2019

Sandia National Laboratories

HIGH PERFORMANCE COMPUTING

SEE IT AND BELIEVE IT

Watch innovation come to life using SNLSimMagic, an augmented reality application developed at Sandia National Laboratories.

Download the SNLSimMagic application on your iOS device

Use your phone to scan images with the associated icon to watch the content come to life

CONTENTS

INTRODUCTION

SUSAN SEESTROM

ASSOCIATE LABORATORIES DIRECTOR

High performance computing (HPC) has long been central to Sandia National Laboratories' national security missions. From 1997's ASCI Red, the first terascale supercomputer, to 2006's Red Storm, which initiated one of the most important lines of supercomputers. Sandia and the National Nuclear Security Administration's (NNSA) Advanced Simulation and Computing (ASC) program have built a legacy of developing and deploying innovative HPC systems uniquely suited to stewarding the U.S. nuclear weapons stockpile. Over the last decade, Sandia's ASC contributions have continued in partnership with Los Alamos National Laboratory with the Cielo and Trinity HPC systems, which provided increased simulation capability focused on improving the efficiency of stockpile stewardship computations through co-design of computing technologies, algorithms, and applications.

The Sandia 2019 HPC Annual Report covers topics from enabling technologies such as HPC prototypes, quantum computing, and advanced machine learning to the impact of HPC on nuclear deterrence, cybersecurity, energy security, and other national security applications. Here, I want to highlight several high-impact achievements.

This year, Sandia and the ASC program took a new step in HPC by deploying the first Vanguard Advanced Architecture Prototype system, Astra. Building on our history of co-design, Astra is the world's first supercomputer based on Arm microprocessors, the type of chips commonly found in cellular phones and the Internet of Things. Also in support of the nuclear deterrent mission, Sandia is deploying new HPC algorithms and software to dramatically speed and enhance time-consuming digital surety and verification calculations for complex high-consequence systems. Likewise, HPC is performing advanced simulations using Sandia's Sierra engineering mechanics software suite to design and analyze tractor-trailers for securely transporting nuclear weapons and materials.

Our national security mission extends far beyond the nuclear deterrence. For example, HPC is contributing to our work on cybersecurity by helping to train machine-learning algorithms that recognize and mitigate threats to industrial control systems. HPC also enables Sandia's unique Emulytics capability to create complex, fine-grained emulations of networked systems with tens of thousands of endpoints to enhance cybersecurity.

To ensure HPC is available for a growing user base, Sandia builds on its hands-on experience to improve HPC systems management, including new approaches to optimizing power consumption, monitoring system performance, and using containers and virtualization for agile application deployment.

Sandia research in computing is exploring entirely new computing paradigms. We are working to realize the potential of neuromorphic computing through tools such as Sandia's Whetstone software package, which trains artificial neural networks to behave more like their natural analogues. Our new Neural Exploration and Research Lab offers researchers a neuromorphic computing testbed. Similarly, we are actively exploring quantum information science including quantum simulation algorithms for materials science, while Quantum Scientific Open User Testbed is deploying a quantum computing testbed using trapped-ion qubit technology.

This work is possible only with the support of our sponsors and partners in HPC, especially our primary sponsors: the NNSA ASC program, the Department of Energy Office of Science Advanced Scientific Computing Research program, and the joint NNSA/Office of Science Exascale Computing Project.

I hope you enjoy this Annual Report and share my excitement about the potential for HPC to tackle some of the greatest challenges facing our nation and the world.

CRAY LEVERAGES SANDIA'S AWARD-WINNING SOFTWARE IN LATEST EXASCALE SUPERCOMPUTER

Jim Brandt

CRAY AND SANDIA have a long-established tradition of partnership that dates back to developing the architecture for one of the first scalable teraflop supercomputers, Red Storm. Sandia and Cray are partnering once again to reach the next milestone in supercomputing – Exascale. Advancing system and application performance is highly dependent on effective use of system resources. The Cray Shasta next generation platform will provide more extensive monitoring services than previous generations in order to facilitate more efficient operation.

To ensure their new monitoring capabilities meet site needs, Cray released a "Software Preview System" to various end users for evaluation and feedback. Engineers from Sandia, Cray, and the National Center for Supercomputing Applications are assessing advanced monitoring capabilities in the Software Preview System, which incorporates Sandia's open source Lightweight Distributed Metrics Service (LDMS) software package. Together these researchers are evaluating this new framework for node- and application-level monitoring for eventual deployment on Shasta. The new Shasta architecture provides a modern challenge to traditional monitoring services as it uses container technology for providing services, including monitoring, and also enables users to run applications both on bare metal and within a containerized environment.

LDMS, an R&D100 winning software, is one of the only high performance computing monitoring tools that provides continuous high-fidelity whole system awareness, allowing system administrators, application developers, and users the ability to explore and understand resource utilization, network congestion, I/O bottlenecks, and associated causes of compute inefficiencies. LDMS works both within and alongside the Cray-provided monitoring framework, which includes log and numeric data collection, as well as a transport bus and published Application Power Interfaces to enable users to easily extract this data from the system.

The Shasta architecture will be the basis of the Department of Energy's (DOE's) first Exascale computing system, Aurora. Aurora is expected to be delivered to Argonne National Laboratories and will consist of over 200 Shasta racks. Aurora will support continued high-fidelity modeling and simulation, and artifical intelligence, throughout the DOE national laboratory complex.

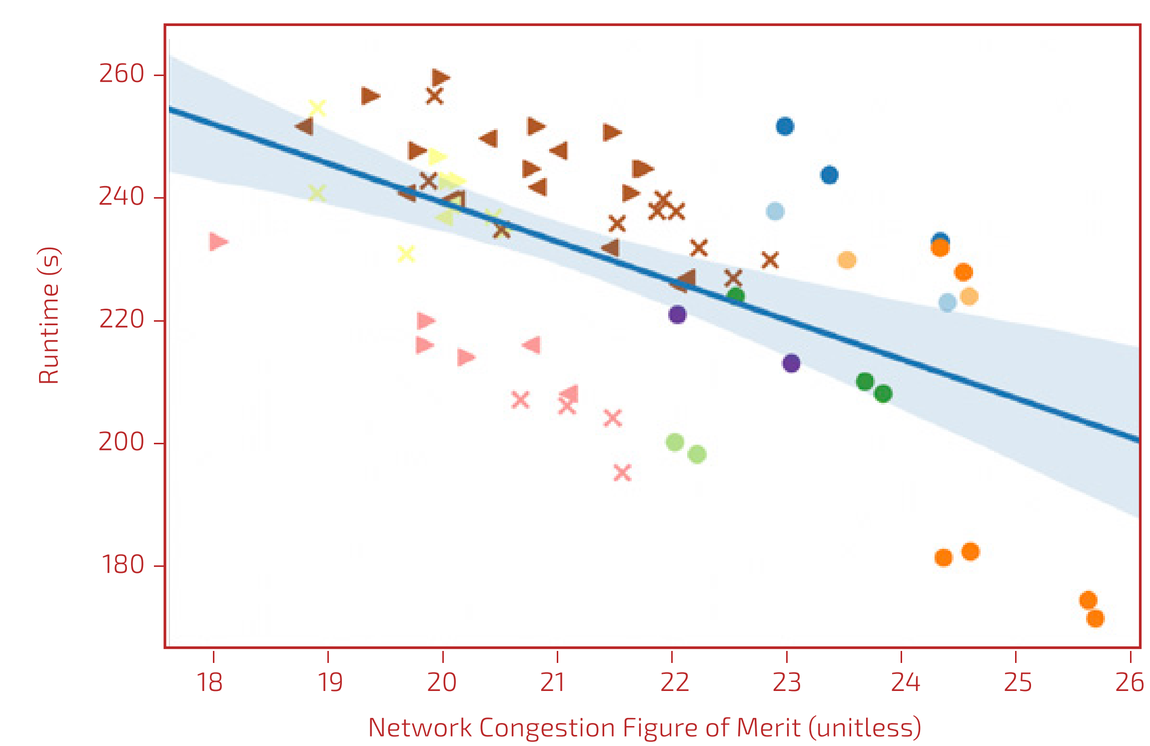

REMS-MBD

Cheryl Lam

REMS-MBD is an enabling capability to reduce the development cycle time for nuclear deterrence systems by using integrated, credible models and simulations to ensure electrical system function and survivability under stressing radiation environments.

RADIATION-AWARE ELECTRICAL MULTI-SCALE MODEL-BASED DESIGN

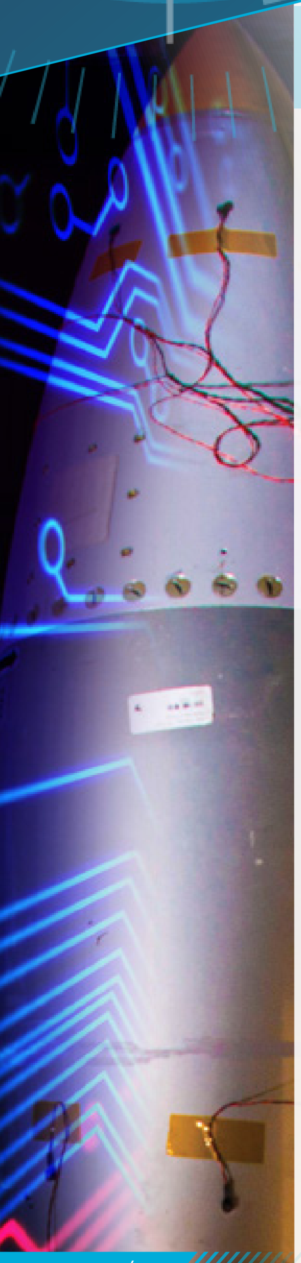

As stewards of the nation's nuclear deterrence stockpile, Sandia must ensure that crucial components of each nuclear weapon can survive and predictably respond to radiation exposure.

Sandia is modernizing the stockpile with radiation-hardened electronics, including Sandia-developed semiconductor and integrated circuit technologies. Alongside continued advances in trusted microsystems and electronics design, Sandia has developed models to predict electrical system response to ionizing radiation and electromagnetic environments. Simulations informed by radiation-aware models are essential to characterizing and predicting the performance of critical electrical circuits within a nuclear weapon.

RAMSES

This project is leveraging simulation tools developed by the Advanced Simulation & Computing (ASC) program. The Radiation Analysis Modeling and Simulation of Electrical Systems (RAMSES) code suite is used to predict the performance and survivability of weapon systems and components when exposed to hostile radiation environments and electromagnetic insults. Its component modules include Xyce (electrical circuit), Charon (electronic device physics), CHEETAH/ITS and CEPTRE (radiation transport, including coupled electron-photon), GEMMA/EIGER and EMPIRE/EMPHASIS and (electromagnetics and plasma), and q (formal verification and analysis for digital systems).

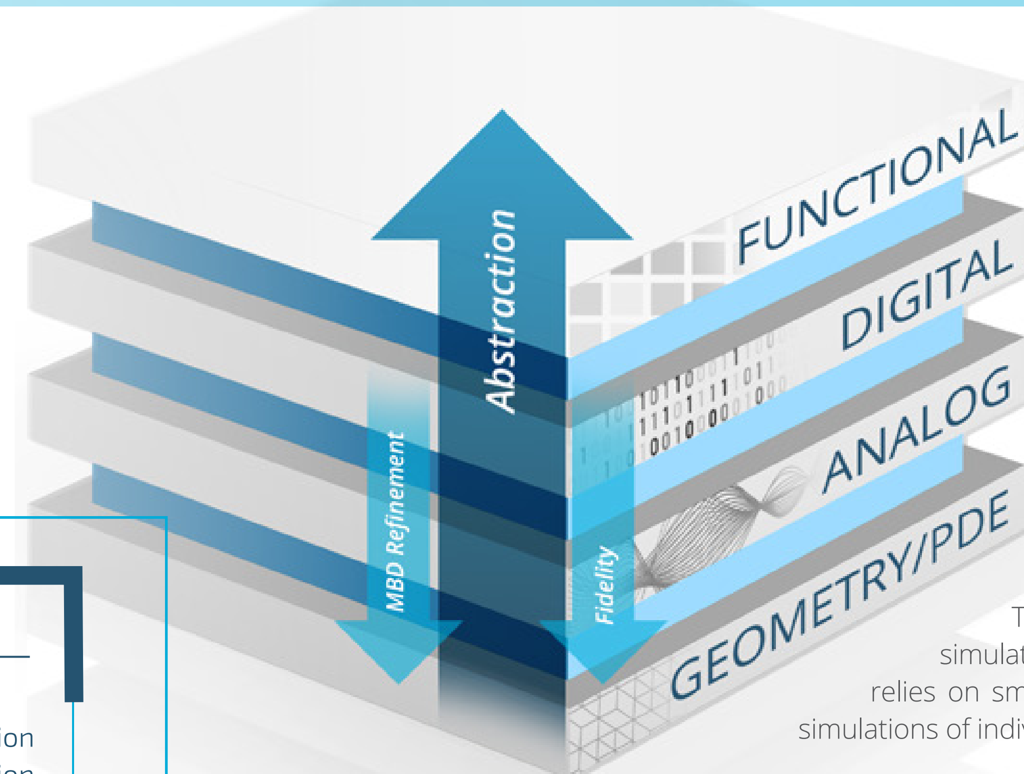

Sandia is focused on further developing these models and integrating them into early system design—enabling more agile and rapid development, design, and delivery of critical electrical systems for national security missions. Integrated models facilitate design refinement and virtual test to inform complex engineering decisions.

See the next article for more information on Model-Based Design

A nuclear weapon can have thousands of components. The computational cost to run integrated high fidelity simulations for the entire system can be very high. So REMS-MBD relies on smaller more abstracted models informed by high fidelity simulations of individual electrical components.

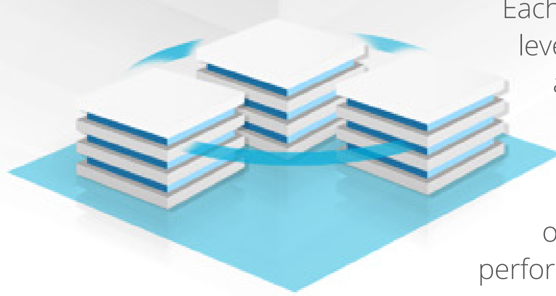

Each component within the system can be modeled at different scales or levels of abstraction (functional, digital, analog, and geometric). The models at each level are needed to capture key physics at the relevant physical scale. The models must then be consistently integrated with each other across multiple scales to simulate the integrated system response. When the component-level models are viewed together on a system-level basis, it provides prediction of radiation effects to assess performance margins and design trade-offs.

FAILURE IS NOT AN OPTION:
FORMAL VERIFICATION IN HIGH-CONSEQUENCE SYSTEMS

Geoffrey Hulette

HUMANS INTERACT with computer programs all the time without thinking about it. Your alarm clock has a program that knows when to wake you up. The car has dozens of complex programs that provide simple conveniences, as well as critical functions like steering or applying the brakes. However, computer programs can, and often do, fail. In most cases, these failures are not catastrophic. But in some cases failure could be life-threatening. But consequences of program failures in high-consequence systems such as airplane safety systems, medical devices, or even nuclear weapons would have graver repercussions.

So how does one ensure the safety and reliability of digital computer programs in high-consequence systems, which unlike others, can't be tested over and over again in "live" scenarios? At Sandia, that is the role of digital surety. Sandia's approach involves a process known as formal verification, which uses deductive logic—the basis of mathematical proofs dating back to Euclid—as the foundation for program assurance. Digital controllers are generally harder to verify than traditional, mechanical controllers because they usually have a very large number of states; that is, distinct configurations of their digital components such as registers, memory, etc. They have so many states, in fact, checking all of them one at a time is time consuming.

Unlike consumer products, Sandia's high-consequence systems cannot be field tested to discover any bugs in the software. Instead, formal verification uses inductive invariants and pre- and post-conditions—essentially, things that can be shown to be true for whole sections of the program—to verify large or even infinite numbers of states all at once. Using formal verification, Sandia's engineers can provide mathematical evidence that their programs meet safety and reliability requirements, effectively guaranteeing the final program will be error free. Furthermore, formal verification does not have to be applied only after the program has been written, it can be used during program development to detect errors early, preventing costly program rewrites.

So, how does one formally verify a computer program? It starts by building a model of the program behavior. This is often done with some form or variation of Hoare predicates (named after computer scientist Tony Hoare who invented the idea in 1969). Hoare predicates characterize program statements by logical pre- and post-conditions—what must be true before the statement executes, and what the statement guarantees will be true after execution. These conditions are expressed in a mathematical language, such as first-order logic with set theory, or a higher-order typed logic. The use of regular mathematics, instead of program constructs, to characterize the behavior of the program allows the behavior to be modeled using powerful abstractions, like quantification over all values a variable might take. Hoare logic is often extended in various ways, depending on what is being verified. For example, separation logic extends Hoare logic with a logical model of separated regions of memory, facilitating verification of dynamic memory management and concurrent algorithms.

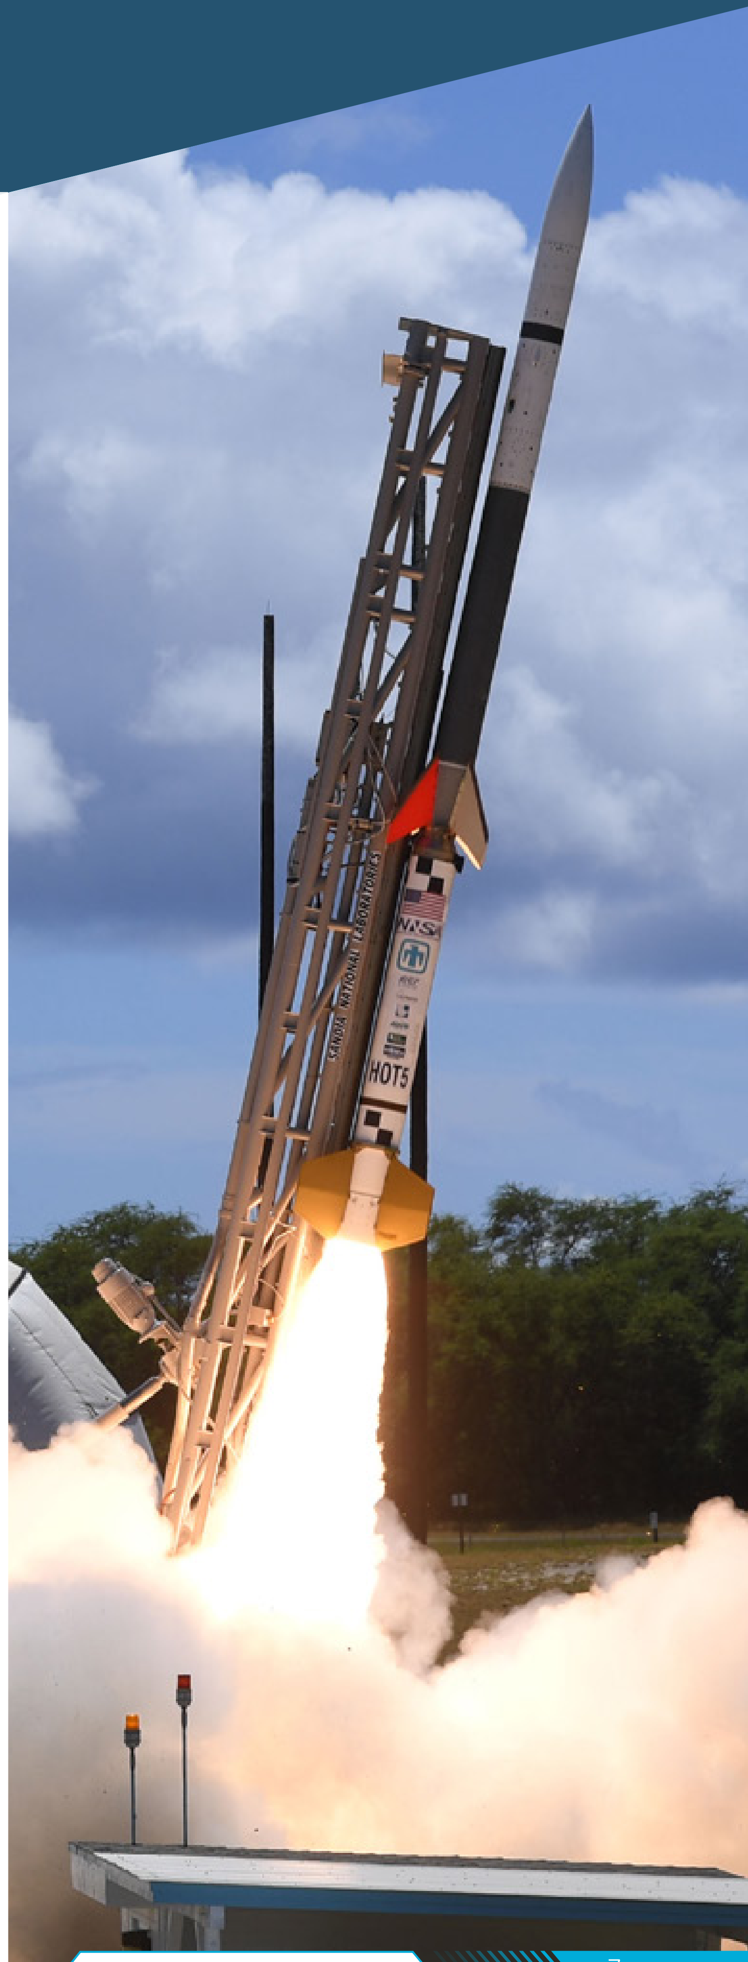

Once the program is characterized this way, deductive logic can be applied to find a proof of a desired property. Because programs tend to be long and complicated, the proofs needed are usually long and complicated as well, hence the need for the power of high-performance computing (HPC). Automated theorem proving tools, such as Boolean satisfiability (SAT) solvers and satisfiability modulo theories (SMT) solvers have advanced enormously in the past two decades, to the point where properties on large programs can be proven automatically, often within seconds. Sometimes, for more complicated properties, a human expert is needed to guide the proof-finding process after which SMT/SAT can be used to fill in the details.

Today, not all formal methods tools generate certificates—in this case, the formal methods tools themselves must be verified, and although this is still evidence, it is weaker evidence than a certificate since the tools themselves might have errors. Because reproducible evidence is so important to Sandia's mission, researchers are extending commercial off-the-shelf (COTS) formal methods tools to help produce proof certificates and to "connect the dots" between certificates produced from different tools.

The process for formal verification described clearly requires a lot of computing power for those tools and algorithms. Even for a very small digital controller, complete formal verification requires finding and checking thousands of proofs, and each proof may involve hundreds of thousands of individual clauses. Sandia researchers are exploring new ways to leverage HPC for this critical process of formal verification. As this process continues to gain traction, the need for greater HPC cycles will increase.

For critical systems to operate safely and reliably, especially as digital systems become more integral to their function, a robust method for assuring their reliability is clearly vital. Formal verification is one way to ensure the reliability and safety of these systems in an empirical manner, providing confidence that they will work how they're meant to work when they are meant to work. This will be a need that will continue to grow as digital elements become further and further integrated into both high-consequence as well as every day systems.

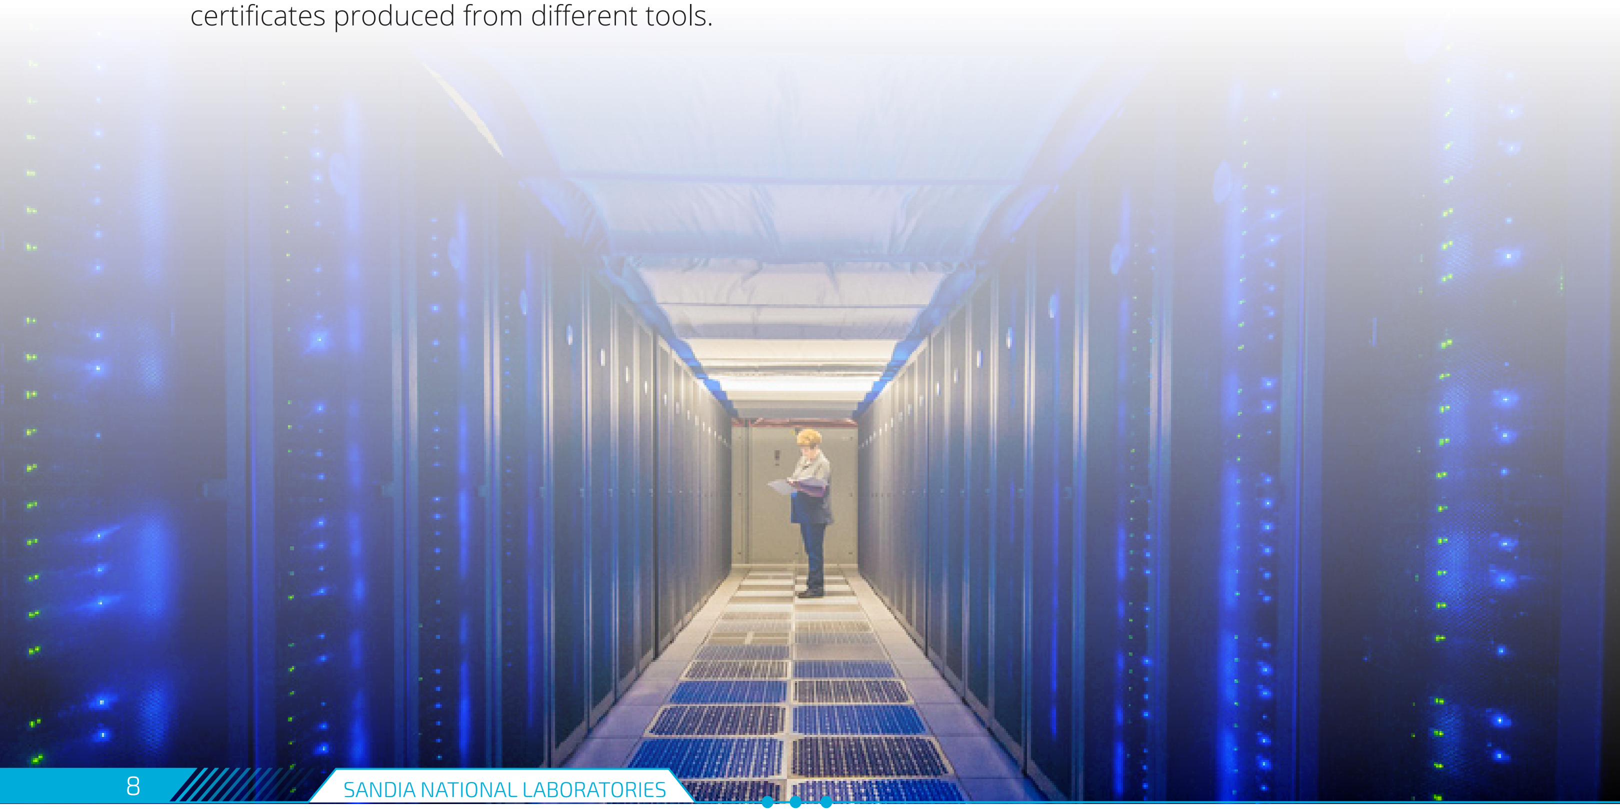

CYBERSECURITY ON HPC SYSTEMS

Aron Warren

WHETHER ONE REMINISCES on 8-bit games, floppy disks, or dial-up connections, it is easy to see how rapidly and substantially computing technologies evolve. Unfortunately, greater complexity also means more vulnerabilities and attack vectors. Just like security patches can be added to our personal computers, high performance computers (HPCs) must also be augmented or reengineered to prevent the more sophisticated attackers. However, upgrading even a minor patch on an HPC system could require dozens of hours of labor and halt the HPC's processing capabilities for upwards of several hours. That's why one team at Sandia must simultaneously apply their technical expertise with their regular investigation of ever-evolving cybersecurity requirements. This combination of focus areas allows Sandia to pinpoint the exact HPC enhancements needed to perpetually maintain the most secure and up-to-date HPC operating systems.

To facilitate constructive upgrades, researchers must first ensure that the HPC operating system and configuration meet all current federal and corporate cybersecurity requirements. Incorporating such vigorous requirements is an especially unique challenge for HPCs, as these requirements are typically written from a non-Linux mindset. Additionally, as the number of requirements increases to combat more sophisticated threats, it becomes more difficult to balance compliance with efficient fulfillment of already well-defined mission objectives. Collaboration with staff at other laboratories incorporates others' experiences to improve efficiency and streamline solutions in such challenge areas. It also confirms that implementations of cybersecurity requirements are compatible within the Tri-Labs enterprise.

When new cybersecurity requirements have been fully defined and scoped, they can then be effectively and efficiently implemented. As a result of meeting the new challenges, additional automated processes have been successfully developed and integrated alongside improved monitoring capabilities using a Splunk infrastructure. These improvements have allowed staff to better distinguish between benign and malicious attacks, thus reducing investigation time.

The timely application of security upgrades is not only beneficial to the security of HPCs themselves—they also save significant expense and labor hours due to non-compliance with federal and corporate requirements. To aid in performing these upgrades this year, researchers worked on automating security patching. Using the internally developed Remote Automated Patching System software framework, Sandia staff were able to deploy security patches to the HPC cluster in a matter of minutes; previously, such work required individual staff investing dozens of hours per month.

Sandia's HPC teams continue to provide a consistent, stable, and secure computing environment to further Sandia's exceptional service in the national interest.

USING EMULYTICS TO SIMULATE THE UNITED KINGDOM

Jacob Fritz, Vince Urias

IN 2019, the Emulytics team from Sandia was invited to the CyberUK conference to share insights on using Emulytics for modeling and simulating on an extremely large scale. The presentation was especially pertinent because one of Sandia's partners, the National Cyber Security Centre (NCSC), was creating an IT model of the U.K.'s infrastructure. The challenge of building a model/simulation of an entire country using Emulytics began over a decade ago.

Emulytics emerged from a series of cyber-attacks that took place in 2007. Around this time, the KVM tool was being merged into the greater Linux environment. Sandia researchers found KVM's abilities to launch virtual machines intriguing. While virtual machines were not new at that time, it was the first time there was a native Linux capability for launching or booting those virtual machines The Sandia team researched how many KVM virtual machines could be booted, and it turned out they were able to boot up around four million.

Following this, Laboratory Directed Research and Development funding was used to build-out the capability around very large-scale virtual machine test beds. The team started exploring the possibility of running Monte Carlo-type questions on those machines, i.e., the possibility of having 10,000 variants of the same experiment run over a short time-frame. Initially, a set of tools called Mega was being used to explore these questions. The tools were reimplemented in 2012 and came to be known as minimega.

Today there is a new tool set to compose virtual machine-based environments, with use cases for this capability ranging from training environments to nation-state sized emulations. The tools can be used to build a copy of a company's network and then be used to train enterprise operations and instant responders, or to examine DevOps questions. Because networks may be very large, perhaps around 60,000 endpoints, a high enough fidelity environment is required to actually test with some confidence.

There are two dimensions to these experiments: the structure of the thing being modeled and the behavior of what is being modeled. Going back to the example of simulating a company's network, to successfully do that, there should be workers conducting ordinary day-to-day functions on their computers so as to establish a behavioral model vis á vis the structural model.

These capabilities represent the current-state for Emulytics at Sandia. There are two different business models for using the emulytic tools. One involves answering a question for a customer using emulytics. The other model focuses on providing the tools as a product; tools are open source and available for anyone to use. In the latter instance, the team delivers the platform to a customer, who builds an experiment, hosts it on their own infrastructure, and tweaks the experiment as they go.

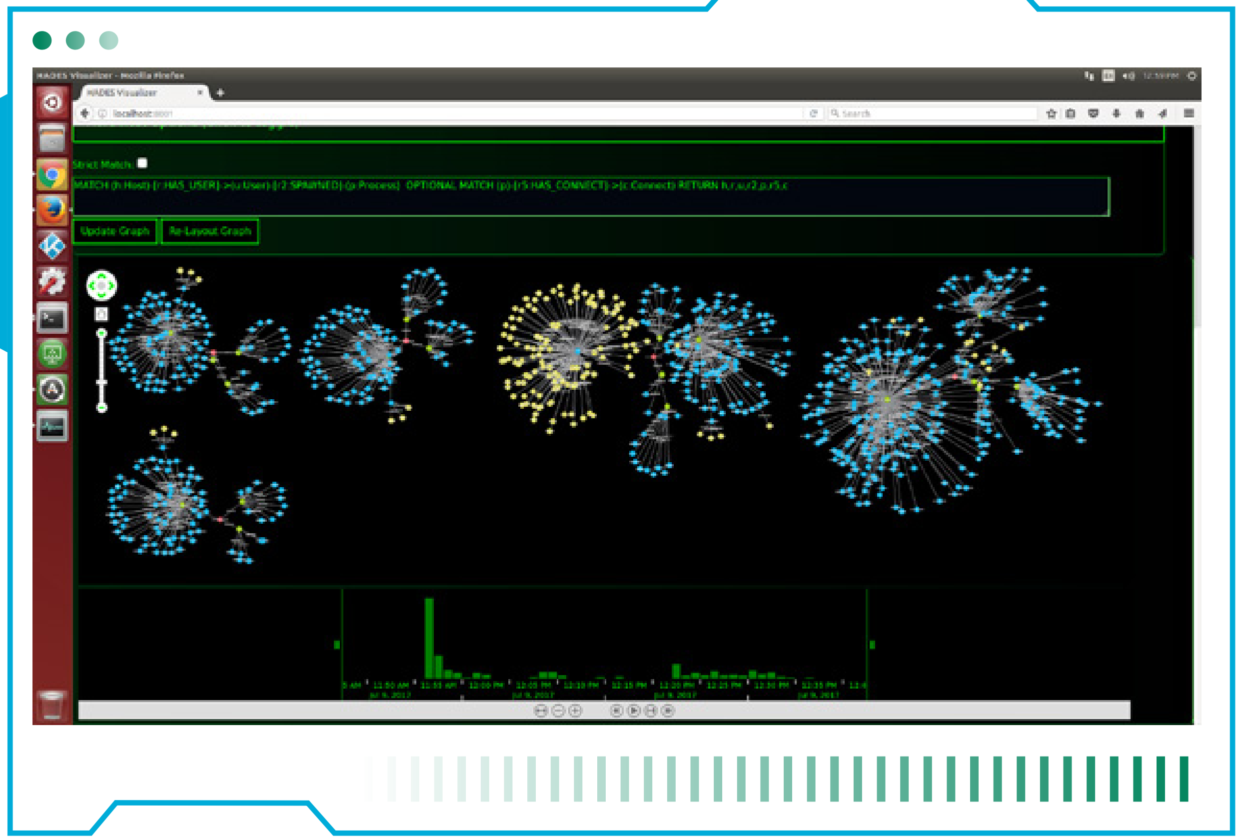

The business models highlight how there are two different research avenues when it comes to emulytics: emulytics as a platform, i.e., researching in emulytics proper and secondly, research that is enabled by emulytics. If the ability exists to model an entire giant network or whole state, what kinds of new questions can be asked? There are few limits if there's enough fidelity.

In 2017, the NCSC approached Sandia and its emulytics team because they had recently received a new charge to model the entire cyber infrastructure of the U.K. They wanted to look at what happens when telecommunications are blended with electric power, with petroleum production, and with other IT utilities. Since the start of this collaboration, NCSC has been using emulytics—specifically minimegas—to build models to accomplish their goal. NCSC recently purchased a cluster that mirrors one of the dedicated Emulytics clusters at Sandia.

RETHINKING TEAMING AND PRODUCTIVITY IMPROVEMENT TO SCALE UP COMPUTE POWER

Elaine Raybourn

THE DEPARTMENT OF ENERGY'S EXASCALE COMPUTING PROJECT (ECP) was established in 2016 to maximize the benefits of high performance computing and accelerate the development of an exascale computing ecosystem by 2023. Central to this project is development of next generation modeling and simulation applications that fully exploit emerging architectures to analyze more data in less time—providing insights and answers to the most critical U.S. challenges in scientific discovery, energy assurance, economic competitiveness, and national security. Interoperable Design of Extreme-Scale Application Software (IDEAS-ECP), led by Sandia National Laboratories' Michael Heroux and Argonne National Laboratory's (ANL) Lois McInnes was established to support the ECP mission by addressing challenges in software development productivity, quality, and software sustainability. IDEAS-ECP addresses these issues by investigating ways to improve practices that enable collaboration and contributions from teams across the Department of Energy (DOE) laboratory-complex.

While historically much scientific software development has occurred in smaller teams, one of the most important outcomes of ECP is the ability to increasingly couple multiple types and scales of physics, science and data, providing much higher fidelity and a deeper understanding through computation. At the same time, the complexity of the applications, teams and interactions also increases, requiring a better understanding and improvement of practices and processes. The IDEAS-ECP effort at Sandia led by Principal Investigator Elaine Raybourn is working with collaborators from Los Alamos National Laboratory (LANL) and other national labs on scaling the successes and agility of smaller teams to networks of larger, geographically distributed teams (which can be challenging given that ECP comprises over 1,000 researchers and scientists from Sandia and across the DOE-laboratory complex). When small teams make their software available to new developers or users, the team dynamics often need to change to form new processes as others want to contribute to software development or adopt codes into theirs. Raybourn's multi-lab team conducts Productivity and Sustainability Improvement Planning (PSIP) with ECP application development and software technology teams, while using a team of teams model, based on Gen. Stanley McChrystal's experience and best-selling book.

To create a team of teams, Raybourn's team strives to understand productivity bottlenecks and improve software development practices through the application of PSIP. PSIP is a lightweight process aimed at mitigating technical risk by aligning software development activities, enabling partnerships, and facilitating adoption of best practices across aggregate teams. One example is with Exascale Atomistic capability for Accuracy, Length and Time EXAALT (EXAALT), a materials modeling framework for accelerated molecular dynamics simulations. EXAALT is a collection of multiple software components forming a team of teams network, each component depending on the others. The team of teams includes EXAALT, Large-scale Atomic/Molecular Massively Parallel Simulator (LAAMPS), developed at Sandia, LATTE and ParSplice, both developed at LANL. Team members from ANL, LANL, and Sandia successfully collaborated with EXAALT to formulate progress tracking cards for the actionable items needed to lay the foundation for continuous integration within the existing software repository.

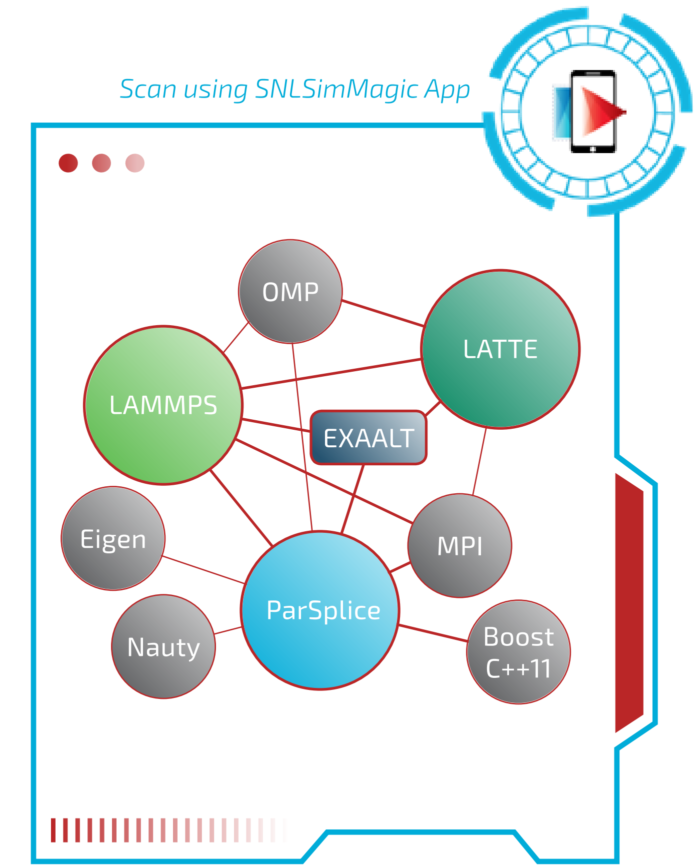

As IDEAS-ECP continues to meet with teams and understands how they function to help them apply methodologies to improve development practices, they hope to increase collaboration across the DOE complex—leading to innovative application design, sustainable software technology, and furthering U.S. leadership in high performance computing as we approach the exascale milestone.

TRANSPORTATION RESEARCH AND ASSESSMENT

John Bignell, Adam Brink

THE NATIONAL NUCLEAR SECURITY ADMINISTRATION'S OFFICE OF SECURE TRANSPORTATION (OST) has been safely and securely transporting nuclear weapons, weapon components, and special nuclear materials in purpose-built tractor-trailers along U.S. highways for decades. The enduring desire to provide the highest practical safety to the public makes advances in transportation system performance increasingly important. Sandia supports OST in this regard by performing evaluations of current transportation systems as well as research into advanced systems for future deployment.

Each transportation system is required to ensure the safety of the public and the security of on-board assets under a wide range of conditions that include normal operation, severe or unusual accidents (e.g., highway crash, large fuel fires, lightning), and hostile environments. Historically, research and assessment activities have relied exclusively on physical testing, including numerous large-scale, sub-system and full-system test scenarios. The complexity of these tests makes them time consuming and expensive to perform, limiting the ability of scientists and engineers to quickly investigate and address concerns, or fully explore the scenario space of interest. Additionally, while system-level tests provide invaluable performance data, they often supply only limited data from which the underlying mechanisms affecting performance can be understood.

New capabilities in high performance computing (HPC) and modeling and simulation (M&S) are being used to address the shortcomings of a purely test-based approach. At Sandia, HPC resources and the Sandia's SIERRA code suite are being utilized in a synergistic manner with selective testing to help OST more efficiently and effectively answer pressing questions. For example, system-level tractor/trailer SIERRA Solid Mechanics (SIERRA/SM) models that comprise approximately 10 million elements have recently been used to explore the crash performance of trailer designs. The simulations, which are computationally intensive and require between 30 and 60 days to complete using 700 to 1,000 cores, are possible only through the HPC resources at Sandia. Simulations are used to investigate various crash scenarios, including the side impact of the trailer by a second tractor/trailer (Figure 1), and results used to inform planning for future test series. Data from the tests will not only demonstrate system performance but will be used to assess and improve the accuracy of the models, which can then be applied to investigate more fully the crash scenario space and more completely understand the underlying phenomena driving the performance of the system.

Sub-system SIERRA Thermal Fluids simulations have been used to evaluate trailer thermal protection system designs. Radiant heat and fuel fire tests of trailer wall sections provide data to calibrate models, which in turn help to evaluate design options (Figure 2). Additional panel tests can validate the performance and confirm the accuracy of the model. Because these simulations are relatively modest in cost, requiring between 1 and 5 days to complete using 16 to 256 cores, they drastically cut the time necessary to arrive at a final design and reduced the number of confirmatory tests that must be performed. With the accuracy of the models demonstrated, system performance may be confidently predicted with a system model (7 million elements, 4 to 6 day run time on 256 cores) with a reduced need to resort to complex, expensive system-level testing.

Leveraging these new HPC and modeling and simulation capabilities at Sandia to more effectively and efficiently support the mission of OST ultimately ensures safety and security for us all.

FIGURE 1

Scan using SNLSimMagic App

Side impact crash scenario investigated analytically using Sandia's high-performance computing resources.

FIGURE 2

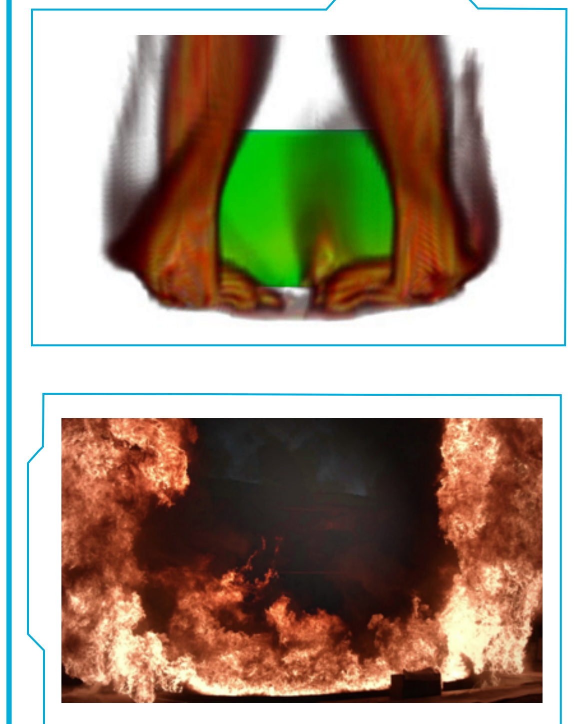

Fuel fire test and simulation.

ARTIFICIAL DIVERSITY AND DEFENSE SECURITY

Adrian Chavez, Jason Hamlet, William Stout

INDUSTRIAL CONTROL SYSTEM (ICS) environments, such as the electric power grid, oil and natural gas refineries, and water pipelines, continue to use predictable communication paths, static configurations, and unpatched software; all of which benefit adversaries. Sandia has developed a technology, named "Artificial Diversity and Defense Security" (ADDSec), that automatically detects and responds to threats within critical infrastructure environments in real-time. The detection component is used to identify anomalous behavior and then classify those anomalies into categories of attacks. Then, based on this classification, the system activates the appropriate mitigation for the threat. The responses include several Moving Target Defense (MTD) strategies that modify the underlying environment to disrupt the reconnaissance phase of an attack and invalidates any knowledge gained by the adversary.

The ADDSec detection mechanism leverages an ensemble of supervised machine learning (ML) algorithms that are trained using over a hundred host-and network-based features. On commodity computing hardware, the training process may take more than ten minutes to complete; on HPC platforms, that time can be reduced to seconds. HPC platforms also make it possible to include more sophisticated analysis techniques. For example, rather than modeling the behavior of individual machines, HPC would permit modeling at the network level in near real-time. Additionally, HPCs could also be used to build models of the physical behavior of the ICS, and then features extracted from these models could be used in conjunction with the host and network features. The ADDSec ML algorithms and features continuously vary over time to create a dynamic framework that allows the ensemble to adapt to new notions of normal and abnormal system behavior, presenting a moving target to adversarial ML attacks. This moving target approach attempts to complicate attacker efforts to coax the ensemble into classifying attacker behavior as normal during retraining intervals of the ML ensemble.

