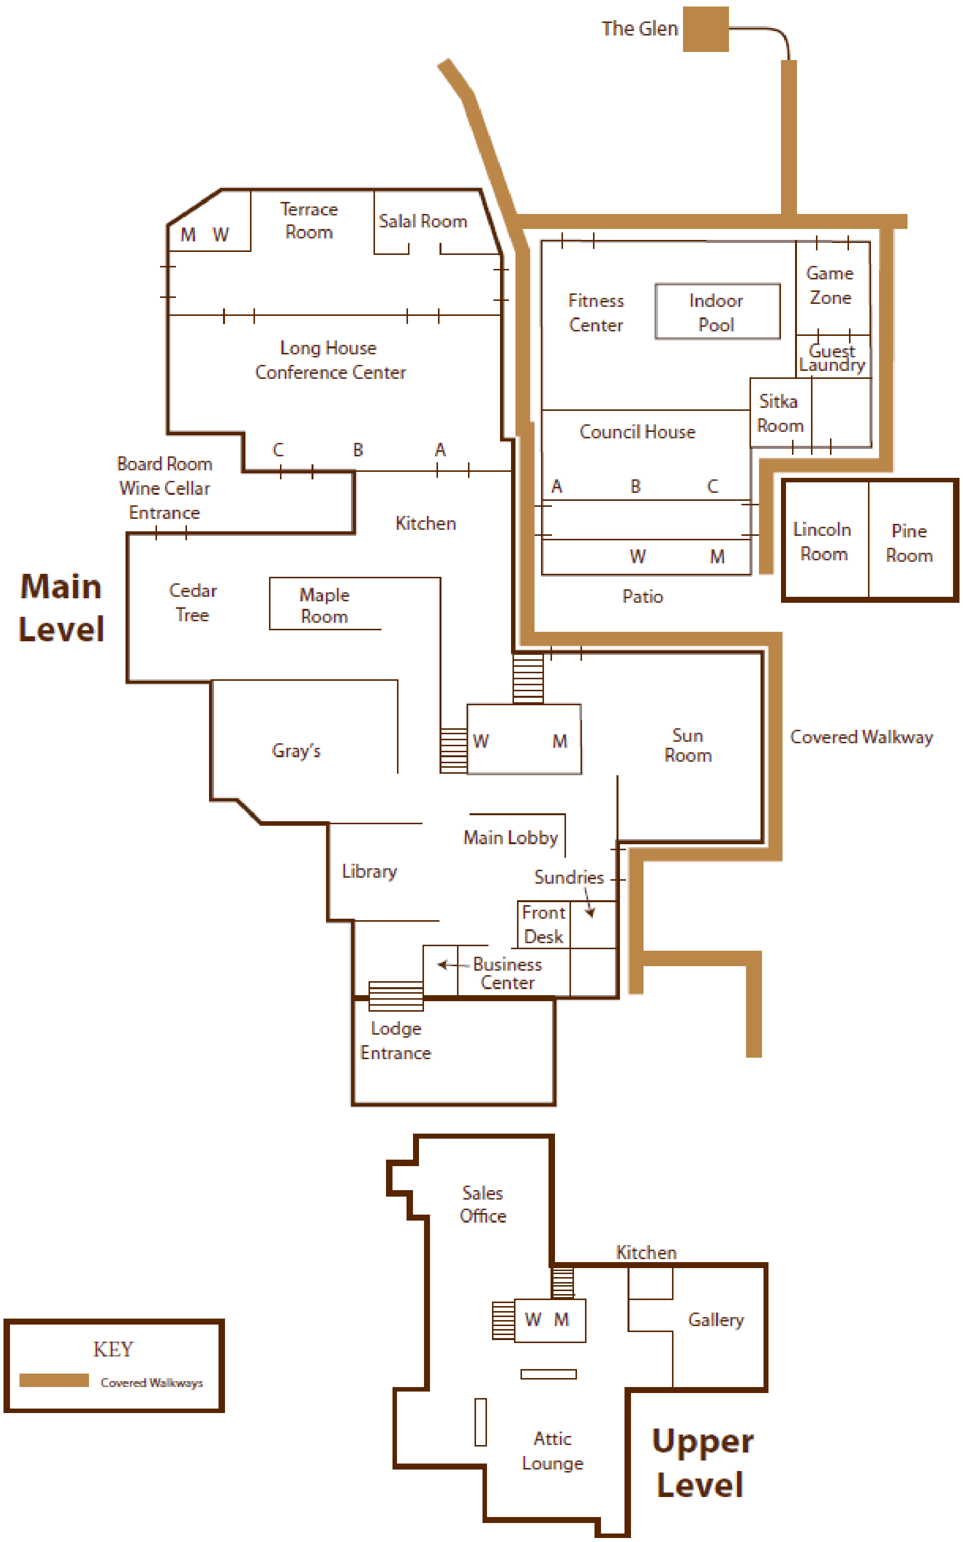

Table of Contents

Welcome and Logistics 1

Lodge Map 2

Sponsorship 5

Conference Theme 7

Conference Program

Monday: Keynote Address 11

Tuesday: Session 1: Maximizing Mission Science per Investment Dollar 12

Session 2: Removing Constraints, Engaging Challenges, Realizing Opportunity: An Era of New Ecosystems 13

Wednesday: Session 3: Defining the HPC Ecosystem 15

Thursday: Session 4: Leveraging Investments from the Past and Optimizing Future Investments in Applications 18

Session 5: Sharing Knowledge: Cross-Talk Between HPC and the Large-Scale Data Center Community 19

Survey 20

Abstracts 21

Attendees 35

Conference Notes 43

THIS PAGE LEFT BLANK
INTENTIONALLY

Sponsorship

The Salishan Conference on High Speed Computing is administered, hosted, and managed by The Association for High Speed Computing (AHSC). Additional sponsorship for the evening portions of our program is provided by the corporations listed here.

