



BERKELEY LAB

Bringing Science Solutions to the World

2019

Computing Sciences Strategic Plan



2019

Computing Sciences Strategic Plan





Kathy Yelick, Associate
Laboratory Director for
Computing Sciences

Computing has transformed nearly every aspect of scientific inquiry — across disciplines and across scales — from the behavior of subatomic particles to the formation of structures in the early universe, from the assembly of the human genome to the evolution of earth systems. Over the past two decades, computing has become an integral part of how Berkeley Lab is “Bringing Science Solutions to the World.” Advances in computing and mathematics have been key, with new mathematical models of complex physical phenomena, new methods for analyzing complex data, new algorithms for accuracy and scaling and sophisticated software

systems that encapsulate these techniques into open, reusable tools. The performance of NERSC computers and the ESnet network have grown by several orders of magnitude, along with our understanding of how to map scientific computations and workflows onto these systems. From research to facility operations, the passion, talent and dedication of the Computing Sciences Area staff has been the cornerstone of our success. The plan outlined in this document describes the next step in a journey to expand the influence and impact of our efforts, building an increasingly connected global enterprise for science that places more powerful instruments in the hands of scientists, along with more powerful methods and tools for modeling, analysis and prediction.

Katherine Yelick



Jonathan Carter, CSA
Deputy for Science



Helen Cademartori,
CSA/CRD Deputy for
Operations



Sudip Dosanjh,
NERSC Division
Director



Inder Monga, Division
Director for Scientific
Networking and ESnet
Executive Director



David Brown, CRD
Division Director




Table of Contents

A Vision for Scientific Discovery	4
Progress on the Computing Sciences 2015 Vision	6
2019: Moving Forward	9
Overview of Berkeley Lab and Computing Sciences	10
Overview of Berkeley Lab	12
Overview of Computing Sciences at Berkeley Lab	12
Computing Sciences Strategic Initiatives	18
• Learning	18
• