The Salishan Conference on

HIGH SPEED COMPUTING

Algorithms Architecture Applications

April 26 – 30, 2021

Virtual Conference
All Pacific Times

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: Entering an Era of Extreme Heterogeneity. 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 contact any of the conference committee members and/or complete the electronic survey at the end of the conference **(https://salishan.ahsc-nm.org/2021Survey.html)**.

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

Conference Committee: Olga Pearce and Ian Karlin, *LLNL*
Ron Brightwell and Simon Hammond, *SNL*
Carolyn Connor and Christoph Junghans, *LANL*

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

Next Conference Dates:	April 25 - 28, 2022	April 24 - 27, 2023	April 22 - 25, 2024

Table of Contents

Sponsorship

The Salishan Conference on High Speed Computing is administered, hosted, and managed by The Association for High Speed Computing (AHSC). We gratefully acknowledge the generous support of all our corporate sponsors over the past 40 years, including the following:

Amazon

Advanced Micro Devices, Inc.

Appro International, Inc.

ARM

Cerebras

Cisco Systems, Inc.

ClearSpeed Technologies, Inc.

Convey Computer Corporation

Cray, Inc.

DataDirect Networks, Inc.

DDN Storage

Dell, Inc.

Dell/EMC

Dell Technologies

D-Wave Systems, Inc.

EMC

Google

Hewlett-Packard Company, Inc.

Hewlett Packard Enterprise

IBM Corporation

Intel Corporation

Sponsorship

LexisNexis Special Services, Inc.

Linux Networx, Inc.

Marvell

Mellanox Technologies

Micron Technology, Inc.

Microsoft Corporation

Netezza Corporation

NVIDIA Corporation

Panasas, Inc.

Penguin Computing

PGI Compilers & Tools

Reservoir Labs, Inc.

Rogue Wave Software, Inc.

Seagate Government Solutions

Silicon Graphics International, Inc.

Sun Microsystems, Inc.

SUSE

The Portland Group, Inc.

TotalView Technologies

University of California: Lawrence Livermore National Laboratory

University of California: Los Alamos National Laboratory

Conference Theme

Strong Legacy, Bold Future

A Celebration of High Speed Computing on Salishan's 40th Anniversary

The story of computing and mission-science has arguably been one of conjoined innovation. It is a story of a continuous, at times integrated and at times turbulent, technology-push and demand-pull dynamic in which science problems have driven computing advancement and computing has advanced science capability. The story is set within a complex ecosystem that includes a rarified collection of performance-driven users, technology and system providers and innovators, and the challenges of a highly unusual market, with each actor navigating the intricate balance between innovation on one hand with system complexity, technical feasibility, and usability on the other. On this 40th anniversary of the Salishan Conference on High Speed Computing, it is fitting to look back on the road we have traveled, to consider knowledge gained and lessons learned, and to celebrate notable achievements, while also gazing forward toward a future filled with grand challenges, technological promise, and scientific advancement.

In the summer of 1976, Nic Metropolis and Jack Worlton organized the historic International Research Conference on the History of Computing, inviting computing pioneers worldwide. The conference covered the evolution of computers from mechanical desk calculators to the 1976 recent past and resulted in seminal conference proceedings, as well as several books, including a famous collection of essays published in 1980. As these pioneers examined the legacy of computing, they also sought to shape its future.

In the late 1970s, the composition of the broader HPC Community was changing as an increasingly diverse set of HPC users (beyond the traditional users found mostly in government) began to buy supercomputers. The rate of performance improvement was slowing. While the US HPC industry appeared to be moving toward evolutionary rather than revolutionary platform development, other countries, particularly Japan, seemed to be focusing more heavily on massively parallel development. It is against this backdrop, and with a similar vision as the 1976 Metropolis Conference, that the Salishan Conference was established in 1981 as a means of convening experts in computer architecture, languages, and algorithms to facilitate discussion and develop strategies to ensure that the long-term needs of government users could be met. The targeted goals of the conference were to improve communications, develop collaborations, solve problems of mutual interest, and provide effective leadership in the field of high speed computing. Over the years these areas continued to be important, and although languages still played a major role in HPC, the conference logo was updated to reflect the applications area as a major aspect of the conference along with architectures and algorithms.

Conference Theme

Over the last four decades, the annual Salishan Conference on High Speed Computing has convened thought leaders of the HPC community, as well as influential policy makers and program managers. Unlike traditional academic conferences that bring together those directly engaged in research, Salishan attendees represent the broader community including experts and luminaries from government labs and agencies, industry, and academia to nurture a unique experience in which all facets of the HPC ecosystem are considered, including market viability and the political considerations of sustaining long-term investment in approaches and technologies. The invitation-only structure of Salishan ensured that the diversity of viewpoints and perspectives, as well as cutting edge information within the community could be represented, considered and developed. As a result of its location, unique structure, and opportunity for access to peer thought leaders, as well as speakers and attendees chosen for their ability to influence and impact the HPC community, Salishan asserted its position as one of the most sought-after invitations by HPC community leaders. In 2021, due to the global pandemic, we celebrate the 40th anniversary and continue the Salishan tradition with a virtual offering of seven conference sessions covering a broad range of topics. Thank you for joining us. We hope that you enjoy the conference.