After the detection mechanism discovers threats in the system, ADDSec can automatically respond to those threats using several MTD techniques. The MTD algorithms' primary goal is disrupting the reconnaissance phase of an attack. The MTD strategies include randomizing network features, communication paths, and application library function locations. The two network-based MTD techniques include randomizing (1) Internet Protocol addresses; and (2) application port numbers. Randomization intervals may be configured to evade adversarial discovery. The randomization algorithms reside at the network layer and are transparent to the end devices. Improved usability and scalability comes from the fact that the randomization algorithms are managed at the network level and do not need to be deployed at every end device. HPC can be further used to provide a flexible simulation and emulation platform to test the scalability of the algorithms and software controllers themselves.

The ADDSec solution can be baked into modern ICS infrastructures, or can be retrofitted into existing networks using commercial or open source software. Compatibility with legacy and modern hardware switches is especially important for ICS systems because ICS systems typically do not replace existing hardware for decades at a time to reduce costs and also to maintain high levels of availability. To date, ADDSec has been applied to a representative industrial control system environment and has been shown to increase the resiliency against attacks by automatically detecting and responding to threats in the initial phases of an attack in real-time. Because attacks on ICS environments could translate into physical actions with consequences that threaten public health and safety, cybersecurity is a necessity.

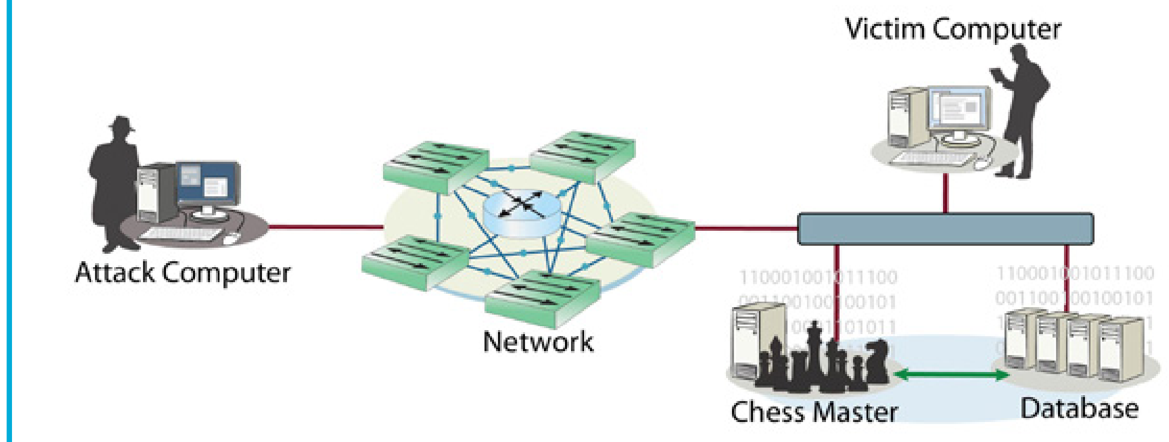

THE POWER OF HPC AND PARTNERING WITH INDUSTRY

Rekha Rao and Dan Moser

HIGH PERFORMANCE COMPUTERS CAN SOLVE bigger, more complex problems in a shorter amount of time—a powerful capability that's particularly useful throughout industry. From complex modeling to detailed analysis, HPC capabilities can help companies advance the state-of-the-art, or help get ideas from the drawing board to the production floor in a shorter amount of time. The HPC for Manufacturing (HPC4Mfg) program connects U.S. manufacturers to the Department of Energy's (DOE's) unique computing resources to advance clean energy technologies and increase energy efficiency while reducing the need for HPC adoption. Few companies can afford an investment in HPCs. Since its formation in 2015, HPC4Mfg has sponsored 67 projects, and Sandia is currently working on several projects.

With the help of Sandia's HPC capabilities, 3M is creating a film for windows that combines a visible/near-infrared mirror to prevent heating during daylight, and a radiative emissive structure for an IR atmospheric window to provide cooling. Using Sandia's HPC resources to model the design and manufacturing process to produce composite, metamaterial films for passive solar cooling could allow 3M to develop low cost, electricity-free cooling options within a 1-3-year window. Without this computational design and process modeling effort, this technology would still be 10 or more years away from realization.

Sandia's HPC capabilities are also being leveraged in the growing field of additive manufacturing. Vader Systems (recently acquired by Xerox) is working to improve their patented MagnetoJet (MJ) metal additive manufacturing technology. To order to increase efficiency, the nozzle needs to be improved to ensure precise control of the melted aluminum wire at higher frequencies. The Aria module of the Sandia-developed Sierra finite element multiphysics suite is being used to develop models that will explore ways of improving the build rate of Vader's MJ process, which could potentially expand the use of additive manufacturing throughout industry by making it cheaper, faster, and more energy efficient.

In the third project, Sandia's expertise in modeling is providing insights into better insulation. Dow Chemical is the world's leading supplier of polyurethane formulations and is working to increase insulation efficiency without increasing, and perhaps even by decreasing, material usage. A new formulation/process-design paradigm informed by predictive modeling shows promise in achieving these goals without the need for slow and costly build-test cycles. For this effort, Sandia is using massively parallel computational fluid dynamics-based foam models, which will be enhanced and customized via Dow's expertise in kinetics and bubble-scale physics. The impact of optimizing polyurethane foam products and their manufacturing will be seen in several areas of practical application, starting with refrigerators and water heaters.

The HPC4Mfg program aims to solve difficult challenges and develop new technologies that will shape the future. Its purpose is to develop new manufacturing processes that are more cost and energy efficient leveraging HPC capabilities resident in DOE laboratories.

HPCART:

THE CRAFT OF HIGHER PERFORMANCE

Anthony Agelastos, Justin Lamb, Doug Pase, Ryan Shaw, Joel Stevenson, Tim Thomas

THE HIGH PERFORMANCE COMPUTING APPLICATION READINESS TEAM, or HPCART, was set up to solve Sandia's toughest high performance computing (HPC) problems; the ones where answers can remain elusive. Grasping the essentials of diverse application domains is critical, as is the team's blend of science, engineering and mathematics skills. The team focuses on the effectiveness of large and complex scientific and engineering applications (simulation and analysis software, or "codes"). The team engages with Sandia researchers in need of support of HPC systems' procurement, deployment, and operation and management needs. They then work closely with code developers and code users to provide comprehensive and ongoing support, enhancement, and acceleration of codes and related workflows.

The team provides both strategic and tactical support for HPC operations. Strategic support involves maintaining installations of and expertise with debugging and profiling toolchains and assisting HPC resource users in their applications. Essentially, HPCART bridges the gap between users, developers, administrators, procurers, and planners of HPC systems and software.

The team has extensive training and experience with parallel systems and applications software and makes use of advanced debugging tools that are specialized to uncover errors in parallel codes running at very large scales on the supercomputers at Sandia. Tools include TotalView and Allinea Distributed Debugging Tool (DDT) A different set of sophisticated tools called performance analyzers are deployed when codes run too slowly relative to mission requirements. These tools are designed to explore performance issues at different levels of detail, including surveyors, profilers, and deep tracers. Performance analyzers don't automatically make code run faster, but they can help an analyst understand the purpose of an application and where a speedup might be most productive. Regular feedback and close collaboration with developers and tool providers ensures that the latest features are incorporated into Sandia's HPC sites.

The bulk of HPC applications at Sandia comprise several applications tied together into multi-pathed data and computational flows. Both software producers and consumers typically use complicated workflows, and HPC resources are used in ever-changing ways. Developer workflows typically parallel-compile, link, and perform low-level verification of HPC programs such as checks for memory leaks. Software component version management, testing, and deployment are all very significant challenges in relation to developer workflows. During the process of mapping their domain-specific problems onto Sandia's large-scale HPC resources, analysts (code consumers) can accidentally adversely impact HPC systems operations. HPCART helps users fit their application workflows into the HPC systems' environments and assists in adapting existing workflows to improve their efficiency, performance, and overall effectiveness. Improvements to workflows are done by applying advance technologies, such as the use of containers and continuous integration.

HPCART developed an advanced workflow management solution called Imperator, a generalized workflow tool that can automatically execute complex user-defined workflows within multiple HPC environments. When a workflow based on Dakota (a Sandia-developed research code) was causing the system to crash, HPCART was able to solve the problem with a temporary work area in a RAM-based disk. Each short job had processed dozens of small input files and produced dozens of small output files, resulting in the rapid creation, opening, writing and reading, closing, and deleting of roughly half a million files every few seconds. This then ended up crashing the metadata server of the underlying parallel, high-performance file system.

HPCART also led an Advanced Simulation & Computing Facility Operations and User Support-directed milestone to introduce a broad set of performance and analysis tools on two Sandia codes, Sierra/Aria and Integrated TIGER Series (ITS). The goal was to use these tools to generate useful scaling, trend, and profile data over multiple code versions, use cases, and HPC platforms. This wide set of data could then be used by application developers to evaluate performance trends between code versions, pinpoint weak areas of their code, and determine performance deficiencies on various HPC architectures. At the same time, the tools were evaluated along with their utility on various platforms. To obtain the required performance results on the selected codes, tools were chosen to cover the range of performance analysis domains such as surveying, profiling, and tracing. HPCART implemented a methodology for diagnosing performance problems that utilized each of these domains.

The methodology included three steps. First, the team surveyed application performance for an overall view of the problem; this information is important because it provides an idea of where to explore for problems. Then, they took a profiling tool and separated the application into its constituent components and attributed time and resource usage to each of these components. Finally, they traced small portions of the application identified by profiling to understand the path of execution. Tracing gives a more precise picture of a sequence of application events and how they interact with each other.

The expectation for the milestone was that the HPCART would be able to use the plethora of performance tools at Sandia, but many of the tools are unable to handle production codes running in production environments. Almost all tools had issues, some of which were severe.

The milestone had major benefits despite the difficulties encountered with the tools. They included an improved understanding of the tools' feature set, enhanced documentation for the tools, and insights for future tool-development activities. One of the key results for the milestone was a set of recommended tools for each performance analysis activity, along with several reported issues and gaps in feature coverage. In addition, the Sierra/Aria and ITS code teams received performance results from a large number and variety of runs with specific instructions on how to use the tools with their codes. The data and recommended instructions will help developers mitigate problem areas and trends to enable their codes to run more efficiently.

Table 1

Profiler	Chama	Sky Bridge	Serrano	Mutrino	
				Phase 1	*Phase 2*
Profiler A		A, I	A, I		
Profiler B	A, I	A, I	A, I		
Profiler C				A, I	A, I
Profiler D	A, I	A, I	A, I	A, I	
Profiler E	A	A, I	A, I		
Profiler F	A, I	A, I	A, I		
Profiler G			I		
Profiler H			I		

A	Denotes profiler was utilized on Sierra/Aria
I	Denotes profiler was utilized on ITS
(grey)	Denotes profiler is unavailable
(red)	Denotes profiler is installed but issues were found

List of all performance profilers and relevant platforms that were utilized.

HPCART members often partner with Sandia researchers on an important large-scale simulation of interest on a new, possibly unstable platform during the early-access period of a campaign. Embedding an HPCART member with the analyst team allows HPCART to coordinate integration with new platforms: finding the issues, communicating to platform integration/operations teams, and developing temporary work-arounds to keep the simulation moving forward, while working with sub-system and integration experts long-term to shepherd the "final" corrections into the production system. The goal is to maximally stress the system in the very early stages (preferably at full-scale) in a planned, organized manner and find/correct the problems that surface. This enables Sandia researchers to have a better user experience.

THE ALLIANCE FOR COMPUTING AT EXTREME SCALE

In the early Alliance for Computing at Extreme Scale (ACES) Cielo campaign period, HPCART partnered with Sandia researchers on a large-scale simulation called "Lightweight, Blast Resistant Structure Development" for Critical Asset Protection & Security. Figure 1 shows an image of ACES Cielo as deployed at Los Alamos National Laboratory (LANL) from 2011 to 2016. The simulation helped Sandia researchers understand the response of structures under severe blast loading conditions so that the robustness of these structures could be improved. The simulation contained 3.2 billion cells, which represents one of the largest ParaView data visualizations of this type in the world. In the simulation, a structure composed of sheet metal is loaded by an explosive blast. Figure 2 shows the velocities of various components of the structure as they are torn apart. Accurately resolving these thin parts required a very fine mesh resolution, which can, at present, only be achieved using very large-scale computing platforms such as Cielo.

For over 3 months, the HPCART documented and tracked platform issues while running the simulation of the first ACES Cielo computing campaign. HPCART partnered with developers of CTH a shock physics code, to evaluate performance. HPCART was responsible for apprising researchers of simulation progress as well as tracking issues and impediments. This led to an improved "user experience" and progress, despite the many issues encountered over that period. The platform integration/operations team implemented fixes necessary to enable a successful long running simulation at 32 K cores, which was subsequently expanded to 64 K cores, achieving the largest runs to date for this class of problem.

Figure 1

New Mexico ACES Cielo Platform 2011-2016.

Figure 2

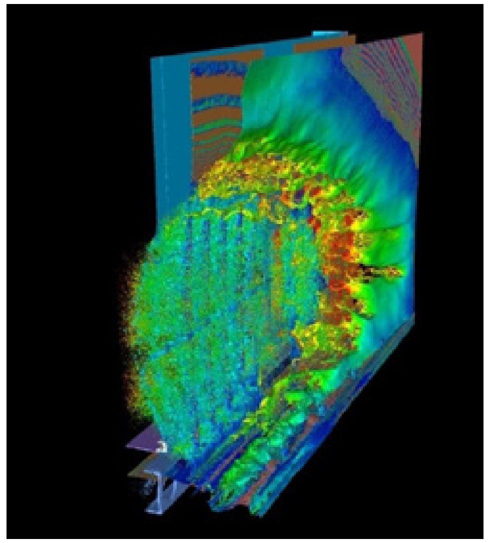

Multilayered thin-walled structure torn apart by an explosive blast.

Spectre and Meltdown are names given to different variants of the same fundamental underlying vulnerability that affects nearly every computer chip manufactured in the last 20 years. If exploited, attackers could access data previously considered completely protected. Technically, there are three variations on the vulnerability, and two of those variants are grouped together as Spectre and the third is dubbed *Meltdown.*

The underlying vulnerability involves a malicious program which gains access by exploiting two important techniques used to speed up computer chips, called *speculative execution and caching.*

Shortly after their discovery hardware vendors pushed out software patches to deal with these weaknesses. The patches generally mitigate the vulnerabilities, the downside is that these features were designed to improve system performance, and working around them slow systems down with initial reports of performance hits up to 30 percent.

Table 2

	Activity	Wall Clock Contribution	Activity Performance Loss	Overall Workload Loss
SNL App 4 Prod. Workload	MPI	76.4%	19.2%	14.7%
	I/O	0.1%	30.7%	0.0%
	Memory Allocation	0.8%	79.1%	0.7%
	Compute	22.7%	0.0%	0.0%
SNL App 6 Prod. Workload	MPI	40.2%	32.0%	12.9%
	I/O	0.7%	-37.9%	-0.3%
	Memory Allocation	0.5%	38.5%	0.2%
	Compute	58.6%	0.0%	0.0%

Impact from Spectre/Meltdown patches of Message Passing Interface (MPI), file input/output (I/O), memory allocation (e.g., malloc), and compute activities for two production application workloads.

HPCART was asked to quantify the production impact from the slowdown due to these patches. Mission-critical work would be prioritized to account for the loss of potential CPU cycles. Within two weeks, HPCART was able to pull together a comprehensive briefing for DOE that looked at the impact of the patches on seven of Sandia's production applications and a micro-benchmark suite. HPCART performed strong scaling studies of the seven production applications, measuring their fine-grained usage of the many central processing unit (CPU) hardware components, and correlating the application slowdown to the affected CPU components. A performance impact mitigation strategy was simplified for a wide audience as, "reduce the node counts of simulations to increase their memory footprint per node."