One of the highlights of the conference are 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.

arm

Cray, Inc.

DDN Storage

Dell EMC

D-Wave Systems, Inc.

Hewlett Packard Enterprise

IBM Corporation

Intel Corporation

Mellanox Technologies

Micron Technology, Inc.

NVIDIA Corporation

Penguin Computing

SUSE

We would like to express our gratitude to these companies for their generous support!

THIS PAGE LEFT BLANK
INTENTIONALLY

Conference Theme

Maximizing Return on Investment for HPC in a Changing Computing Landscape

History

The Salishan Conference on High Speed Computing was first established in 1981 to be a forum for the exchange of technical information and lessons learned among the DOE Weapons Labs regarding how to make effective use of the *new* Cray X-MP vector supercomputers. While these halcyon days were well before the time of any members of the organizing committee, we can envision how the development of capabilities like the Cray Time Sharing System (CTSS) benefited from the interdisciplinary discussions among our predecessors in the Long and Council Houses.

The Accelerated Strategic Computing Initiative (ASCI) and its follow on Advanced Simulation and Computing (ASC) Program has been investing in Applications, Software Environments, and Platforms for over two decades. These investments have always been guided by mission impact on the DOE and National Nuclear Security Administration (NNSA) Defense Programs. The cumulative Tri-lab ASC investment in our application portfolio is well over $1B at each of the Tri-Labs. This investment supports efforts to develop our Integrated Multi-physics and Engineering Analysis Codes, and Physics and Engineering Models; and to address confidence in our mission-critical applications through independent validation and verification of the ASC application portfolio. Mission impact led to the decision to move away from dependence on the custom Cray vector supercomputers to a strategy that was based on the integration of commercial, commodity computing components into large scale, integrated massively parallel processor systems. During the heyday of Moore's Law, this strategy really paid off well for ASCI/ASC.

The Changing Computing Landscape

The end of Moore's Law has two components. Over a decade ago, the end of Dennard Scaling imposed a constraint on the maximum clock rate, which led to the advent of multi-core CPUs and many-core designs like IBM's BlueGene line and GP-GPU accelerators. The main implication of this change is that realizing commodity performance improvements requires increasingly disruptive changes to application codes. The second component of the end of Moore's Law is that there are only a couple more generations of feature size shrink left. While there may still be additional transistors for computer architects to work with, the cost per transistor no longer significantly cheaper for each of these last feature size generations.

Return on Investment

Maximizing return on investment (ROI) requires targeting areas where impact can be best realized and then optimizing across several distinct areas (e.g., data movement, memory-to-flops ratios, and energy/power/water requirements, to name but a few). These, in turn, are affected by critical cost factors, including the cost of open source system software; hardware costs and capability tradeoffs (for both commodity and custom/semi-custom hardware design/development options); operating costs (e.g., energy, water, infrastructure); and the cost of refactoring applications for new architectures. The benefit of different approaches must also be considered, as we can all recognize that many times the lowerst cost option can actually cost more in the long run, or in a bigger picture perspective. Another important framework consideration is how far the HPC community should go with Total Cost of Ownership (TC), as there is an indication that European and Japanese HPC sites are increasingly focused on developing TCO models. The five sessions for the Salishan 2018 conference will explore different perspectives/dimensions on the ROI for future HPC directions.

Session 1: Maximizing Mission Science per Investment Dollar

Measuring return on investment in the world of HPC is a challenging prospect. The Top500 list has existed for years as a measure of how fast computers can run an optimized Linpack. We all recognize that the correlation between this metric and the impact of any particular machine toward a scientific objective is limited at best and a fool's errand at worst. So how do we assess the impact of the HPC dollars we invest? Is it the number of research publications generated? Is it the completion of mission objectives on schedule and on budget? Can it be obtained by surveying users and assessing their satisfaction with an architecture or code base? While many organizations concentrate on minimizing capital expenses and operating expenses (these things are measurable), the harder piece of the equation is knowing how to do so while simultaneously maximizing mission science capability. This session features speakers from a variety of mission spaces using High Performance Computers to meet their specific objectives. Discussion during this session will provide an opportunity to share insights on this important topic.

Session 2: Removing Constraints, Engaging Challenges, Realizing Opportunity: An Era of New Ecosystems

Some have argued that choices in the past have led us collectively down a path in which we no longer have a vibrant collection of architectures. What are the real costs of a non-vibrant hardware ecosystem? In terms of maximizing ROI for the HPC community, what should vibrancy look like moving forward? ARM and RISC-V ecosystems provide foundations for open source hardware technology development. These ecosystems enable users to develop their own design innovations to directly meet their needs, as compared to the more traditional model and/or technology roadmap in which innovation is driven largely by manufacturers. China presents a case study. Chinese computers have led the TOP500 list rankings for several years, and in the latest listing, China overtook the US in aggregate performance, as well. While the merits and flaws of this measure can and should be argued as part of a larger ROI discussion, China's TaihuLight supercomputer, which has maintained its number one ranking since topping the list in 2016, is uniquely different: it was built with Chinese-designed and manufactured processors, demonstrating the efficacy of China's System on a Chip (SoC) processor design investments. Is China operating under a fundamentally different set of technical constraints? If so, how do they differ? Have removing and/or changing constraints enabled fundamentally different hardware (and software) choices and/or development cycles and thereby led to fundamentally different system outcomes and capability advances? What far-reaching improvements in electronics performance, including considerations of architectures, design, and materials and integration, are possible well beyond the limits of traditional scaling? Looking beyond Moore's law, quantum computing systems are both similar and completely different from normal CMOS-based computing systems. What promise does quantum computing bring to HPC? What will be required to create a high productivity quantum computing ecosystem?