Session 1: Platforms: Achievement Propelling Opportunity

Computer systems form the bedrock on which scientific computing innovation happens. Their ever-expanding capabilities spur scientific discovery. Their evolving architectures keep algorithm, application and programming model researchers busy adapting to new hardware paradigms. This session brings together experts who helped shape our platforms by measuring their capabilities, ranking them, and spurring efficient use through the most prestigious HPC performance award. It also brings together a look back at the history of key platforms of the last 40 years and the ASCI program, which procured many number one systems.

Session 2: Software Innovation: Building Blocks for Success

The software stack and associated tools are responsible for transforming a massive collection of interconnected complex hardware components into a highly effective resource for developing and executing parallel science and engineering applications. The HPC community has been on a never-ending quest to provide software components and application development tools that maximize performance and scalability on an ever-expanding set of hardware, while simultaneously minimizing the impact on the programmer. This session examines several aspects of the software stack and application development environment, including the challenge of mapping application parallelism to machine parallelism, optimizing data layout and data representation, novel methods of increasing programmability using non-traditional HPC languages, and an approach to the software life cycle aimed at improving the quality of scientific software for the US DOE Exascale Computing Project (ECP).

Conference Theme

Session 3: Emerging Directions in High Speed Computing

Salishan has been a conference that has witnessed both change and been part of making change happen. Areas where progress has been most keenly felt has been in the size and scale of our high speed computers and in the design of their architectures. Cloud-based systems, custom-silicon, open source instruction sets and more point to a continued and exciting future for our community. This session brings a collection of industry and lab experience in this domain to present novel directions for the future of our computing platforms. Many of these topics will become the mainstay of future Salishan presentations and themes, so we welcome discussion on how these emerging directions can be harnessed to deliver the exceptional performance required by our algorithm, system-software and application users.

Session 4: Simulations: Code Innovation for High Performance Computing

Over the last decades, innovations in code have made it possible to study problems in the mulitphysics domain that were previously unthinkable to simulate. Improvements on all algorithmic, as well as computer science, levels were needed to achieve this. In this session these innovations, ranging from new multiphysics developments over novel parallel algorithms for radiation transport to using machine learning approaches to bridge scale for materials will be explored. Additionally, the interplay between these developments and performance portability will be discussed.

Session 5: Applications: From Machines to Mission Science

Computational science is arguably the interface where pressing science meets computing platform. Thus, this session will further explore the technology-push and demand-pull relationship between computing and mission science through the lens of the application. We will consider both the aspirations of science-based prediction and the experience gained through the national investment in the Stockpile Stewardship Program with an eye toward the future, by exploring a range of topics including the sustainability of science-based prediction in the nuclear security mission, the benefits and challenges for concurrent V&V (verification and validation) development, key lessons learned from pre-ASC and ASC era code development, and where discovery science with respect to multiphysics applications is headed.

Conference Theme

Session 6: Exascale and Co-Design

As the continents race to exascale many strategies have arisen, with Fugaku, Japan, casting its net wide, aiming for applicability with the 7nm ARM CPU at its core. The EU is also targeting a broad set of applications, aiming to deliver several energy-efficient exascale systems. Meanwhile, the US focuses on power requirements, reliability, billion-way parallelism, and data movement for its own upcoming GPU- and CPU-based exascale machines. All of the systems have been co-designed with the applications they aim to enable to drive scientific discovery. Meanwhile, Quantum computing is closer than ever to becoming reality.

Session 7: Reflections on Salishan's Impact Past and Future

Please join us for a fireside chat reminiscent of Salishan and the Northwest coast. We conclude our 2021 conference with an opportunity for the Salishan community to share memories with friends and colleagues and to celebrate our 40th Anniversary. Our panel will discuss a wide range of topics, including how and why the conference was established; the important role that the annual convening of thought leaders at the Salishan Conference played over the years in providing a venue for creating opportunity and addressing HPC community concerns, both technical and non-technical; and how the conference has evolved over time, with ample time for audience Q&A, participation, and discussion.