Figure 3

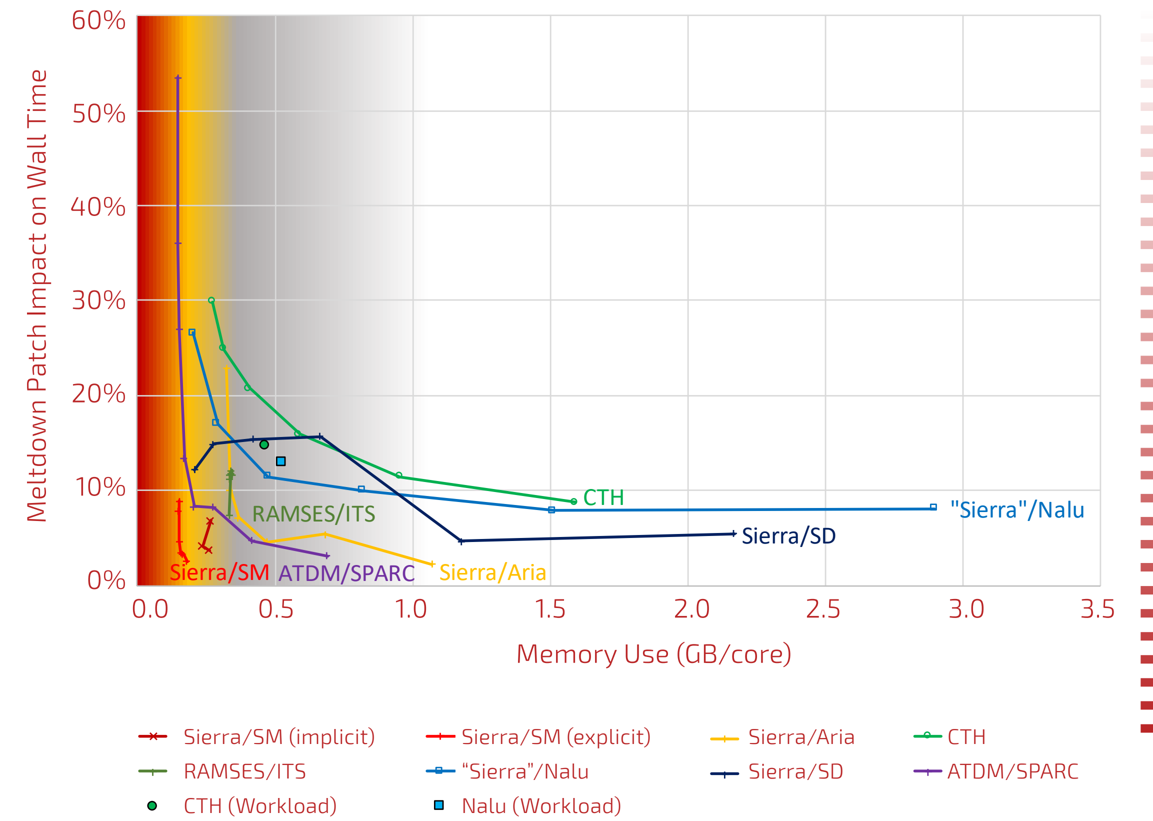

Strong scaling impact from Spectre/Meltdown patches of seven production applications.

LDXPI

LDPXI (pronounced "load pixie") is named for its method of operation—it uses the Linux LD_PRELOAD facility to inject extremely light-weight eXecutable Instrumentation into a program. It was invented and developed at Sandia and is designed to give a first glimpse into how an application performs and how it uses cache, MPI networks, memory, etc. This view can be very useful to software developers and analysts attempting to triage a performance problem, system architects trying to understand what kind of cluster will benefit the lab, and others. It is widely used by the HPCART for cases such as a level-two milestone to assess applications on HPC systems, to triage performance problems when users have issues, to survey applications to find out what portions of the hardware are heavily used such as network or memory performance, and in Spectre/Meltdown to measure what was different before and after the patches were installed.

LOOKING FORWARD

SCOTT COLLIS
DIRECTOR, CENTER FOR COMPUTING RESEARCH

Sandia and our partners at national laboratories, academia, and industry are hard at work on the next milestone in supercomputing capability–exascale computers that will achieve exaflop performance (a billion billion floating point calculations per second). Sandia is a core member of the Exascale Computing Project, a joint Department of Energy Office of Science and National Nuclear Security Administration effort to develop and deploy applications and software technologies that fully exploit exascale computing. The first U.S. exascale systems are planned to be available starting in 2021 and Sandia is deeply involved with cutting-edge exascale software technologies such as the Kokkos programming model and exascale applications in climate, combustion, wind energy, and materials. Simultaneously, Sandia is helping to define the frontiers of computing beyond exascale.

The next 10 to 15 years will see exascale systems become the norm in supercomputing, enabling breakthroughs across a wide range of science, energy, and security applications. Simultaneously, future computing systems will be better optimized for data analytics and machine learning, helping us make sense of the flood of data from social media, embedded sensors, the Internet of Things, and other sources. To efficiently convert data into actionable decisions requires innovation in computing and a holistic perspective on the computer-human interface, with codification of human intuition through brain-inspired computing. We must also research means to distribute computing to data and sensors, often requiring dramatic reductions in electric power consumption. Finally, the National Quantum Initiative brings the prospect of significant new U.S. investment in quantum information science with Sandia poised to play a pivotal role in this arena.

Sandia is also helping shape and understand the impending shift to humans enhanced with augmented and virtual reality technologies. We will soon be literally and routinely connected to our devices, potentially disrupting socio-economic systems and creating new risks and opportunities in national security. How will we stay abreast of augmented human reality? What are the roles of computing research in facilitating or managing these changes? These are some of many research questions Sandia is exploring.

Meanwhile, traditional computational simulation will remain essential as we increase the use of computing for the design, certification, and life-cycle management of engineering systems. Sandia is helping the U.S. computing industry move beyond Moore's law with innovations in architectures, devices, programming models, and algorithms. Our goal is to advance modeling of physical and engineered systems by building on Sandia's uniquely high-consequence mission to produce full-system, full life-cycle computational models. However, we must not lose sight of users' need for shorter turnaround time, improved usability, and more robust, higher-confidence results.

The breadth of these challenges is immense. Computing research is central to Sandia's solutions, and to succeed we must partner across the national labs on technologies and across mission users to build computing and information capabilities while leveraging and expanding our outreach to universities and computing industry. Our path is challenging, but it also presents opportunities and requires us to imagine new technologies and new ways to manage complexity through creative partnerships. We are excited and energized by these opportunities and hope you glimpse this future in this year's report.

WHETSTONE:

TRAINING SPIKING DEEP NEURAL NETWORKS FOR NEUROMORPHIC PROCESSORS

Brad Aimone, William Severa

APPLICATIONS REQUIRING SIGNAL PROCESSING and sensor analysis increasingly rely on artificial neural networks (ANNs), inspired by cortical processing within the human brain, to make inferences and to identify and discriminate information. The advent of low-power neuromorphic hardware has permitted the use of deep-learning capabilities in remote or constrained environments. These brain-inspired chips communicate using spikes and achieve much of their power advantage by communicating using a binary value of one or zero between neurons rather than continuous values. However, while ground-breaking, most deep learning networks currently do not use this kind of binary communication and tuning a conventional network to the spiking functions of a neuromorphic chip would require considerable customization.

Whetstone, a software package for training conventional ANNs to behave as spiking neural networks, has been developed by Sandia National Laboratories. Whetstone is unique among spiking neural network approaches in that it directly uses existing deep-learning neural network tools to make neural networks compatible with neuromorphic hardware.

OVERVIEW

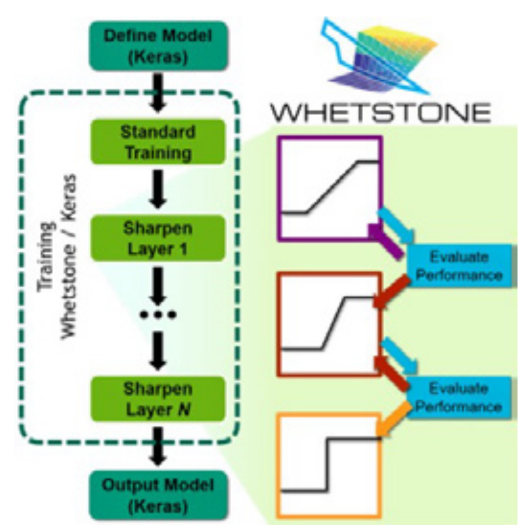

Whetstone leverages and integrates with popular deep-learning tools for easy integration into workflows. Network adaptation via the Whetstone library happens gradually and in parallel to the standard deep-learning training phase.

Whetstone works with many new chips and does not require the application developer to master their programming, thus minimizing the barrier of entry to neural network application developers who may have no familiarity with spiking algorithms or neuromorphic hardware.

Whetstone operates by generating a spiking neural network directly in the training process of deep-learning neural networks. Neural networks begin the training process as standard neural networks, with full-precision activation functions. During the training process, the network is morphed to a binary activation neural network. This process occurs gradually and one layer at a time. Each layer is progressively “sharpened” to a spike-like activation function (thus, the name “Whetstone”) in a manner that preserves network accuracy greater than or equal to that of full-precision networks.

Scan using SNLSimMagic App

DEEP NETWORK

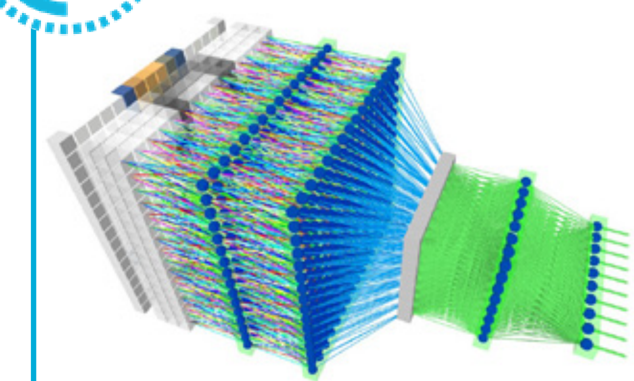

Neural networks, particularly Convolutional Neural Networks as shown, are built by configuring thousands or millions of parallel processing nodes (neurons). A hierarchy of neurons can provide complex problem-solving ability. Whetstone produces neural networks preconfigured for low-power, binary activation or neuromorphic computing hardware.

While there are methods for converting conventional ANNs to spiking neural networks (SNNs), these transformations often diminish the benefits of spiking. In contrast, the new approach offered by Whetstone is to train an ANN while producing an SNN. Specifically, if training procedures can include the eventual objective of low-precision communication between nodes, the training process of an SNN can be nearly as effective as for a comparable ANN. Further, the binary nature of communication offered by Whetstone makes the resulting networks compatible with both spiking neuromorphic hardware and conventional hardware.

BIOLOGICAL NEURONS

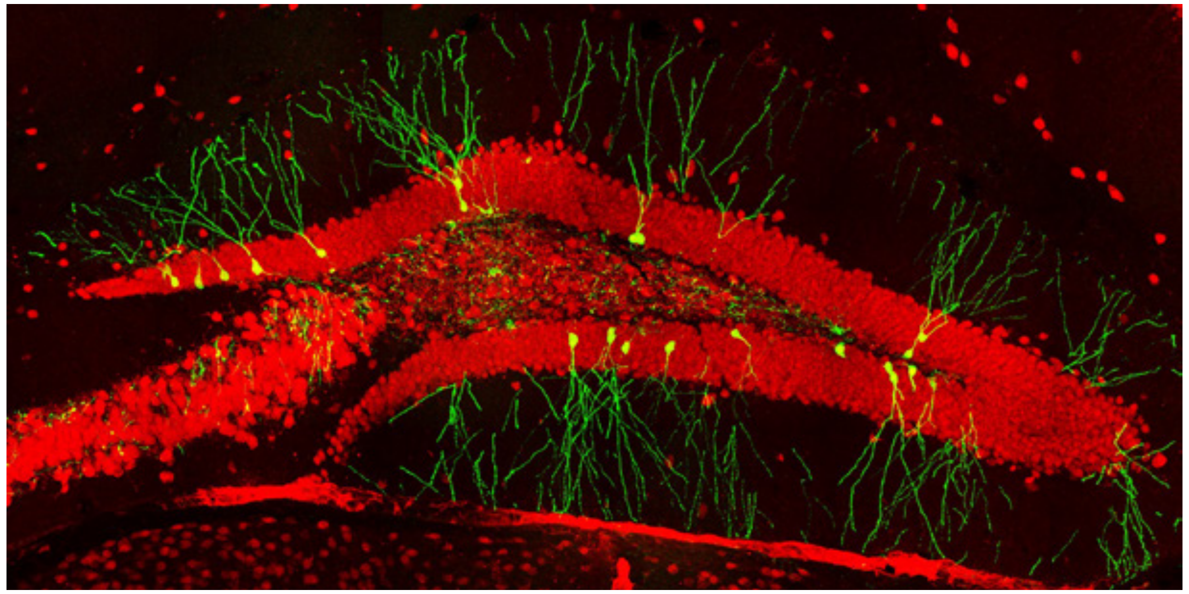

In contrast to most artificial neural networks, large-scale biological neural networks use a discrete signaling method neuroscientists call "spikes." This communication strategy may be responsible for the brain's incredible energy efficiency and serves as inspiration for the Whetstone network.

Despite increases in efficiencies and renewable energy there is a growing need to reduce the energy footprint in computing, even when using already thrifty devices such as neuromorphic hardware. Devices enabled with artificial intelligence that send calculation requirements to data centers will contribute their share to the more than 2,500 terawatt hours of projected consumption by data centers worldwide. However, neuromorphic hardware using Whetstone functions has an estimated efficiency 100 times greater than that of conventional hardware. Whetstone gives neuromorphic hardware a path by which it can be used for commercially impactful computing, resulting in a potential future that can radically reduce the energy footprint of advanced artificial intelligence operations.

The Sandia Whetstone software package provides two distinct options compatible with most feed-forward network topologies. Once these simple modifications are in place, the sharpening process can be largely automated by the software. The key to Whetstone's functionality is its ability to generate spiking neural networks into an existing network structure. This innovation opens a world of possibilities for the deployment of computing resources in places and circumstances not previously imagined while unlocking the potential for energy conservation in the future of computing.

POWER API:
THE POWER TO SAVE

Ryan Grant

DATA CENTER POWER USAGE is growing rapidly; in the next decade, the U.S. Department of Energy predicts data center energy use will rise to comprise 20% of all energy consumption in the U.S. Using the technology of Power Application Programming Interface (API) creates energy savings for data centers that could add up to $8.4 billion a year by 2030. Power API makes this possible by providing a standardized method of measuring and controlling power on modern computer systems, from individual laptops to the largest supercomputers in the world. Power API allows users and administrators of the world's largest supercomputers to measure power consumption, to better understand the behavior of applications and other software systems running on their hardware, and then to optimize them for efficient energy consumption. This means doing the same amount of work but using less energy, or getting more work done without using any extra energy.

Power API provides a scalable, distributed interface that works with all types of hardware, allowing tools and applications to be developed that can save energy while being used across many different systems; for example, an energy-optimized code for an Intel x86 system with Power API can seamlessly translate to an ARM or IBM power based system without modification (e.g. Sandia's Astra and Oak Ridge National Laboratories' Summit). Even within a single system, there are multiple energy/power controls, one for each component that comprises a computer (CPU, memory, storage, network). Power API provides a single interface for all of these components.

Power API also serves as a framework for optimizing runtimes that are developed for systems and job coordination software. Users can utilize a single portable interface with custom optimizations and middleware transparently operating behind the scenes, improving overall system efficiency. While this underlying system software may improve between different generations of systems, the difficult-to-develop power management code base remains the same and benefits from these improvements. In this manner, Power API is the first portable efficiency software interface for computing systems. Many successful APIs in the past have been performance portable, but never before has any product provided portable energy efficiency. Power API is part of the HPC Open Compute Project (OCP) software stack and is available on commercial systems running OCP. It is also a part of the Tri-lab Operating System, making it available on many National Nuclear Security Administration computing systems.

A truly useful power measurement and control interface must be based on a deep understanding of measurement accuracy and frequency. Power measurement has multiple measurement layers, with only the top layers exposed to the user. For example, a device might sample the hardware and take measurements very quickly, but expose averaged values to the user less frequently to avoid delivering large volumes of data. A measurement natively sampled at 1 Hz might not be fast enough to capture the shorter power fluctuations that are important to understanding the system's actual power use. However, if the measurements are the average of the underlying 10 kHz samples collected and aggregated by the measurement hardware, the resulting measurements may be accurate enough for some use cases. In this case, the hardware's power is sampled, but the results of these samples are only available intermittently. Exposing the measurements is a costly operation for high frequency power measurements, but averaging them is a viable alternative. To accommodate these differences, Power API associates metadata with each measurement point to inform the user about the underlying sampling methodology and its accuracy. The user can then determine whether the measurement capability is an effective choice for a particular use case.

Power API is specified primarily as a C API because C is the preferred language for low-level software for data centers and supercomputing systems and is universally supported. However, alternative language bindings, such as Python, are provided. Power API implementations can use whatever language is most convenient; for example, Sandia's reference implementation is written primarily in C++, with user-visible C interfaces provided externally.

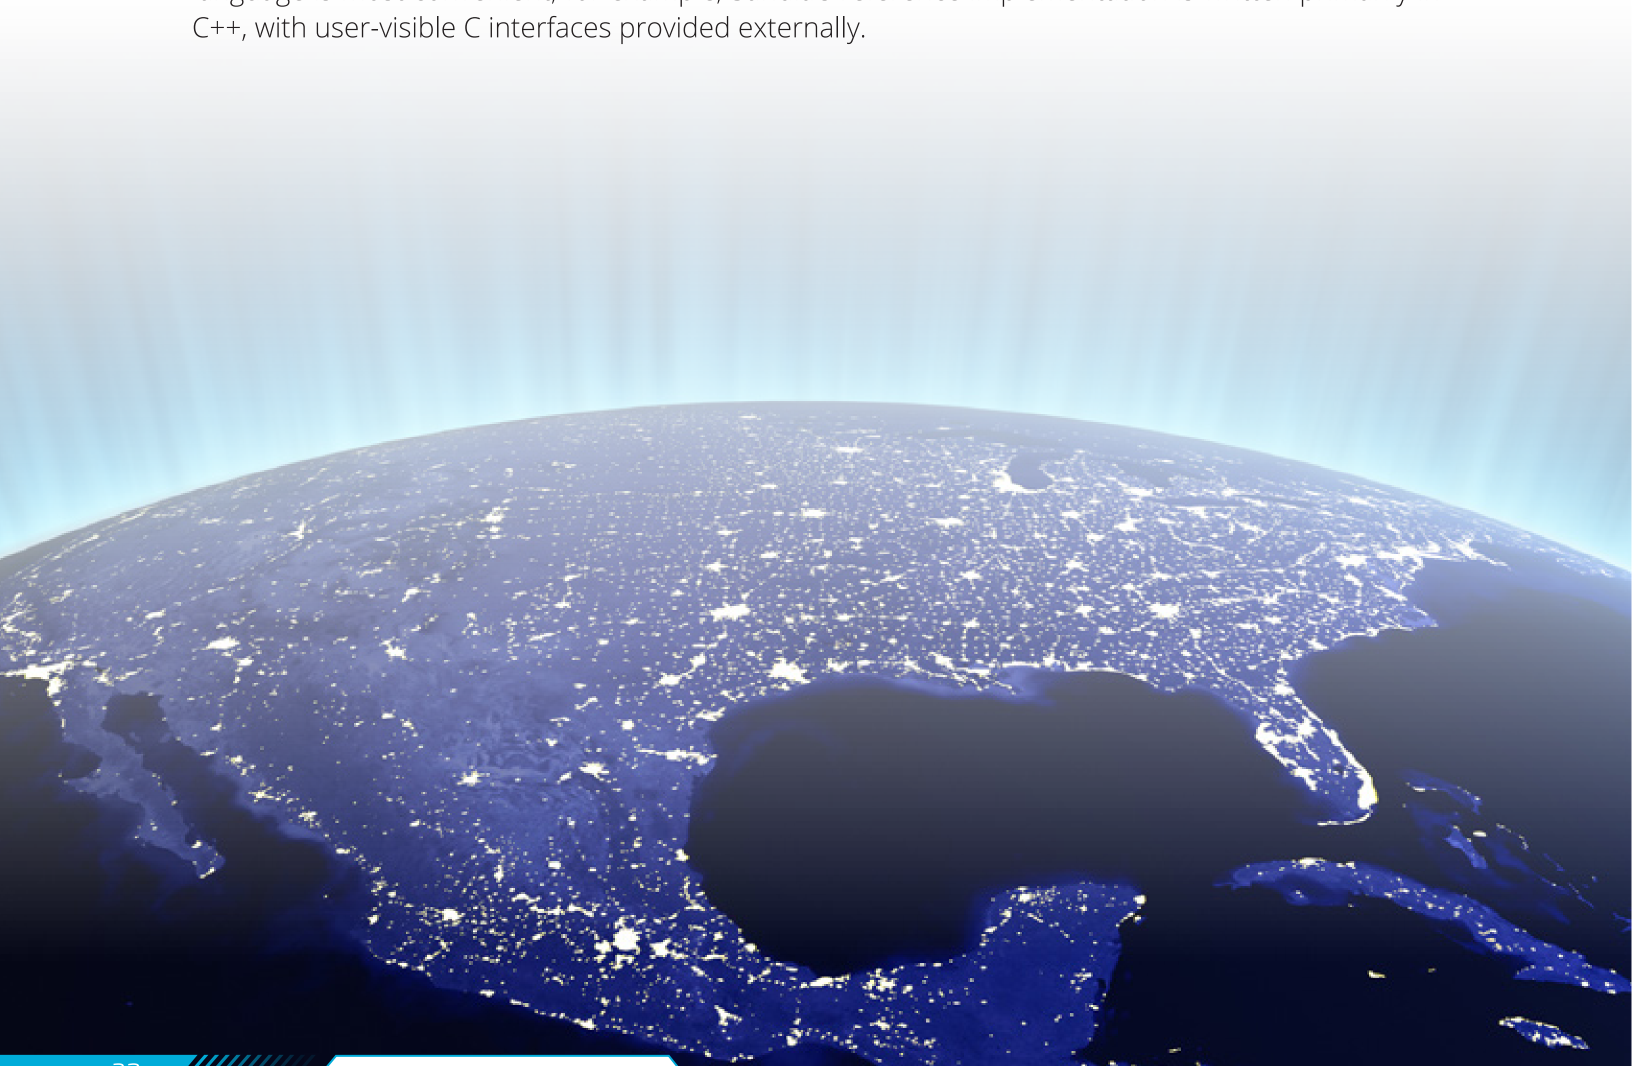

The API itself has great implementation flexibility, but current implementations work using plugin architectures that allow for easy, low-level hardware plugins to be written that support several different underlying hardware monitoring and control systems simultaneously. Power API has several interfaces designed specifically for different use cases; called "roles" in Power API terms. These roles were identified through an extensive use case study for power management that pre-dated Power API. Some of the possible Power API roles are: Application, Monitor and Control, Operating System, User, Resource Manager and Administrator. Some of these roles are meant to serve software or hardware, while others are meant to interact with a human actor. For example, an Application can use Power API to adjust its power usage during execution and can read sensors for feedback on current system operating conditions. Likewise, an Operating System can use Power API to balance power for all applications running on a server for a machine local power management scheme. The resource manager functionality allows power management at a job and whole system level for computer clusters, ranging from small two node systems to the largest supercomputers in the world.

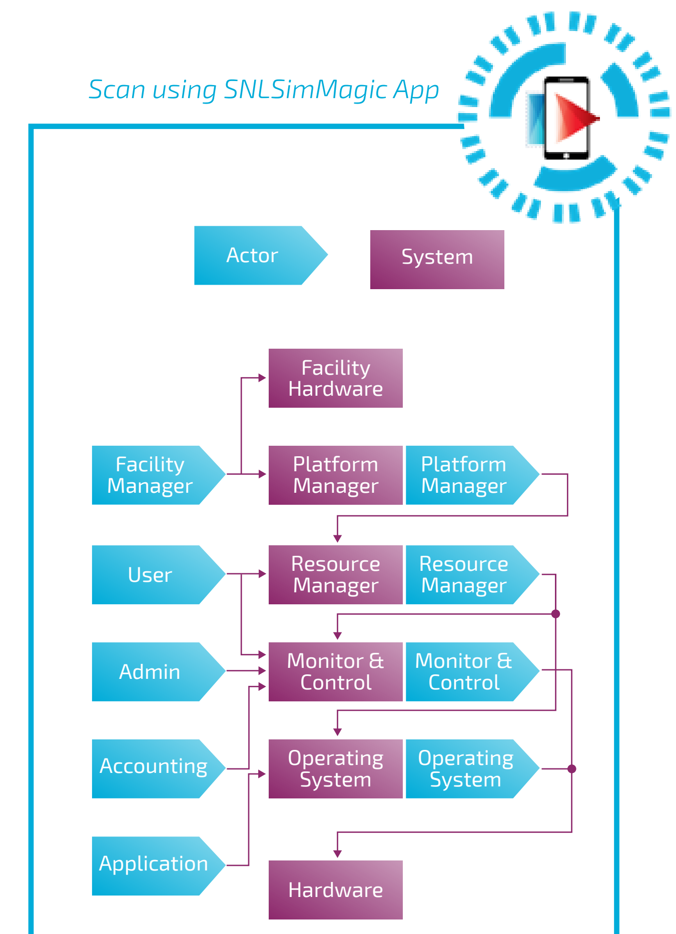

Basic relationships between Power API interfaces.

Power API can save energy on computing systems in the order of 5-15% of total system energy by using the power controls on existing platforms. Given budgets in the order of tens of millions of dollars per year for DOE supercomputers, this saves millions of dollars per year throughout the DOE supercomputing centers. Saving even 10% of wider U.S. data center consumption (20% of the total power consumption in the U.S.) means Power API can save on average 2% of all electrical generation capacity in the U.S. per year with no reduction in work performance. According to World Bank statistics, this is equivalent to 83.3 billion kilowatt hours (83.3 kilowatt hours) or approximately $8.4 billion per year at average U.S. consumer power prices. That is the equivalent of doubling the efficiency of all solar generation panels in the United States as of 2017. In terms of CO2 emissions, it is equivalent to taking 20.6 million cars off of the road. Power API is being deployed around the world, providing even more impact to energy savings and carbon reduction globally.

QUANTUM COMPUTING

John Aidun

THE UNIQUELY CHALLENGING technical problems arising in stockpile stewardship have long driven advances in computational science. They are now beginning to propel the development of quantum information processors (i.e., "quantum computers," QCs), to perform quantum simulation.

Quantum simulation refers to a class of algorithms for solving the many-body Schrödinger equation using a QC and is the most promising application of QCs to stockpile stewardship. The Schrödinger equation governs the electrons and ions that comprise chemical and materials systems, as well as the quarks and gluons that comprise nuclei. Classical solutions to the Schrödinger equation are already used to generate materials models for engineering codes. However, these classical approaches force a choice between those with unquantifiable accuracy, which have a computational "cost" that increases polynomially with the number of particles in the simulation, and those with systematically improvable accuracy at a cost that increases exponentially with system size. Quantum simulation will achieve an exponential advantage of systematically improvable accuracy at a polynomial cost; a QC running a quantum simulation algorithm will generate more accurate materials models with quantifiable uncertainties. It is not uncommon for material properties to be poorly constrained or unavailable for extreme conditions, due to the cost and difficulty of performing experiments. These truly-first-principles calculations might be used to provide values for such poorly known or experimentally inaccessible quantities for a variety of essential properties of materials that are vital to stockpile stewardship.

The impacts of such a capability will be broad and may help resolve outstanding problems in chemistry, materials science, and nuclear physics. However, the number of qubits and quantum operations that are required for these calculations to surpass classical supercomputers require error rates dramatically lower than those expected of current or near-term qubit technologies. Hence, mission-relevant calculations will not only need many more qubits than are planned for near-term QCs, but also much more reliable qubits than currently exist.

Sandia's newly initiated work in the Department of Energy/National Nuclear Security Administration (DOE/NNSA) Advanced Simulation and Computing (ASC) program has two thrusts: adapting a quantum simulation algorithm to calculation of key stockpile materials and estimating the QC resources required to reliably execute these algorithms. This effort is complemented by four new projects at Sandia funded by the DOE/SC Office of Advanced Scientific Computing Research (ASCR), including two involving QC testbeds.

The Quantum Scientific Open User Testbed (QSCOUT) project has set out to build a QC testbed comprising 5 to 15 trapped-ion qubits. The low error rates and full connectivity provided by trapped-ion qubits make the QSCOUT testbed an ideal system for testing algorithms on a real-world processor. But the more immediate value of QSCOUT is to provide data on speed and fidelity of quantum operations ("gates") on real qubits. This data will inform the ASC quantum simulation project, improving the precision in the QC resource estimates.

The other new ASCR project, Q-PERFORMANCE, will produce techniques to evaluate emerging QCs, like QSCOUT, and assess every aspect of their performance to provide understanding of the quantum processor's performance and track the progress of improvements. It will develop cutting-edge models and assessment protocols for primitive properties of quantum processors, invent and deploy holistic benchmarks that capture high-level performance from multiple perspectives, and focus particularly on application-specific performance, to predict suitability of a processor (or its successors) for DOE mission-relevant tasks.

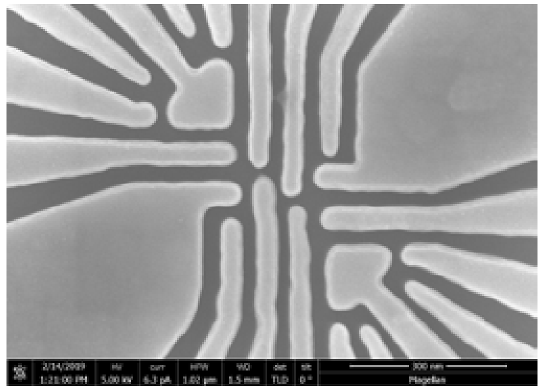

SEM image of surface electrodes for a double quantum dot structure for electron spin qubits in silicon.

SEM image of surface electrodes for a quadruple quantum dot structure for electron spin qubits in silicon. Red dots indicate the locations of the four quantum dots.