Session 3: Defining the HPC Ecosystem

Twenty years ago it was relatively straightforward to describe the field of high performance computing. Highly specialized systems pushed the boundaries of scale and performance to enable science and engineering breakthroughs. There were key capabilities and characteristics that defined and distinguished HPC platforms and the HPC environment . It's not clear that this is still true today. HPC has begun to adopt more technologies developed for other markets, such as video games, and the emergence of large data centers that support cloud computing services are now scaling at and beyond the largest HPC systems. Companies deploying these large data center systems are now developing their own custom hardware solutions, which provide significant performance improvements that enable breakthroughs in data analysis through machine learning. As the HPC community adopts more technologies developed for other markets and the broader computing community develops solutions to what used to be HPC challenges, what really defines HPC today? This session focuses on identifying differentiating technologies for the HPC ecosystem and whether or not these capabilities will continue to be differentiating in the foreseeable future. This session will address the following questions: Where is HPC converging or diverging with the broader computing community? Is the list of differentiating capabilities for HPC growing or shrinking? What technologies should HPC be trying to adopt? Is HPC becoming more specialized? Does the focus on high performance scalability really continue to differentiate HPC?

Session 4: Leveraging Investments from the Past and Optimizing Future Investments in Applications

The Accelerated Strategic Computing Initiative (ASCI) was established in 1995 and was formalized as the Advanced Simulation and Computing (ASC) Program in 2004. The success of this program is unquestionable, with its first-of-a-kind supercomputers, uniquely scalable scientific applications, and strong partnerships with academia. Nonetheless, there are investments made outside of the ASC Program and the DOE universe that can and should be leveraged, with an understanding that adopting external solutions is never "free". During this session, we will discuss advancements in large-scale simulations from industry and academia along with contributions from national laboratories into the external scientific computing community. We will address strategies for engaging with the external scientific applications community and how our path toward exascale computing can benefit. We will also discuss where ASC challenges are unique and some of the realities of trying to rely too heavily on external solutions. What is the right balance between external software products and internal solutions? What are the costs and the benefits of relying on external software, in Advanced Technology Systems (ATS) and Commodity Technology Systems (CTS) environments? What contributions should national laboratories make to the large-scale, scientific computing ecosystem, and how can that affect the ASC Program?

Session 5: Sharing Knowledge: Cross-Talk Between HPC and the Large-Scale Data Center Community

The large-scale data centers and the high performance computing community deal with very similar problems, for example when it comes to storage of large data sets and virtualization of applications. In this session, we will explore how large-scale data center technologies have inspired solutions in the HPC stack and how these technologies can act as a research platform for HPC. Specifically, we will discuss how virtualization technologies, like docker, can be used in HPC and how one can build a large-scale HPC platform with them. On the storage side, we will look how successful commercial products have to be developed to fit into the HPC stack. Last, but not least, some earlier experiences in using cloud computing for HPC related applications will be shared. Where is HPC storage different from large-scale data centers? Why can’t all compute happen in the cloud? How do commodity solutions have to be modified to fit in the HPC stack? What features are missing in the cloud?

Keynote and Dinner Speakers

This year’s conference will be kicked off with a keynote address by Mark Anderson, the new Advanced Simulation and Computing (ASC) Program Office Director, who will be able to share his vision for the program. Our dinner talk will be given by Andrew Connolly, a professor of astronomy from the University of Washington, who will speak about the Large Synoptic Survey Telescope (LSST) project and some of their large data analysis challenges.

Conference Program

Maximizing Return on Investment for HPC in a Changing Computing Landscape

Monday
April 23, 2018

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

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

Title: **Computing Challenges for the Third Decade of Stockpile Stewardship**
Speaker: Mark Anderson,
National Nuclear Security Administration

7:00 pm **Reception and Informal Discussions**
(Immediately following the Keynote in the Council House)

Tuesday
April 24, 2018

7:30 am **Registration opens (Salal Room)**

Breakfast available (Terrace)

8:30 am **Session 1:** **Maximizing Mission Science per Investment Dollar**
Chair: Robin Goldstone

Title: **HPC and its Role in Maximizing Weapons Program Investments**
Speaker: Chris Clouse, *Lawrence Livermore National Laboratory*

Title: **Maximizing the Impact of DoD HPCMP Investments**
Speaker: Roy Campbell, *Department of Defense HPC Mod*

Title: **Can the Tension Between Projectization and Agility/Innovation for Apps and Software Development Motivate Better ROI?**
Speaker: Doug Kothe, *Oak Ridge National Laboratory*

10:00 am **Break**

Refreshments available (Terrace)

10:30 am **Title:** **30 Years of HPC from Gigaflops to Exaflops: No Science Applications Left Behind or Thinning the Herd?**
Speaker: Steve Plimpton, *Sandia National Laboratories*