Monday – All Pacific Times
April 26, 2021

7:30 – 8:15 am	**Welcome/Keynote Address**	
	Title:	**Salishan (1:20): The Transformative Years of the HPC Revolution**
	Speaker:	Gordon Bell *Microsoft*
	Session 1:	**Platforms: Achievement Propelling Opportunity**
	Chairs:	**Ian Karlin and Carolyn Connor**
8:15 – 8:45 am	**Title:**	**HPC: Where we are Today and a Look into the Future**
	Speaker:	Jack Dongarra *University of Tennessee*
8:45 – 9:00 am	**Break**	
9:00 – 9:30 am	**Title:**	**The Causes and Origins of ASCI**
	Speaker:	Dona Crawford *Livermore Lab Foundation*
9:30 – 10:00 am	**Title:**	**Accelerating the Future: 40 Years of Tri-Lab Platform Development**
	Speaker:	Nic Lewis *Los Alamos National Laboratory*
10:00 – 10:30 am	**Panel Discussion**	
10:30 – 10:45	**Break**	

Monday – All Pacific Times
April 26, 2021

	Session 2:	**Software Innovation: Building Blocks for Success**
	Chairs:	**Ron Brightwell and Si Hammond**
10:45 – 11:15 am	**Title:**	**Programming at Extreme Scales**
	Speaker:	Bill Carlson *IDA Center for Computing Sciences*
11:15 – 11:45 am	**Title:**	**How Programming Systems Meet the Needs of New Architecture Classes - Past, Present and Future**
	Speaker:	Mary Hall *University of Utah*
11:45 – 12:15 pm	**Title:**	**E4S: Toward an Ecosystem for HPC R&D Software**
	Speaker:	Mike Heroux *Sandia National Laboratories*
12:15 – 12:30 pm	**Break**	
12:30 – 1:00 pm	**Title:**	**Better Tools for Better P3**
	Speaker:	Michael Wolfe *NVIDIA*
1:00 – 1:30 PM	**Panel Discussion**	

Tuesday – All Pacific Times
April 27, 2021

	Session 3:	**Emerging Directions in High Speed Computing**
	Chairs:	**Ian Karlin and Si Hammond**
7:30 – 8:15 am	**Title:**	**Repealing "Amdahl's" Law**
	Speaker:	Danny Hillis *Applied Invention*
8:15 – 8:45 am	**Title:**	**Riding the Deep Learning Tsunami**
	Speaker:	Garth Gibson *Vector Institute*
8:45 – 9:15 am	**Title:**	**Accelerating HPC Flywheel: Lessons from Machine Learning**
	Speaker:	Nafea Bshara *Amazon Web Services*
9:15 – 9:30 am	**Break**	
9:30 – 10:00 am	**Title:**	**RISC-V: The Open Era of Computing**
	Speaker:	Calista Redmond *RISC-V International*
10:00 – 10:30 am	**Panel Discussion**	

Wednesday – All Pacific Times
April 28, 2021

	Session 4:	**Simulations: Code Innovation for High Performance Computing**
	Chairs:	**Christoph Junghans and Olga Pearce**
7:30 – 8:00 am	**Title:**	**Mission Driven Innovation: A Case Study of Multiphysics Code Development**
	Speaker:	Rob Rieben *Lawrence Livermore National Laboratory*
8:00 – 8:30 am	**Title:**	**Parallel Sn Algorithms**
	Speaker:	Jim Morel *Texas A&M University*
8:30 – 9:00 am	**Title:**	**Bridging Scales of Materials Modeling Using Machine Learning**
	Speaker:	Gowri Srinivasan *Los Alamos National Laboratory*
9:00 – 9:15 am	**Break**	
9:15 – 9:45 am	**Title:**	**Performance Portable Applications - Is Portability Even a Possibility?**
	Speaker:	John Levesque *Hewlett Packard Enterprise*
9:45 – 10:15 am	**Panel Discussion**	

Thursday – All Pacific Times
April 29, 2021

	Session 5:	**Applications: From Machines to Mission Science**
	Chairs:	**Carolyn Connor and Christoph Junghans**
7:30 – 8:00 am	**Title:**	**Is the Role of Science-Based Prediction Sustainable in the Nuclear Security Mission?**
	Speaker:	Raymond Juzaitis, Retired *Lawrence Livermore National Laboratory* *Los Alamos National Laboratory*
8:00 – 8:30 am	**Title:**	**A Short History of Software Development During LLNL's ASC Program**
	Speaker:	Teresa Bailey *Lawrence Livermore National Laboratory*
8:30 – 9:00 am	**Title:**	**Concurrent V&V/UQ and Code Capability Development: Hypersonic Reentry Analysis with SPARC**
	Speaker:	Sarah Kieweg *Sandia National Laboratories*
9:00 – 9:15 am	**Break**	
9:15 – 9:45 am	**Title:**	**A Look at Trends in Computational Methods: The Tie Between Machines and Mission Science**
	Speaker:	Aimee Hungerford *Los Alamos National Laboratory*
9:45 – 10:15 am	**Panel Discussion**	
10:15 – 10:30 am	**Break**	