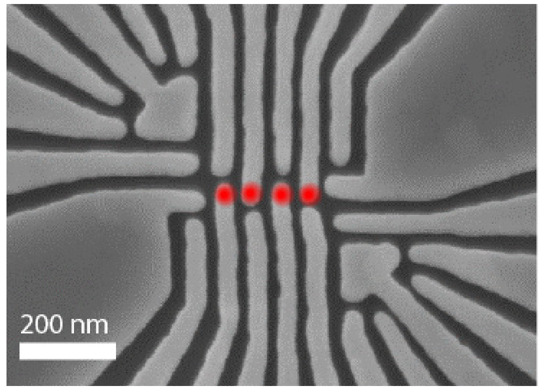

ECP SUPERCONTAINERS:
FROM TESTBED TO PRODUCTION

Andrew Younge

CONTAINER-BASED COMPUTING has revolutionized the way many industries and enterprises develop and deploy software and services. Containers enable developers to specify the exact way to construct a working software environment for a given code. This includes a base operating system configuration, system libraries, third party libraries, environment settings, and how to compile their application at hand. As a lightweight collection of executable software, containers encapsulate the software ecosystem needed to run a single specific task on a computing system, aiding in the total application development and deployment time.

For these reasons, containers have caught the attention of high performance computing (HPC) researchers. Traditionally, the boutique nature of HPC systems has led to significant development time and resource usage for applications to adapt to new hardware and system software, requiring heroic efforts to deploy on each new system. This often leads to increased time spent debugging and small-scale manual testing, which focuses on only the immediate HPC deployment at hand. Time spent on these efforts could be better spent running applications at scale, rather than reproducing build operations over again.

What is the ECP?

The Exascale Computing Project (ECP) is chartered with accelerating delivery of a capable Exascale computing ecosystem to provide breakthrough modeling and simulation solutions to address the most critical challenges in scientific discovery, energy assurance, economic competitiveness, and national security.

This role goes far beyond the limited scope of a physical computing system. ECP's work encompasses the development of an entire Exascale ecosystem: applications, system software, and hardware technologies and architectures, along with critical workforce development. With the goal to launch a U.S. Exascale ecosystem by 2021, the ECP will have profound effects on the American people and the world.

The ECP is a collaborative effort of two U.S. Department of Energy organizations–the Office of Science (DOE-SC) and the National Nuclear Security Administration (NNSA). The Supercontainer project is part of the ECP effort.

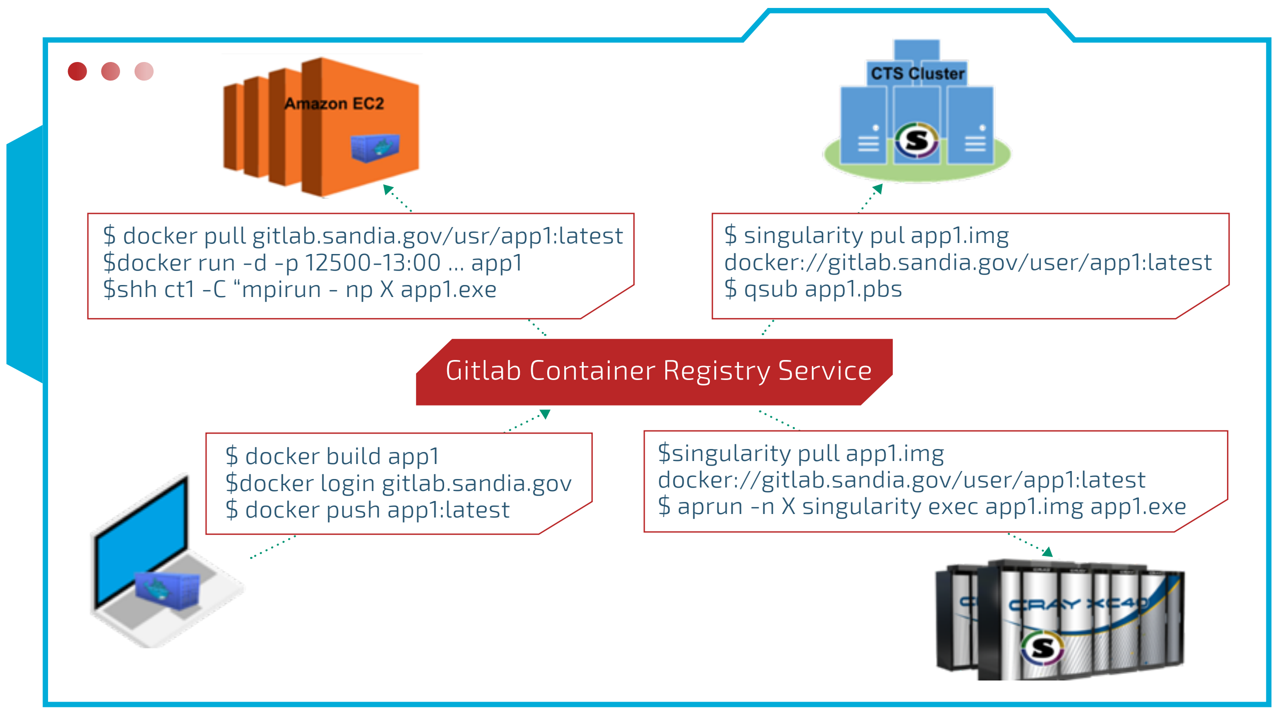

The idea of using containers to import a pre-determined program directly to an HPC resource through new development and operations (DevOps) mechanisms is exciting news. With containers, researchers and developers can write a program on their laptops, then can run that same container on an HPC system with no massive conversion required, and no wasted time on the HPC that would be available for other operation runs. Everything needed to run a program on that laptop was included in the container to run on the HPC.

However, different hardware and system architectures have led to increased complexity when deploying new software on large-scale HPC platforms. With the reality of Exascale computing on the horizon, this complexity and extreme scale will only increase. While containers provide greater software flexibility, reliability, ease of deployment, and portability for users, there are still several challenges for using containers at Exascale that must be addressed.

Researchers at Sandia, in collaboration with Los Alamos National Laboratory, Lawrence Berkeley National Laboratory, and the University of Oregon, have joined forces to enable "Supercontainer" deployments from laptops to Exascale computers. The Supercontainer project looks to accelerate adoption of container technologies for Exascale, and ensure that HPC container runtimes will be scalable, interoperable, and integrated into Exascale supercomputing across DOE. This multi-pronged effort includes creating portable solutions (containers that can run on different platforms), working with multiple container runtimes (not focusing on just one approach), and working for multiple DOE facilities at multiple scales (different supercomputers run different hardware, and the ECP Supercontainers project will work with them all).

In addition to greater flexibility and agility for users, other benefits include the potential for greater software portability between multiple users and different systems. The goal of moving an HPC application container from initial development on a laptop and scaling to Exascale presents a powerful and tractable paradigm shift in HPC software development. Sandia's team of container developers is hoping their next iteration of Supercontainer development answers these challenges.

NEURAL EXPLORATION & RESEARCH LAB

Craig M. Vineyard

INSPIRED BY the computational feats brains perform, researchers have long explored creating hardware and algorithms mimicking the neural principles of the human brain. Interest in neuromorphic computing parallels much of the progression of the field of computing at large; however, general purpose computing advances achieved in parallel to the pursuit of neuromorphic approaches have historically overshadowed neural-inspired computing.

Recent algorithmic advances in machine learning, artificial intelligence, and deep learning have spurred interest in accelerating neural network algorithms, but computer performance laws are being reached with device technologies approaching atomic scales. Computing can no longer rely upon scaling to enable advances. Moore's Law combined with Dennard scaling can no longer enable exponential improvements in computing power. To address these limitations the field of neuromorphic computing is pursuing novel innovative algorithms and architectures employing different computational principles. Rather than computing with gates as the core computational unit, neuromorphic approaches use a neuron as the fundamental computational unit. Computation among neurons is intrinsically parallel and often event driven. Some of the more biologically inspired approaches communicate via discrete spikes (single bits) rather than complex packets. And some neuromorphic approaches explore employing novel devices or computing in analog. Effectively, neuromorphic computing has the potential to bring advanced computational power to resource constrained environments which otherwise might require a supercomputer to compute. Additionally, enabling advanced machine learning capabilities in HPC has the potential to provide revolutionary computational ability.

The Neural Exploration and Research Lab (NERL) at Sandia enables researchers to explore the boundaries of neural computation. The research conducted in the lab evaluates what is possible with neural hardware and software for both the advancement of basic research and national security applications. The lab provides a testbed facility for comparative benchmarking, new architecture exploration, and neural algorithm development.

A key tenet of the research conducted at NERL is the interplay between hardware, software and algorithms. These three components must be designed cooperatively for effective, impactful adoption of a new technology. As such, NERL models much of its work on the successful history of traditional computing architectures and HPC platforms, from working to establish meaningful benchmarks to systematic analysis of architecture/algorithmic design.

Many compelling neuromorphic architectures exist using technologies of today, but novel devices, algorithms, and architectures are also being explored for future generation neuromorphic architectures. Design tradeoffs focus upon different features making them better suited for different applications and enable different performance characteristics. However, interest and progress in neural-inspired computation are generally linked by two key components (1) energy costs limit compute capability and (2) data-driven algorithms provide data insight. Consequently, neuromorphic computing approaches are showing promise to impact a wide spectrum of computing domains from high performance computing to resource constrained environments for remote sensing.

High performance computing harnesses massively parallel computational infrastructure to enable advanced simulations and analysis capabilities central to scientific computing. Unfortunately, energy consumption (and the ability to cool large machines) becomes an increasing burden as the efficiency of processors stagnate and the scale of these machines increase. As such, machine learning and neural network approaches are being increasingly sought-out for in situ analysis, fast surrogate models, and system maintenance. Accelerator hardware and graphics processing units are increasingly being incorporated into HPC infrastructure distributing these technologies across compute nodes, and neuromorphic approaches may usher in the next wave of "smart HPC."

Remote sensing and extreme environments present a unique and critical algorithm and hardware tradeoff due to extreme size, weight, and power operational constraints. Consequently, in many applications, systems favor centralized computation over remote computation. However, in some real-time systems (e.g. a Mars rover) latency and other communication bottlenecks force on-board processing. With traditional processor performance at a plateau, researchers look to brain-inspired, neuromorphic architectures to enable future capabilities, such as event detection/tracking and intelligent decision making.

ASTRA: TO THE STARS

James Laros, Simon Hammond and Kevin Pedretti

IN THE FALL OF 2018, Sandia's brand-new data center in Building 725E became host to Astra, bringing to life a new paradigm in high performance computing (HPC), even employing a new hardware platform: Hewlett Packard Enterprise's (HPE) Apollo 70. Astra's development was the result of the Astra Supercomputer Team, who worked on an aggressive schedule to obtain and deploy Sandia's fastest computer to date. Because the facility and the platform were both new, the Astra team had multiple issues to resolve. Working closely with HPE and Sandia Facilities, the team optimized thermals for Astra and 725E and developed an early HPC-capable software environment for the new system while working to stabilize first-of-a-kind hardware.

Astra, with 2,592 compute nodes, uses the ThunderX2 processor, developed by Marvell, as a fully featured, server-class Arm design that provides exceptional memory bandwidth contributing to a more optimized balance between compute and data throughput requirements. Through their fiscal year 2019 milestones, the Astra team is helping answer the vital question of how well the peak performance of Astra's architecture translates into real performance for National Nuclear Security Administration (NNSA) mission applications. For example, the processors used in Astra feature twice the memory channels of the current Trinity production platform developed jointly by Sandia and Los Alamos National Laboratory (LANL). As a result, Astra can deliver almost twice the memory access rates, which accelerates some of the performance critical components for Sandia workloads by more than 85%.

Astra is the world's largest, as well as the first petascale, HPC system based on Arm architecture. In October/November of 2018, Top500 High-Performance Linpack (HPL) and High-Performance Conjugate Gradients (HPCG) runs took place on a partial system to demonstrate initial system operations. By November 2018, on the Top500 list of the fastest supercomputers, Astra placed 204th for HPL, with 1.529 petaflops on 2,238 compute nodes and 36th on HPCG, with 66.94 teraflops on 2,238 compute nodes. For Top500's latest list, released June 2019, Astra placed number 156th at 1.758 PFlops/s and 29th HPCG at 90.92 TFlop/s, representing a 36% improvement from November 2018.

The Astra supercomputer is the first deployment of Sandia's larger Vanguard program. Vanguard is tasked to evaluate the viability of emerging HPC technologies in support of the NNSA's mission to maintain and enhance the safety, security and effectiveness of the U.S. nuclear stockpile.

In conjunction with Astra, Sandia is leading an effort through a collaboration with the NNSA's Tri-Labs (Sandia, LANL, and Lawrence Livermore National Laboratory [LLNL]) to develop a well-integrated and optimized HPC software stack for Astra and future Advanced Simulation and Computing (ASC) platforms. Astra's stack, called ATSE (Advanced Tri-Lab Software Environment), provides a set of foundational software components needed to build and execute ASC mission computing workloads efficiently at scale.

Recognizing there is no one-size-fits-all software stack, ATSE is being designed with the flexibility needed to integrate with multiple base operating system layers, including the Tri-Lab Operating System running on Astra, and to allow vendors to contribute value-added components such as optimized communication libraries and compilers. Assembled together, each release of the ATSE stack is being packaged in native, container, and virtual machine formats to provide end users with reproducible build environments that can be deployed across multiple platforms, from supercomputers, to clouds, to laptops. An overall goal is to leverage ATSE to build a broad, sustainable, and multi-vendor HPC software stack ecosystem for Astra and other advanced architecture ASC platforms.

SYSTEM PERFORMANCE OF ASTRA OVER EXISTING TRINITY PLATFORM

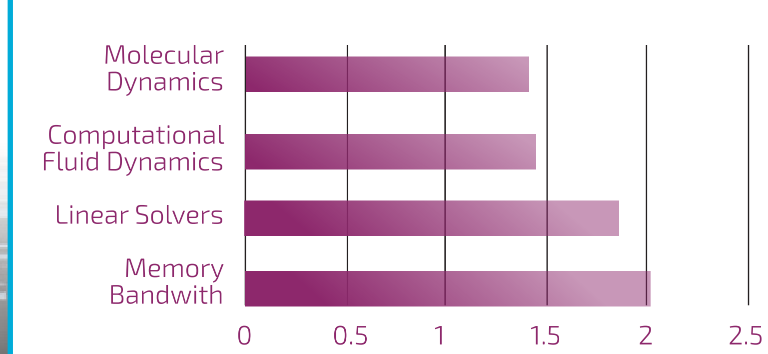

Sandia is working with LANL and LLNL to collaborate on using Astra. Applications cover particle transport modeling, hydrodynamics, and laser implosion modeling for the National Ignition Facility.

SLYCAT VIDEO ENSEMBLE ANALYSIS:
VIDEOSWARM

Patricia Crossno

GIVEN THAT EXASCALE ARCHITECTURES are expected to constrain the size of simulation outputs, researchers are investigating the use of image and video outputs as an alternative to conventional finite element meshes. Additionally, scientists are increasingly exploring problem solution spaces by iteratively running their simulations with varying parameter inputs, generating sets of related simulations known as ensembles. Ensembles contain large, complex sets of results, which consist of high-dimensional, multi-variate, temporally-varying data. Outputs are typically stored on the high performance computers where they were generated.

Ensemble analysis is a form of meta-analysis that looks at the combined results and features of an ensemble to understand and describe the underlying domain space. By looking at groups of runs in combination, ensemble-wide patterns can be seen despite variations in the individual runs. Sandia's Slycat is a framework for ensemble analysis that integrates data management, scalable analysis, and remote user interaction through a web-based interface. Slycat supports a number of analysis types, shown in Figure 1, which are referred to as models (not to be confused with physics models, such as the simulations themselves) including: canonical correlation analysis for evaluating correlative relationships between a set of input parameters and a set of scalar outputs; time series analysis providing agglomerative clustering of temporally-changing variables; and the parameter space model, which does not actually perform any analysis, but instead provides an exploratory tool centered on a scatterplot representation with interactively selected variables for the x and y axes, point color-coding, media variables (e.g., images and videos) for retrieval and viewing during hover over points, and per variable filtering to reduce the number of visible points in the scatterplot.