11:00 am **Panel Discussion**

Tuesday
April 24, 2018

12:00 pm	**Lunch (Council House)**	
1:30 pm	**Session 2:**	**Removing Constraints, Engaging Challenges, Realizing Opportunity: An Era of New Ecosystems** **Chair: Carolyn Connor**
	Title:	**International Trends in the HPC Market and an Update on the ROI from HPC Investments with a Focus on Chinese, European and Japanese Exascale Plans**
	Speaker:	Earl Joseph, *Hyperion Research*
	Title:	**Opportunities and Challenges of the COTS IP Model in HPC**
	Speaker:	Eric Van Hensbergen, *arm*
	Title:	**Building a Universal Silicon Compiler**
	Speaker:	Andres Olofsson, *DARPA*
3:00 pm	**Break**	
	Refreshments available (Terrace)	
3:30 pm	**Title:**	**Quantum Computing from the Computer Science Perspective: What Will the Software Ecosystem Look Like?**
	Speaker:	Mark Heiligman, *IARPA*
4:00 pm	**Panel Discussion**	

Tuesday
April 24, 2018

6:00 pm	**Working Dinner/Speaker (Council House)**	
	Title:	**Surveying the Sky with the LSST: Software as the Instrument of the Next Decade**
	Speaker:	Andrew Connolly, *University of Washington*
8:00 pm	**Reception and Informal Discussions** *(Immediately following the Working Dinner in the Cedar Tree Room)*	

Wednesday
April 25, 2018

7:30 am	**Breakfast available (Terrace)**	
8:30 am	**Session 3:**	**Defining the HPC Ecosystem** **Chair: Ron Brightwell**
	Title:	**The HPC Ecosystem at Amazon Web Services**
	Speaker:	Brian Barrett, *Amazon Web Services*
	Title:	**What is a Supercomputer and Why Do We Get Confused About How to Define Them?**
	Speaker	Bill Kramer, *University of Illinois/ National Center for Supercomputing Applications*
	Title:	**Interesting Times Ahead**
	Speaker:	Bob Lucas, *University of Southern California/ Information Sciences Institute*
10:00 am	**Break**	
	Refreshments available (Terrace)	
10:30 am	**Title:**	**HPC Workflows – Ecosystem Impact and Challenges**
	Speaker:	Robert Clay, *Sandia National Laboratories*
11:00 am	**Panel Discussion**	

Wednesday
April 25, 2018

12:00 pm **Lunch on your own**

1:30 pm **No Scheduled Sessions**

5:00 pm **Random Access (Long House)**

We invite Salishan 2018 attendees who want to propose a Random Access talk to sign-up online (maximum of one talk per participant) at: **http://salishan.ahsc-nm.org/2018RandomAccess.html.** Only a talk title and abstract tweet (max. 280 characters) are needed at time of sign-up. Alternatively, sign-ups can be done in person with the conference organizers. Sign-ups close Tuesday at 8:00 pm. *In fairness to at-large attendees, invited session speakers are requested to refrain from submitting proposals for Random Access talks*.

On Wednesday morning, a comprehensive list of talk titles and tweet abstracts will be published online for voting (an email will be sent to attendees with the link to vote). Every attendee will be allowed three equally-weighted votes until polls close on Wednesday at 1:00 pm. The talks with the most votes will be presented in the Random Access session starting on Wednesday at 5:00 pm. The final agenda of talks will also be posted online by Wednesday at 3:00 pm.

Wednesday
April 25, 2018

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

Student Poster Session (Council House)

This conference selects and hosts students from various universities, inviting them to present posters and discuss their research with our Salishan participants. All conference attendees are encouraged to visit with this year's students.

Thaleia Dimitra Doudali, *Georgia Technical Institute*
"Data Management in Heterogeneous Memory Systems"

Ivo Jimenez, *University of California, Santa Cruz*
"Reproducible Evaluation of Computer Systems"

Nicholas Lewis, *University of Minnesota*
"Subversive Architectures: Roadrunner and the Path to Exascale"

Massimiliano Lupo Pasini, *Emory University*
"Deterministic and Stochastic Acceleration Techniques for Richardson-type Iterations"

Vinay Ramakrishnaiah, *University of Wyoming*
"Facilitating the Scalability of ParSplice for ExascaleTestbeds"

Nicholas Stegmeier, *South Dakota State University*
"Parallel CFD with Regent"

Chad Wood, *University of Oregon*
"SOSflow: A Scalable Observation System for Introspection and In Situ Analytics"

Michael Wyatt, *University of Delaware*
"PRIONN: Predicting Runtime and IO using Neural Networks"

Thursday
April 26, 2018

7:30 am	**Breakfast available (Terrace)**	
8:30 am	**Session 4:**	**Leveraging Investments from the Past and Optimizing Future Investments in Applications** **Chair: Katie Lewis**
	Title:	**Leveraging Past and Current Investments for LANL's Future Multi-Physics, Large-Scale Simulation Capability**
	Speaker:	Jerry Brock, *Los Alamos National Laboratory*
	Title:	**How ASCI Changed the Face of Modern Computing**
	Speaker:	Patrick Miller, *Quantlab Financial*
	Title:	**An Applications-Driven Path from Terascale to Exascale with the AMT Uintah Framework**
	Speaker:	Martin Berzins, *SCI Institute/University of Utah*
10:00 am	**Break**	
	Refreshments available (Terrace)	
10:30 am	**Title:**	**DOE's Open Source Leadership in Scientific Visualization**
	Speaker:	Cyrus Harrison, *Lawrence Livermore National Laboratory*
11:00 am	**Panel Discussion**	