Thursday – All Pacific Times
April 29, 2021

	Session 6:	**Exascale and Co-Design**
	Chairs:	**Olga Pearce and Ron Brightwell**
10:30 – 11:00 am	**Title:**	**History and Solutions to the Four Exascale Challenges**
	Speaker:	Al Geist *Oak Ridge National Laboratory*
11:00 – 11:30 am	**Title:**	**Towards an Integrated, Dynamic and Resource Aware Software Stack for the EU Exascale Systems**
	Speaker:	Martin Schulz *Technical University of Munich*
11:30 – 12:00 pm	**Title:**	**Fugaku: The First Exascale Supercomputer and its Innovative Arm A64FX CPU**
	Speaker:	Satoshi Matsuoka *RIKEN*
12:00 – 12:15 PM	**Break**	
12:15 – 12:45 pm	**Title:**	**Quantum Computing - Why, Where and When**
	Speaker:	Robert (Bo) Ewald *ColdQuanta*
12:45 – 1:15 pm	**Panel Discussion**	

Friday – All Pacific Times
April 30, 2021

	Session 7:	**Reflections on Salishan's Impact Past and Future**
	Moderator:	**Dona Crawford**
7:30 – 7:40 am	**Title:**	**Conference History**
	Speaker:	**Dona Crawford** *Livermore Lab Foundation*

Panel Members

Robert (Bo) Ewald
ColdQuanta

Manuel Vigil, Retired
Los Alamos National Laboratory

Jim Ang
Pacific Northwest National Laboratory

Don Dossa, Retired
Lawrence Livermore National Laboratory

Don Tolmie, Retired
Los Alamos National Laboratory

7:40 – 9:30 am Panel Discussion

2021 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/2021Survey.html

Abstracts

Keynote Address

Salishan (1:20) The Transformative Years of the HPC Revolution

Gordon Bell
Microsoft

Salishan (1) began with the mono-memory single processor Cray 1 vector FORTAN computer c1960-1980. The developments during Salishan (1:20) helped enable programs to run across millions of independent, coupled computers that now operate at o(trillion) megaflops will be recalled. Besides the construction of such computers, the effort to exploit multicomputer systems has turned out to be quite challenging. In 1987, as the founding head of the NSF computing directorate, CISE, and believing in the difficulty, I offered an annual prize to acknowledge the efforts…that has hopefully been useful to acknowledge the extraordinary progress.

HPC: Where We Are Today and a Look into the Future

Jack Dongarra
University of Tennessee

In this talk, we examine how high performance computing has changed over the last 40 years and look toward the future in terms of trends. These changes have had, and will continue to have, a significant impact on our numerical scientific software. A new generation of software libraries and algorithms are needed to effectively use (wide area) dynamic, distributed, and parallel environments. Some of the software and algorithm challenges have already been encountered, such as management of communication and memory hierarchies through a combination of compile-time and run-time techniques, but the increased scale of computation, depth of memory hierarchies, range of latencies, and increased run-time environment variability will make these problems much harder.

Session 1
Platforms: Achievement Propelling Opportunity

The Causes and Origins of ASCI

Dona Crawford
Livermore Lab Foundation

With the end of the cold war and the associated peace dividend, DOE budgets for the nuclear weapons laboratories were in sharp decline. The Galvin Commission was looking to close a laboratory or two and the Clinton administration stated unequivocally there would be no more underground nuclear tests. This was the backdrop against which ASCI was formulated and developed. This talk will explore the conditions and players that led to the creation of ASCI and the initial years of getting it off the ground.

Accelerating the Future: 40 Years of Tri-Lab Platform Development

Nic Lewis
Los Alamos National Laboratory

The Department of Energy's (DOE) Advanced Simulation and Computing Program (ASC) is now over 25 years old. Founded in 1995, ASC was a response to the demands that a rapidly changing geopolitical and technical environment placed on the core mission of the DOE weapons labs — Los Alamos, Sandia, and Lawrence Livermore. The conjoined end of the Cold War and beginning of a moratorium on full-scale nuclear testing in the US meant that the weapons labs could no longer rely on their traditional means of certifying the performance and safety of the nation's nuclear stockpile. Much of that crucial process had to move to the *virtual* laboratory of computer simulations. In the mid-1990s, many hurdles stood in the way of producing simulations at the extreme levels of detail and accuracy needed by the DOE labs to care for an aging stockpile. Those hurdles included the limitations of existing supercomputer platforms, and the disarray of the computer industry that produced them. This talk describes some of the ASC program's most impactful accomplishments in the advancement of HPC platform development, and its role in the transition from the vector supercomputer era to the dawning of exascale in less than three decades.

Session 2
Software Innovation: Building Blocks for Success

Programming at Extreme Scales

William Carlson
IDA Center for Computing Sciences

As extreme scale platforms, both for HPC and Cloud Computing, continue the decade-long shift towards an architectural configuration that greatly favors local, small memory footprint computation, it is time to consider spending a small amount of these excellent local compute and memory resources on efforts to increase programmability and provide for a wider programmer base. This is especially the case for computations which inherently require either significant non-local interaction and/or large memory footprints. This talk will highlight the work of the speaker and others on computational techniques using novel approaches to remote data manipulation, such as Actors and Conveyors; using techniques that employ aggregation and asynchrony, in programming languages such as JavaScript and Rust.

How Programming Systems Meet the Needs of New Architecture Classes: Past, Present and Future

Mary Hall
University of Utah

According to Bell's Law, a new computer architecture class forms roughly each decade (cheaper, smaller, more capable) and establishes a new industry. In this talk, we first review the technology trends and resulting architectures from each decade, from the 1980s to the present, and how programming systems met the needs of these architectures. We make the case that the current class of architectures demand increased focus on *reducing data movement.* How data is laid out in memory and representations that compress data (e.g., reduced floating point precision) have a profound impact on data movement. Moreover, the cost of data movement in a program is architecture-specific, and consequently, optimizing data layout and data representation must be performed by a compiler once the target architecture is known. With this context in mind, this talk will provide examples of data layout and data representation optimizations and call for integrating these data properties into code generation and optimization systems.

Session 2
Software Innovation: Building Blocks for Success

E4S: Toward an Ecosystem for HPC R&D Software

Mike Heroux
Sandia National Laboratories

Open source, community-developed reusable scientific software represents a large and growing body of capabilities. Linux distributions, vendor software stacks and individual disciplined software product teams provide the scientific computing community with usable holistic software environments containing core open source software components. At the same time, new software capabilities make it into these distributions in a largely *ad hoc* fashion.

The Extreme-scale Scientific Software Stack (E4S, https://e4s.io), first announced in November 2018, along with its community-organized scientific software development kits (SDKs), is a new community effort to create lightweight cross-team coordination of scientific software development, delivery and deployment and a set of support tools and processes targeted at improving scientific software quality via improved practices, policy, testing and coordination.

In this presentation, we introduce E4S, discuss its design and implementation goals and show examples of success and challenges so far, including its important role in the US Exascale Computing Project. We will also discuss our connection with other key community efforts we rely upon for our success and describe how collaboration around E4S can be realized.

Better Tools for Better P3

Michael Wolfe
NVIDIA

The P3 (performance, portability, productivity) problem has faced high performance computing for decades. The perfect solution would be a single version of an application that gives high performance across a wide range of target systems and is easy to develop and maintain. Actual solutions give up some level of performance, or portability, or productivity, or all three. Here we review three points in the past 65 years when the P3 problem had good solutions, even when introducing new architectural styles. Today's problems are more challenging, as the number of dimensions of parallelism keeps increasing and the importance and values for each dimension vary across different systems. I argue that the path to better P3 solutions requires good tools to help identify parallelism and, just as important, to map that to the appropriate dimension of machine parallelism.

Session 3
Emerging Directions in High Speed Computing

Repealing "Amdahl's Law"

Danny Hillis
Applied Invention

In the 1980s, most supercomputers had only a small number of processors because supercomputer users and designers believed that increasing the number of processors would yield diminishing returns in performance. This myth was not widely considered to be a matter of opinion, but a mathematical certainty, often referred to as "Amdahl's Law." I will discuss the history of "Amdahl's Law," what was right about it, why it was misleading, and how it delayed the widespread adoption of highly parallel computers.

Riding the Deep Learning Tsunami

Garth Gibson
Vector Institute

A little more than a dozen years ago Roadrunner heralded the Petascale era and the Graphics Processing Unit acceleration era. Less than 10 years ago, AlexNet, a Convolutional Neural Network implemented on a gamer's GPU, nearly halved the error rate for automated image classification by deepening the network to explore far more potential data features automatically, launching the Deep Learning era. Deep Learning's ability to search a large space of possible data features to find an accurate model fitting data has transformed the lives of much of the world's population. Together, HPC, GPUs and Deep Learning are building models of human communication that, for many practical purposes, understand questions and discussions and provide useful answers and interactions. As an operating system software designer, drawn into HPC two decades ago by the challenges of scaling storage systems to peta-, then exa-scale, then drawn into the distributed system challenges of scaling machine learning a decade ago, I am now drawn into the human scaling challenges of accelerating the creation of new machine learning and Deep Learning technologies, accelerating the training of the human workforce skilled in AI and accelerating the adoption, deployment and management of AI technologies by industry and society. This talk will explore emerging trends and technologies from the perspective of a research lab full of Deep Learning innovators whose statistical and modeling skills far exceed their software engineering training.

Session 3
Emerging Directions in High Speed Computing

Accelerating HPC Flywheel: Lessons from Machine Learning

Nafea Bshara
Amazon Web Services

Machine learning, as a segment of HPC, has enjoyed exponential growth in business adoption - in numbers of practitioners, and in hardware and software investments. This has resulted in a flywheel effect: wider business adoption drove more engineers to join the field. More engineers and more business brought more investment. More investment created more power-efficient and cost-effective products, which meant even more business adoption. In this talk we will describe how this machine learning flywheel evolved, sharing a unique perspective from Amazon. We are a business that adopted machine learning and helped democratize it through investments in silicon, hardware, compilers, and service model. We will finish with some insights about how this could help speed up the flywheel for other HPC cases.

RISC-V: The Open Era of Computing

Calista Redmond
RISC-V International

RISC-V is a free and open RISC Instruction Set Architecture (ISA) enabling a new era of processor innovation through open standard collaboration. RISC-V is growing rapidly around the world from smart watches to national architectures, and mobile to high performance computing. Across multiple domains and in areas of fast growth, we are seeing innovation and collaboration in the open. RISC-V has brought a new, unrestricted approach to processors that accelerates and opens innovation while reducing necessary investment. RISC-V ISA delivers a new level of free, extensible software and hardware freedom on architecture, paving the way for the open era of computing. This talk will highlight the strategic importance, global momentum, and tremendous opportunity ahead.

Session 4
Simulations: Code Innovation for High Performance Computing

Mission Driven Innovation: A Case Study in Multiphysics Code Development

Rob Rieben
Lawrence Livermore National Laboratory

LLNL is developing next generation multiphysics simulation capabilities for national security applications under the Multiphysics on Advanced Platforms Project (MAPP). MAPP is part of the overall tri-lab ATDM NextGen code development effort taking place at LLNL, LANL, and SNL to meet the challenge posed by exascale computing. Developing multiphysics codes capable of meeting the various simulation needs of the NNSA, supporting robust and accurate simulation workflows with increased automation, enabling innovations such as ML and AI to be easily explored and incorporated, all while scaling to pre-exascale and exascale computers is a daunting task. To accomplish this ambitious goal, we have embraced two key themes: use of high-order numerical methods and a modular approach to code development. Each of these core pieces of technology were originated under previous R+D efforts at LLNL, requiring a careful transition of research-based methods and code into a production environment. In this talk, we take a historical perspective on the development of MAPP, starting from the foundational research at low technology readiness levels and how these were integrated and built upon over time to form the complex, integrated set of simulation capabilities that we have successfully deployed at scale to users. In addition, the LLNL effort was heavily influenced by the lessons learned from the ASCI program [PostKendall2004], launched in 1996, which brought forth the massively parallel simulation era and today's petascale class production codes. We review some of the key lessons from this era which have proven true for the ATDM era as well. Finally, we reflect on our vision of how these new codes in MAPP will enable a more responsive and insightful experience for our users.

Session 4
Simulations: Code Innovation for High Performance Computing

Parallel SN Algorithms

Jim Morel
Texas A&M University

We give a brief overview of parallel Sn solution algorithms for both structured and unstructured meshes. In the early days of parallel computing, solution algorithms for the Sn (discrete-ordinates) equations of radiation transport were often identified as being inherently serial. There is always a set of steps in the solution process that must be carried out sequentially, but in spatially multidimensional calculations, there is parallelism within those steps that can be exploited. On 2-D rectangular meshes, this parallelism maps to a one-dimensional domain, and on 3-D rectangular meshes, it maps to a two-dimensional domain. This unusual property demands special domain decompositions for optimal performance. Algorithms for rectangular meshes are mature and well understood, but algorithms for unstructured meshes remain a very active area of research. There are unstructured-mesh algorithms that rely on the "cutting" of standard finite-element meshes to obtain polyhedral meshes with non-reentrant processor domains. The domains are rectangular, but the meshes within each domain are unstructured. The equations can then be solved in parallel on such meshes using rectangular mesh algorithms. Special forms of mesh cutting can also be used to load balance, but perfect load balancing can actually be suboptimal due to the introduction of Sn-specific inefficiencies in the solution process. A current focus area is unstructured-mesh algorithms that accommodate standard Par Metis domain decompositions.

Session 4
Simulations: Code Innovation for High Performance Computing

Bridging Scales for Materials Modeling using Machine Learning

Gowri Srinivasan
Los Alamos National Laboratory

Machine learning methodologies are being used increasingly in many branches of science. One key advantage is the ability of machine learning algorithms to learn and replicate very complex physical behavior once properly trained. This feature is particularly useful in characterizing materials at the continuum scale, where several observable phenomena live. Properties of materials at lower length scales play key roles in governing the dominant physics in such systems of interest but can only be known statistically. Current models either ignore or idealize sub-macroscale information at these larger scales because we lack a framework that efficiently utilizes it in its entirety to predict macroscale behavior. We propose a method that integrates computational physics, machine learning and graph theory to make a paradigm shift from computationally intensive high-fidelity models to machine learning emulators that preserve accuracy yet speed up calculations by orders of magnitude.

Session 4
Simulations: Code Innovation for High Performance Computing

Performance Portable Applications - Is Portability Even a Possibility?

John Levesque
Hewlett Packard Enterprise

Portability across different vendors today is a tremendous headache for application developers. We do have standards in C, Fortran, C++, OpenMP and OpenACC; however, standards do not seem to help. One of the biggest issues with the Cray Programming Environment (CPE) is that we strictly adhere to the Fortran 2018 standard, the C11 standard and C++17. The xRage application at Los Alamos National Laboratory has been developed with the Intel compiler and tested on the GNU compiler. When we ported that application to CCE, we found numerous non-standard Fortran constructs in both the Intel and GNU compilers. If compilers extend the language on their own, why do we have standards?

Another huge issue is support of both OpenMP and OpenACC is fraught with compilers not supporting the latest version of the standard. Nvidia has stated that they will not support OpenMP beyond version 4.5. This means that an OpenMP offload application running on Frontier would not be able to be compiled with the Nvidia OpenMP compiler. Cray currently supports OpenACC 2.0 and the OpenACC standard is now at 2.7.CPE cannot today compile applications that used features beyond 2.0. What is the solution?

HPE is investigating extending CPE to:

1) Accept non-standard features of Intel and GNU compilers where possible – some have no clear solution. For example, argument mismatch is not allowed in LLVM.
2) Accept compiler options used by Intel and GNU where possible.
3) Pick up the development of OpenACC to support version 2.7 and beyond.
4) Support ARM, Intel, AMD, Nvidia.

At least one compiler could be portable across all available platforms.

Session 5
Applications: From Machines to Mission Science

Is the Role of Science-Based Prediction Sustainable in the Nuclear Security Mission?

Raymond Juzaitis
Retired, Lawrence Livermore National Laboratory,
Los Alamos National Laboratory

Over 25 years ago in 1995, the Stockpile Stewardship Program was chartered on the belief and hope that science-based prediction would provide a means to certify the continued performance, safety, and effectiveness of nuclear weapons in the Nation’s stockpile without relying on underground nuclear testing. The unique collaboration between the computing industry and the nuclear weapons program brought forth an unprecedented culture change in the weapon design community. Throughout the Cold War, weapon designers had relied on the direct empirical record provided by data from full-scale nuclear tests in order to certify the performance of weapon designs. The paradigm change was enabled by significant advances in computational physics, but also included some basic assumptions and complementary methodological and program modifications that would ensure that “ground truth” was still governing contemporary design certifications. Twenty-five years later, many intervening changes have taken place in the global and national security environment. The weapons complex is facing a major modernization program that promises to stretch the Stewardship capabilities to the limits. This talk will review the original constraints and assumptions under which the Stewardship program was launched, providing a framework for discussing the prospects for sustainability of the Stewardship paradigm in this new era.

Session 5
Applications: From Machines to Mission Science

A Short History of Software Development during LLNL's ASC Program

Teresa Bailey
Lawrence Livermore National Laboratory

The Science-Based Stockpile Stewardship era, which began in the mid-1990s, has relied on the Advanced Simulation and Computing (ASC) Programto provide computational tools that are used to underwrite stockpile assessments. The goal of these assessment is to ensure that the United States' nuclear weapons stockpile is safe, secure, and reliable. LLNL's ASC Program has developed and matured multiple production codes, which have become foundational elements for Stockpile Stewardship. We will discuss lessons learned from pre-ASC code development at LLNL, the evolution of our code projects during the ASC era, and potential future directions for software development.

Concurrent V&V/UQ and Code Capability Development: Hypersonic Reentry Analysis with SPARC

Sarah Kieweg
Sandia National Laboratories

The Sandia hypersonic reentry V&V and code development teams have demonstrated the benefits of doing concurrent V&V (verification and validation) with code development This presentation will give an overview of the V&V approach and credibility processes used. The V&V team has developed and demonstrated a sensitivity analysis, uncertainty quantification (UQ), and quantitative validation assessment workflow for computationally expensive simulations in coupled environments. During the FY20 ATDM L1 Milestone, the team demonstrated the capability to robustly perform the Dakota-driven analysis with SPARC and Aria on the HPC ARM computing platform (Astra). This presentation will give examples of these quantitative approaches using SPARC and a validation assessment from a published double-cone wind tunnel experiment. Finally, this presentation will summarize the benefits and challenges for concurrent V&V development, including (1) limited validation data, (2) creation and maintenance of V&V suites, and (3) V&V approaches in coupled analyses of combined environments.

Session 5
Applications: From Machines to Mission Science

A Look at Trends in Computational Methods: The Tie Between Machines and Mission Science

Aimee Hungerford
Los Alamos National Laboratory

The field of computational science has been invaluable in driving understanding of complex systems across a broad set of mission spaces. The methods and algorithms that are the tools in the computational science toolbox are impacted both by the nature of pressing questions in those mission spaces, but equally by trends in computing technologies. Computational methods essentially serve as the matchmaker between machines and mission science and are not always a match made in heaven. In this 40th anniversary of the Salishan conference, I will speak to some of the trends in computational methods that have influenced discovery science in mulitphysics applications over the years, including a look at where we are and where we are going at LANL.