FIGURE 1

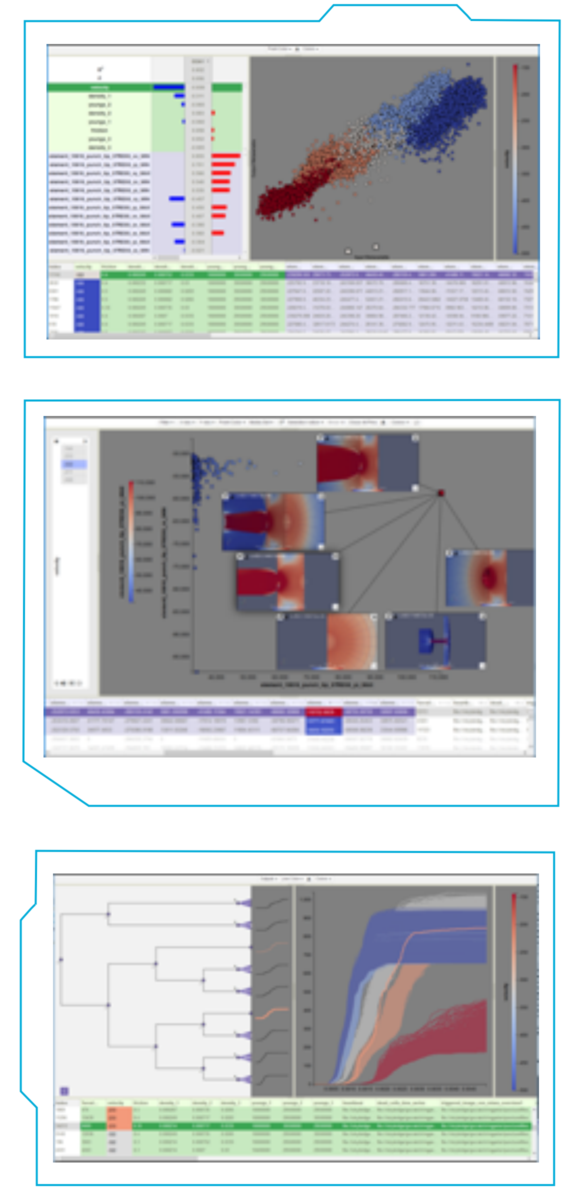

Earlier Slycat™ analysis types (starting in upper left): canonical correlation analysis; time series, and parameter space models.

When ensembles include video outputs, analysts are challenged to understand and evaluate the results. If the ensemble size grows into hundreds or thousands of runs, viewing and comparing the set of video outputs directly will not scale. Instead, our goal is to provide an approach that will be akin to how text analysis operates on documents, but for videos. Just as text analysis summarizes content without requiring one to read all the documents, video ensemble analysis should reveal video similarities and anomalies without users having to view them all.

The Slycat team has recently released VideoSwarm, a new video analysis capability that provides various data abstractions through each of four views, as shown in Figure 2: (B) a scatterplot of points (one point per run) projected using multi-dimensional scaling (MDS) to reveal frame similarity through point proximity; (D) a graph of video trajectories (one line per run) showing evolving video similarities between runs over time; (F) a quantitative metadata table of input parameters and output metrics (one row per run); and (A) a handful of selected videos retrieved from the high performance computer for synchronized comparative playback. The white line in the trajectories pane (E) acts as a slider to control time in the scatterplot and the videos. The controls (C) provide single-frame stepping, and other synchronized playback controls for the videos.

The example in Figure 2 shows an ensemble of 250 simulations. It is a material fracturing problem where a punch impacts a plate. In each run, the input values vary for parameters such as the initial punch velocity, punch and plate density, and plate thickness. Each run generates a 40-second video output, where the punch and plate are rendered using Von Mises stress values to color the cells. Although points and lines in the VideoSwarm views use the same color palette as the videos in the scatterplot and trajectory views, the color encodes the initial velocity of the punch (the input with the greatest correlation to result behaviors). Three runs (highlighted by the enlarged points in the scatterplot and saturated lines in the trajectories) have been selected for video comparison on the right. These videos are retrieved from the cluster and synchronously viewed. The scatterplot and trajectory views reveal that these runs are distinctly different from the runs encoded in red, which represent runs with the slowest initial velocity value. In the trajectory view, the red runs are fundamentally different from the rest. They are grouped at the top of the graph, with no overlaps of the other initial velocities. Loading a sample from the red simulations and comparing with the others reveals the reason. In the slowest velocity case, the punch does not fully penetrate the plate.

FIGURE 2

Scan using SNLSimMagic App

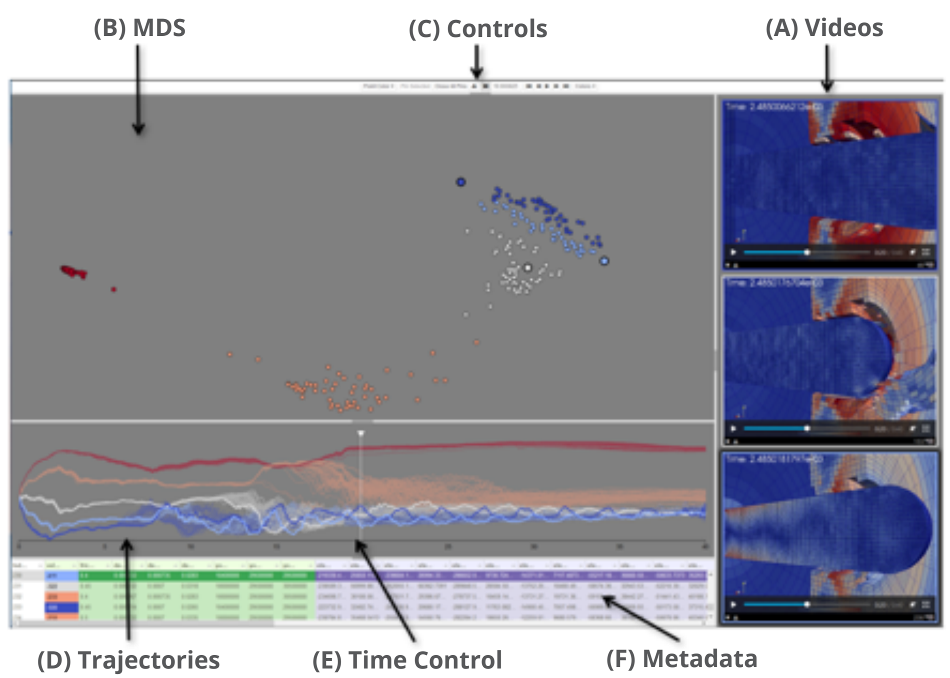

VideoSwarm running in Firefox on a Windows desktop: A) simulation video data selected by the user for playback/examination; B) MDS coordinates of the video in the ensemble, where each point corresponds to a particular video and the coordinates change when the user changes the time in the trajectories pane; C) general controls allowing selection of metadata, colormaps, and video playback options; D) video trajectories showing the progress of each video over its lifetime; E) time selection controls allow the user to navigate the ensemble in time, with corresponding updates to the MDS coordinates and video displays; F) metadata table showing information available for each video in the ensemble, including physical parameters relevant to the analysis.

REMEMBERING STEVE ATTAWAY

OUR FRIEND AND COLLEAGUE

On February 28, 2019 Sandians lost a long-time colleague and friend, Steve Attaway. Steve spent over thirty years at Sandia, and during that time he helped bring big, seemingly impossible ideas into realization: from parallelizing the PRONTO code to make it possible for thousands of cores to run it at a single time to pioneering advances in modeling transient dynamics phenomena. Throughout his career he was recognized with awards and accolades from the Department of Energy, the NNSA, and from the Deputy Secretary for Counterterrorism and Counterproliferation. And yet, Steve never focused on these, significant though they were. He was instead mindful of the fact that his tenure at Sandia would eventually end, and others would need to continue the important work being done here. Thus, he invested his time in training the next generation of big thinkers and problem solvers. He spent as much time mentoring those around him as he did relishing opportunities to tackle the "big problems" brought to Sandia for solving. The impact and legacy Steve left behind at Sandia wasn't just about his technical achievements, but how he was always willing to help, instruct, guide, direct, lead and serve others.

It's never easy to summarize someone's career and influence in the space of a few words. Nonetheless, Steve's vital role in the Sandia HPC community—not just in innovations, but also how he impacted people—is an important legacy he left behind. We share here a selection of tributes from two of Steve's mentees, Jason Wilke and John Korbin, which reflect the thoughts of the many who knew Steve.

STEVE AS A MENTOR

JASON–Steve was a guy who was willing to go out of his way to help mentor people. He was always willing to give career advice, and if you ever had a technical problem or were just stuck, he made time to come and help. He was also willing to help you find the right person to talk to elsewhere in the laboratories to help you solve your problem. As someone new here, I did not know who to call at times, but he was willing to help me make those connections and help introduce me to the right person.

JOHN–Steve was an outstanding mentor. He was focused on bringing up the next generation because he was aware that he wouldn't be here forever. One of the best things he did was working lunches. This might sound like a burden few would want, but it was truly a privilege. Around the time people start thinking about lunch, Steve would gather a group of students and peers to walk with him across the base to the BX to get a sandwich. As we walked, Steve would guide conversations about work, life, and the world in general. This was an opportunity for students, some of them just days into their careers at Sandia, to spend quality time with a Senior Scientist who was willing to give them career advice, mentor them, or just talk with them about their hobbies. It impressed me that where some people might just warm something up for lunch, close their doors, and separate themselves from the world, Steve would do the opposite. He purposefully set out to gather newer members of staff to get to know them and to mentor them, even when he was "off the clock."

STEVE AS A PERSON

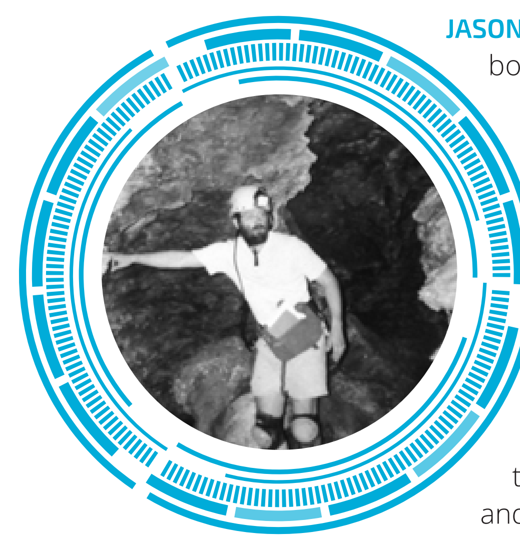

JASON–Steve was enthusiastic and always willing to push the boundaries on what could be done technologically as far as what we could get the codes to do or how big of a problem we could run on the machines. He really enjoyed tackling those immediate problems that came through the door and had to be answered quickly.

JOHN–Intellectually curious is the best way to describe him. I remember working with Steve on a project dealing with additive manufacturing with laser cured resins. We wanted to show that a process developed for jewelry manufacturing could be used to fabricate a special part we needed. Steve had the required tooling and machinery at his workshop at home, and over the weekend he designed a 3D printable part that tested the precision and accuracy of the machine, while experimenting with commercially available resins, and exploring the capabilities of the 3D printer. Over a two-day period, Steve produced a concise 10 page memo outling the problems of the machine and detailed notes on how to solve each problem, including formulation of the resin, adjustments in the printer software, modifications to the hardware. Steve sent his notes to the manufacturer of the 3D printer for free, telling them he hoped it would help them improve their product.

STEVE'S LASTING LEGACY

JASON–Steve's work on parallelizing the PRONTO code, and running it on thousands of cores at the same time, as opposed to a single core, vastly increased computing speeds. What impressed me most about Steve was he would go out and do things that people said couldn't be done. He would just think through how to do it, envision a way to get it done, and then he would work step-by-step to get there. He would encounter lots of dead-ends, would have to try various methods to get things to work, but with persistence and a little bit of luck he would routinely do what others said couldn't be done. He could be down in the weeds technically working on a problem, but he would never lose that thirty-thousand-foot view of the problem he was working on. The other part of his lasting legacy is the people he mentored and worked with that will still be here in the future. He didn't just provide people with technical skills, but he also helped them see the value of mentoring other people to pass on their knowledge.

JOHN–Steve had a box full of awards and trophies that he kept hidden in the corner of his office under his desk, but he only had two hanging on the walls of his office. The first was a plaque recognizing his contributions to high-performance computing. He did a lot of the foundational work in unlocking the potential of parallel computing by proving that algorithms could parallelized. The second was a plaque given to him by his mentees. It says a lot about someone when they don't feel the need to brag about their accomplishments. He never wanted the limelight and was happy to serve the nation when it needed him, and then return to life as usual. One of the best and most lasting lessons Steve taught me was acceptance when others would claim my ideas as theirs, since according to Steve "you are truly successful when someone steals your ideas." Steve felt that in those instances you need to realize that you have made your mark on the community and now other people will carry the torch. I think he liked that part the best, because it meant he was free to let that idea go and move on to create something new.

The Attaway

In 2008, the President of the United States asked engineers to assess the feasibility of shooting down an orbiting US satellite from a moving Navy vessel in the open ocean. It was called Project "Burnt Frost," a six-week priority mission to address a potential National Security threat. Among the engineers assigned to this project was Steve Attaway; a Sandia National Laboratories high-performance computing (HPC) expert. Without his expertise and the use of HPCs, the mission wouldn't have been successful.

To honor the late Steve Attaway and all his contributions to the HPC community, a new machine is being developed which will be named after him, and just like the man did in the past, the machine will tackle the most difficult challenges of the future.

The Attaway will be a bleeding-edge machine that will add to Sandia's HPC capacity system. Sandians who require HPC resources will be able to use The Attaway for solving big problems like Project "Burnt Frost."

HARDWARE

- 1,488 compute nodes (8 scalable units)
  - 1.93 PFLOPs peak
  - Designed to be the fastest production cluster at Sandia
- 2.3 GHz processors
  - Nodes have dual sockets with 18 cores each
  - Intel Xeon Gold 6140 "Skylake"
  - ~1.3 TFLOPs per node!
- 192 GB RAM per node (5.3 GB per core)
- Intel Omni-Path high speed interconnect
- Penguin Tundra™ hardware platform
- 480 VAC direct to Rack power shelf
- Next generation water cooling:
- Delivery in 2019

SOFTWARE

- TOSS 3.x.x (Red Hat EL 7.x)
- Lustre 2.x client
- Intel Fabric Manager
- Diskless boot over gigE private LAN
- Slurm for user job scheduling

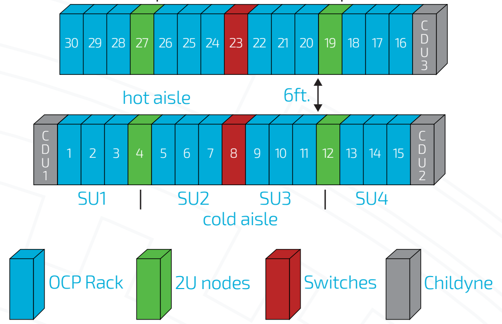

Acknowledgements

Editor
Yasmin Dennig

Writing & Editing
Johann Snyder
Sarah Johnson
Megan Davidson
Laura Sowko
Whitney Lacy
Anna Cairo
Janet Chapman
Michael Church
Laura Perlichek
Zane Brown

Design
Stacey Long Reynolds

Video
Regina Valenzuela

SNLSimMagic
Daniel Gomez

Production
Randy Bill

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.
SAND2019-11345 O