Thursday
April 26, 2018

12:00 pm	**Lunch (Council House)**	
1:30 pm	**Session 5:**	**Sharing Knowledge: Cross-Talk Between HPC and the Large-Scale Data Center Community** **Chair: Christoph Junghans**
	Title:	**Building a Production Testbed for Systems Research from Commodity Software**
	Speaker:	Kate Keahey, *Argonne National Laboratory*
	Title:	**Charliecloud Containers for Fun and Profit in HPC**
	Speaker:	Reid Priedhorsky, *Los Alamos National Laboratory*
	Title:	**Can Economical Object Storage Really Replace Parallel File Systems for HPC Workloads?**
	Speaker:	Tony Barbagallo, *Caringo, Inc.*
3:00 pm	**Break**	
	Refreshments available (Terrace)	
3:30 pm	**Title:**	**A University HPC Center Perspective on HPC and Cloud Providers**
	Speaker:	Dan Stanzione, *The University of Texas at Austin*
4:00 pm	**Panel Discussion**	
5:00 pm	**Reception and Informal Discussions (Council House)**	

2018 Salishan Conference on High Speed Computing
SURVEY

WE VALUE YOUR INPUT!

Thank you for participating in this year's conference.
It would be appreciated if you can take a few minutes to complete this brief online survey:

http://salishan.ahsc-nm.org/2018Survey.html

Abstracts

Keynote Address

Computing Challenges for the Third Decade of Stockpile Stewardship

Mark Anderson
National Nuclear Security Administration

The first two decades of Science-Based Stockpile Stewardship (SBSS) were marked by a relative stasis in high-performance computing (HPC) technology. The HPC landscape included heavy cores, abundant memory per core, and the ubiquitous Message Passing Interface. The scale of computing during this period increased four orders of magnitude from O(1) TeraFLOPS to O(10) PetaFLOPS. As we advance toward the Exascale computing era, the technology landscape is far less certain. Lightweight cores, accelerators, less memory per core, multi-tiered message passing schemes, and cognitive computing technologies are complicating the ability to leverage HPC for SBSS. The mission space is also evolving, from a world in which nuclear weapons are playing a stable deterrent role to one riddled with uncertainty. The keynote address will describe the steps the Advanced Simulation and Computing program at the National Nuclear Security Administration is taking to confront these challenges.

Session 1
Maximizing Mission Science per Investment Dollar

HPC and its Role in Maximizing Weapons Program Investments

Chris Clouse
Lawrence Livermore National Laboratory

This talk will discuss the role of high performance computing in the weapons program and how HPC is being leveraged to provide cost savings over historical methods used to accomplish mission science goals. This includes improved efficiency in operations, guidance as to how to spend investment dollars to achieve maximum return and collateral benefits to broader communities. A mix of both qualitative and quantitative arguments will be made.

Maximizing the Impact of DoD HPCMP Investments

Roy Campbell
Department of Defense

This presentation will provide an overview of the DoD High Performance Computing Modernization Program's efforts to increase the impact of its supercomputing investments on critical defense science topics such as armor, electromagnetic railgun, and hypersonic flight. Strategic changes to the HPCMP's investment approach will be discussed: (a) the introduction of fully integrated partitions capable of machine/deep learning, (b) the adoption of job containers, (c) the adoption of modern workflows (e.g. Jupyter notebooks), (d) the specialization of architectures (initially designing systems for either capability or capacity workloads, and later specializing systems across a broader set of workload types), (e) the increased breadth of classifications and compartments, (f) the introduction of reimbursable HPC cloud offerings to offload less impactful work, (g) the revision of vendor penalties (i.e. for late delivery, lack of performance, and unreliability), and (h) the addition of a fifth year production option in light of slowing technology trends.

Session 1
Maximizing Mission Science per Investment Dollar

Can the Tension Between Projectization and Agility/Innovation for Apps and Software Development Motivate Better ROI?

Doug Kothe
Oak Ridge National Laboratory

Modern algorithm, methods, and software development - whether it spans low-level system software to high-level applications development environments, including the software infrastructure to support large-scale data management and data science for end use applications, advanced mathematical libraries and frameworks, extreme-scale programming environments, tools, and visualization libraries – must follow an agile approach that actually welcomes changing requirements, even late in development. Agile approaches, as opposed to what many believe is a now obsolete waterfall model, harness change for the customer's competitive advantage. This is especially true for the R&D nature of application and software development supporting DOE's mission space where the "whats" (functional requirements) may be known but typically the "how" (design requirements) and "how well" (quality requirements) are not known well or at all. Yet how do we foster, support, and protect this model in our current projectized world of "what have you for me lately"? Many of the "agile manifesto" principles have been realized in DOE app and software development over the past few decades, yet the sustained project and programs necessary to feed (fund) this activity fall within rigid and inappropriately matched project management processes. Can a balance be struck with this tension? Should we fight to change it or try to fit our model with it? How well are we at measuring ROI and other critical performance metrics to make the case for what we think is best?

30 Years of HPC from Gigaflops to Exaflops: No Science Application Left Behind or Thinning the Herd?

Steve Plimpton
Sandia National Laboratories

I develop applications and algorithms for large HPC machines. Shockingly, my career now spans 30 years (ok, I'm old). Motivated by our session theme of "maximizing mission science," I will share some observations and numbers as to what has changed over that time frame, where we're headed, and how well I think we're enabling science as an HPC community. Spoiler alert: it's not all good news. All of this will be from an application-centric perspective (really, what else matters for science), so hopefully others will disagree.

Session 2
Removing Constraints, Engaging Challenges, Realizing Opportunity: an Era of New Ecosystems

International Trends in the HPC Market and an Update on the ROI from HPC Investments with a Focus on Chinese, European and Japanese Exascale Plans

Earl Joseph
Hyperion Research

The presentation will show Hyperion Research's latest information on the HPC market around the world, with an analysis of the investments being made in China, Europe and Japan. It will highlight the key driving factors that are causing a growing interest in HPC/supercomputing around the world. It will show the expected Exascale plans for the major countries compared to the US plans. And it will include an update on our research on the ROI from investments in HPC.

Opportunities and Challenges of the COTS IP Model in HPC

Eric Van Hensbergen
arm

A commercial off the shelf intellectual property (COTS IP) model for silicon design such as those provided by Arm Ltd or SiFive have a rich set of options allowing system-on-chip (SoC) designers high degrees of customization at both a component and composition level. These degrees of flexibility include detailed parameters such as number of functional units or cache sizes, coarser grained compositional features such as the number and type of processing cores or accelerators, as well as the opportunity to incorporate custom logic and accelerators with existing available designs. The fact that all of these different components and their variety of configurations are available as modular building blocks significantly lowers the barrier to entry for companies seeking to build custom SoCs either for their own use or to sell as components to others allowing custom configurations to target particular problems or market segments.

However, there are still significant hurdles and costs associated with realizing these designs — particularly for silicon targeting leadership performance, large core counts, high speed peripherals, and in aggressive technology nodes. Compounding this is the incredible importance of the software ecosystem, tools, and compilers for these custom architectures. This talk will explore the opportunities and challenges of realizing the benefits of IP ecosystems based on the past 5 years of ARM's experience in entering the high-performance computing market.

Session 2
Removing Constraints, Engaging Challenges, Realizing Opportunity: an Era of New Ecosystems

Building a Universal Silicon Compiler

Andres Olofsson
DARPA

As the complexity of chips has rapidly increased in line with Moore's law predictions, recent years have seen an explosion in the cost and time require to design advanced SoCs, SiPs, and PCBs. DARPA is addressing these challenges through two new EDA research programs: IDEA (Intelligent Design of Electronic Assets) and POSH (Posh Open Source Hardware). Together, the two programs form the foundation of an intelligent hardware compiler. The aim of the research effort is to create a hardware compiler capable of automatic generation of production ready GDSII directly from source code and schematics. In essence, the program is trying to create the equivalent of a software compiler. Achieving this ambitious goal will require advancing the state of the art in machine learning, optimization algorithms, and expert systems. This session will review the POSH and IDEA EDA research programs and provide analysis of its potential impact on the current semiconductor ecosystem.

Quantum Computing from the Computer Science Perspective: What Will the Software Ecosystem Look Like?

Mark Heiligman
IARPA

A quantum computer is both a physics experiment and a computing engine. While a lot of attention has been paid to the potential of quantum computing to solve classically intractable problems, and a lot of research has looked at the physics of quantum computing and how quantum algorithms will work, somewhat less attention has been paid to how a quantum computer will actually be programmed. This talk will examine the various steps it will take to turn an algorithmic idea on paper into the actual computer controlled physics experiment that executes the algorithm.

The stack starts with a quantum program language that needs to be compiled into idealized quantum operations that then need to be turned into fault tolerant error corrected operations that then need to be turned into actual control signals that are sent to the physical qubits. It is interesting that the rules of quantum mechanics put significant constraints on the high level programming language. A real time classic feedback loop is needed to run the physics experiment with quantum error correction, which can impose additional constraints. Quantum control theory also needs to be incorporated as a key ingredient in order to achieve adequate fidelity of quantum operations for the error correction to work. Estimating the actual resources, both quantum and classical needed to make a quantum computer work is a complicated task that depends on both the underlying technology and algorithm being implemented.

Dinner Speaker

Surveying the Sky with the LSST:
Software as the Instrument of the Next Decade

Andrew Connolly
University of Washington

The development of a new generation of telescopes, large-scale detectors, and computational facilities has led to an era where it is now possible for deep optical surveys to survey a large fraction of the visible sky. One of the largest of these surveys, the Large Synoptic Survey Telescope (LSST), will comprise an 8.4 m primary mirror with a 9.6 square degree field-of-view and a 3.2 Gigapixel camera and begin operations at the end of this decade. Over the ten years of its operation, the LSST will survey half of the sky in six optical colors, discovering 37 billion stars and galaxies and detecting about 10 million variable or transient sources every night. In this talk I will give an introduction to the LSST, its potential for understanding the nature of dark matter and dark energy, for measuring the properties of our Galaxy, and for creating a census of our Solar System. I will show how advances in computational techniques and resources, as well as changes in how we think about data might enable the LSST to achieve and exceed its science objectives.

Session 3
Defining the HPC Ecosystem

The HPC Ecosystem at Amazon Web Services

Brian Barrett
Amazon Web Services

Amazon Web Services provides on-demand cloud computing platforms for a variety of customer use cases, including High Performance Computing. Customers from a variety of backgrounds, including automotive, aerospace, medical, and finance are using AWS for their HPC workloads. AWS's elastic architecture presents a number of challenges and opportunities for customers: it's possible to create large clusters almost instantly, but difficult to achieve the raw performance of an on-premise, purpose built machine. In this talk, we'll talk about AWS-unique capabilities (compute infrastructure, filesystems, etc.) customers are using to run their HPC workloads, how customers are adapting to the elasticity/performance tradeoffs, and investments AWS is making to improve performance and scalability in its compute infrastructure.

What is a Supercomputer and Why Do We Get Confused About How to Define Them?

Bill Kramer
University of Illinois/National Center for Supercomputing Applications

Like other things, and to paraphrase a Chief Justice, we may not know exactly how to define what a supercomputer is but we know one when we see or use it. This talk will discuss what are the sustaining key characteristics of High Performance Computing and Data Analysis Systems (aka Supercomputers) that differentiate them from other computing and data processing technologies. The talk will then discuss the risks we run with using "simple" definitions of what a HPC system is and how narrow definitions can lead us to miss important technology transitions and/or make poor investments. Finally the talk will conclude with key concepts that must be present in our future High Performance Computing and Data Analysis systems in order to keep the nation at the frontiers of science and research, security and economic competitiveness.

Session 3
Defining the HPC Ecosystem

Interesting Times Ahead

Bob Lucas
University of Southern California/Information Sciences Institute

For the past three decades, the HPC community has used commodity components, leveraging billions of dollars of investments designed for other markets to deliver an exponential growth in capability for science and engineering. Where systems are inefficient, their low cost allows one to simply acquire more of them. This strategy will be increasingly less appealing as we look to a post-Moore's Law future. An alternative path forward would be to exploit system-on-a-chip technology to engineer systems designed for specific aspects of the scientific and engineering workload. This talk speculates on what such a path forward might look like, and the constraints upon our imaginations that will be imposed by application software, which today is often more valuable than the ephemeral systems on which it runs.

HPC Workflows – Ecosystem Impact and Challenges

Robert Clay
Sandia National Laboratories

Workflows are now common practice in design and analysis, where parallel simulation runs substitute for physical experiments. Coupling the upstream design models with the discretized analysis models, along with all the requisite model-building tools and post processing analysis tools via modern workflow engines (e.g., SAW at Sandia or ICE at ORNL) is replacing script-based 'workflows'. This allows more complex workflows to be created, run, and shared. Additionally, parametric runs are now the norm, and in general the simulations and analysis codes are run on large-scale parallel computers. This design-through-analysis coupling gives rise to a new set of problems, as we now span computational scale and environments (e.g., a CAD design is run on a Windows workstation, the output of which is used to mesh the problem on a Linux scientific workstation, and the mesh is then used to run a large-scale parametric simulation/analysis on a cluster or supercomputer). This trend of easily organizing tasks and data together into workflows will surely continue, giving rise to a new sort of problem in terms of managing the information and processes. In this talk we will examine the current SotA practice and discuss some of the key R&D challenges we face as we scale out workflows in physics and engineering analysis.

Session 4
Leveraging Invesments from the Past and Optimizing Future Investments in Applications

Leveraging Past and Current Investments for LANL's Future Multi-Physics, Large-Scale Simulation Capability

Jerry Brock
Los Alamos National Laboratory

Early in the nuclear era, the US relied mainly upon theory and experiment to design, demonstrate and deploy with confidence its nuclear weapons stockpile. From the Manhattan Project to the end of the 20th century, simulation science played an important, but limited role in shaping the US stockpile. At the cessation of the integrated-testing era in the early 1990s, the US elevated the role of simulation to complement theoretical and experimental sciences in support of its nuclear security enterprise. The ASCI/ASC Program was the manifestation of this new investment in simulation science supporting DOE/NNSA, providing novel computing platforms together with multi-physics simulation tools. Today, the ASC Program is challenged to retain the physics fidelity and modeling approaches embodied in the existing multi-physics codes in a new generation of simulation tools. Moreover, these new tools must be adaptable and performant on the many novel architectures expected during the ongoing evolution of high performance computing platforms. This presentation will discuss the challenges in leveraging past and current investments for LANL's future multi-physics, large-scale simulation capability.

How ASCI Changed the Face of Modern Computing

Patrick Miller
Quantlab Financial

The purpose of ASCI/ASC was to demonstrate that modeling incredibly complex systems was not just feasible, but practical and effective. The program provided a willing "first-mover" that let vendors build systems at (then) unprecedented scales for computation, storage, and networking. Those components were key to the program's success, but they had interesting side-effects. The ground work done made it tractable for non-government users to pick up the tools and solve their own problems. Clusters became commonplace rather than something used only by fringe groups. More importantly, high performance computing and the associated techniques became accepted and expected practice for all forms of mathematical thinking. HPC ideas, techniques, software, and hardware started showing up everywhere! Resource management and storage formats/models/systems/hardware flowed from vendors to high end users outside the government labs. GPGPU became so commonplace that it dropped the GP. People and ideas streamed out from the Labs to ignite innovation up and down the tech stack. I'll talk about how I've seen industry transition from a 1980's style throughput computing model to a real HPC model and discuss some of the consequences for the Labs as these new "commodity" HPC actors come of age.