Session 6
Exascale and Co-Design

History and Solutions to the Four Exascale Challenges

Al Geist
Oak Ridge National Laboratory

A decade ago, the HPC community identified four challenges that threatened to make it impossible to build a usable exascale computer. These were system power requirements, reliability, billion-way parallelism, and data movement given the growing memory wall. In 2021, the first exascale computer, called Frontier, will be delivered to Oak Ridge National Laboratory, followed by exascale computers at Argonne National Laboratory and Lawrence Livermore National Laboratory in 2022-2023.

This talk describes the history and solutions to the four exascale challenges that make these computers possible. It begins with the many false starts such as Nexus/Plexus and SPEC/ABLE, then describes the attempts to solve the challenges through Fast Forward and ECP. The talk shows how Frontier's design overcomes the original exascale challenges.

Session 6
Exascale and Co-Design

Towards an Integrated, Dynamic and Resource-Aware Software Stack for EU Exascale Systems

Martin Schulz
Technical University of Munich

The EuroHPC Joint Undertaking is a large European initiative to establish a broad high-performance computing ecosystem among its 32 European member states. It targets both hardware efforts, most prominently the European Processor Initiative (EPI) developing a new generation of energy efficient processors, and software activities towards an integrated HPC development and application environment. Ultimately, all EuroHPC efforts will combine into several European exascale class systems hosted by EuroHPC to be deployed in the next few years. A significant part of the software efforts target the development of an integrated software stack supporting the broad set of workloads targeted by EuroHPC and its user communities. In particular, the focus is on providing new capabilities to support more dynamic application and system behavior caused by complex application workflows and new workloads on one side, and increased variability due to power and energy constraints on the other. Dealing with such scenarios requires new concepts in operating systems and programming models, all the way to system-wide resource management. In this context, I will highlight two newly formed projects: DEEP-SEA, covering a comprehensive software stack for the first European exascale systems including new additions to MPI to support malleable applications, and REGALE, providing adaptive resource management across workflows, with a special focus on power and energy management. Combined, they will address these challenges for malleability and dynamic resource management and with that tackle some of the major obstacles we face on our path to exascale and beyond.

Session 6
Exascale and Co-Design

Fugaku: The First Exascale Supercomputer and its Innovative Arm A64FX CPU

Satoshi Matsuoka
RIKEN

Fugaku is the first exascale supercomputer in the world, designed and built primarily by Riken Center for Computational Science (R-CCS) and Fujitsu Ltd., but involving essentially all the major stakeholders in the Japanese HPC community with the 'Applications First' philosophy.

‘Fugaku’ is an alternative name for Mt. Fuji and was chosen to signify that the machine not only seeks very high performance, but also a broad base of users and applicability at the same time. The heart of Fugaku is the new Fujitsu A64FX Arm processor, which is 100% compliant to Aarch64 specifications, yet embodies technologies realized for the first time in a major server general-purpose CPU, such as 7nm process technology, on-package integrated HBM2 and terabyte-class SVE streaming capabilities, on-die embedded TOFU-D high-performance network including the network switch, and adoption of so-called ‘disaggregated architecture’ that allows separation and arbitrary combination of CPU core, memory, and network functions. Fugaku uses 158,974 A64FX CPUs in a single socket node configuration, making it the largest and fastest supercomputer ever created, signified by its groundbreaking achievements in major HPC benchmarks, as well as producing societal results in COVID-19 applications.

Quantum Computing - Why, Where, and When

Robert (Bo) Ewald
ColdQuanta

This talk will explore the motivation for developing quantum computing, starting with a talk given by Richard Feynman at Los Alamos’ 40th Anniversary in 1983. We’ll briefly review the international investments in things quantum and discuss the potential breadth of quantum-based products. The different technologies being developed for quantum computing will be reviewed and we will present a survey of the state-of-the-art of quantum computing. Finally, we’ll briefly look at the application domains where quantum computing may fit and discuss when it will be available for real-world applications.

Session 7
Reflections on Salishan's Impact Past and Future

	Session 7:	**Reflections on Salishan's Impact Past and Future**
	Moderator:	**Dona Crawford**
7:30 – 7:40 am	**Title:**	**Conference History**
	Speaker:	**Dona Crawford** *Livermore Lab Foundation*

<u>Panel Members</u>

Robert (Bo) Ewald
ColdQuanta

Manuel Vigil, Retired
Los Alamos National Laboratory

Jim Ang
Pacific Northwest National Laboratory

Don Dossa, Retired
Lawrence Livermore National Laboratory

Don Tolmie, Retired
Los Alamos National Laboratory

7:40 – 9:30 am	Panel Discussion