Session 4
Leveraging Invesments from the Past and Optimizing Future Investments in Applications

An Applications-Driven Path from Terascale to Exascale with the AMT Uintah Framework

Martin Berzins
SCI Institute/University of Utah

The Uintah software was designed in the first ASCI program at the University of Utah. The asynchronous many task-based approach that it now employs was inspired by the Utah SCIRun framework and by the potential for such architectures at petascale and beyond. Through building upon the novel design, it has been possible to port broad classes of applications from the original terascale machine through to today's petascale architectures.

At the same time algorithmic work related to a series of challenging driving problems has made it possible to develop underlying methods such as the Material Point Method, block-structured AMR, and the radiation solvers of the Arches combustion module. Additionally, these methods are applied to a variety of challenging and very different applications than those that motivated the original work, including an exascale-ready industrial design problem tackled by the Utah PSAAP2 Center.

Nevertheless, the move to new and different architectures requires that the approach used has not only the ability to execute tasks asynchronously but to deal with memory hierarchies and the issue of portability from standard i86 architectures to GPUs and to a broad range of other possible architectures. The approach based upon the Kokkos portability library that enables this process in Uintah will be described. The challenge of porting Uintah to a very different sort of architecture is described through a feasibility study on the Sunway TaihuLight. Lessons for porting to future exascale architectures will be considered.

Session 4
Leveraging Invesments from the Past and Optimizing Future Investments in Applications

DOE's Open Source Leadership in Scientific Visualization

Cyrus Harrison
Lawrence Livermore National Laboratory

To meet the challenge of digesting and understanding massive data sets from HPC physics simulations, DOE's Advanced Simulation and Computing Initiative invested heavily in scientific visualization. These investments gave rise to two widely used and vibrant open source projects: VisIt and ParaView, each with large users bases and impact beyond DOE. This talk will provide perspectives of the trials and successes of the HPC scientific visualization community and how this community is working together to address the challenges from the evolving HPC landscape. It will also dive into how perceptions and use of open source solutions have evolved over the lifetime of these tools.

Session 5
Sharing Knowledge: Cross-Talk Between HPC and the Large-Scale Data Center Community

Charliecloud Containers for Fun and Profit in HPC

Reid Priedhorsky
Los Alamos National Laboratory

Supercomputing centers are seeing increasing demand for *user-defined software stacks* (UDSS), instead of or in addition to the stack provided by the center. These UDSS support user needs such as complex dependencies or build requirements, externally required configurations, portability, and consistency. The challenge for centers is to provide these services in a usable manner while minimizing the risks: security, support burden, missing functionality, and performance.

This talk presents Charliecloud, which uses the Linux user and mount namespaces to run containers, including industry-standard Docker containers, with no privileged operations or daemons on center resources. Our simple, open-source approach avoids most security risks while maintaining access to the performance and functionality already on offer, doing so in just 900 lines of code. Charliecloud promises to bring an industry-standard UDSS user workflow to existing, minimally altered HPC resources.

A University HPC Center Perspective on HPC and Cloud Providers

Dan Stanzione
The University of Texas at Austin

HPC providers and Cloud providers set out to solve very different problems, but in the end the solutions (so far) have a great deal in common – occasionally too much to keep track of which is for what. Nevertheless, there is enough in common that many technologies have "crossed over" from on to the other. Most valuably, the Cloud has provided, at least in the public consciousness, an "As-a-Service" concept that can translate well into many scientific domains. In this talk, I will give the perspective from the Texas Advanced Computing Center on what we've learned from the cloud, what we use the cloud for, what we still think they need to learn from us, and where we think things might go from here.

Session 5
Sharing Knowledge: Cross-Talk Between HPC and the Large-Scale Data Center Community

Building a Production Testbed for Systems Research From Commodity Software

Kate Keahey
Argonne National Laboratory

Computer Science experimental platforms have significantly different requirements than what a traditional scientific datacenter provides today. Systems researchers work on projects ranging from developing new operating systems, experimenting with programmable networks, or testing new security solutions. This type of work cannot be advanced by submitting jobs to a scheduler. Instead, it requires support for a set of much broader and complex set of capabilities including full stack control – modifications to operating system kernel, firmware, or BIOS – as well as isolation, to prevent the experiments from one user impacting others. Those capabilities need to be implemented on a range of diverse hardware – from commodity processors at scale to FPGAs – to support a wide range of experiments. This provokes many trade-offs: complexity versus control trade-offs to accommodate beginner as well as advanced users, or innovation versus maintenance trade-offs to strike a balance between supporting the broadest possible range of experiments on one hand but keep the maintenance cost reasonable.

In this talk, I will describe the choices, adventures, and lessons learned in building the Chameleon testbed (chameleoncloud.org) that leverages a commodity technology, the OpenStack system, to provide a Computer Science testbed for systems experimentation. I will discuss the overall development strategy as well as specific problems and trade-offs we faced and how we addressed them. I will also give examples of research projects currently running on Chameleon and how they use the capabilities it offers. Finally, I will discuss the path forward including broadening our capabilities, packaging our approach, and providing support for repeatability.

Session 5
Sharing Knowledge: Cross-Talk Between HPC and the Large-Scale Data Center Community

Can Economical Object Storage Really Replace Parallel File Systems for HPC Workloads?

Tony Barbagallo
Caringo, Inc.

We've heard that object storage can be a cost-effective archive repository for original source data as well as results from HPC analysis, but could object storage actually be leveraged as the primary data warehouse for active HPC initiatives?

In this presentation we'll explore two vastly different real-world examples. The second one may be a bit surprising: With the right network setup and software performance characteristics, is it possible to replace parallel file systems for certain HPC use cases?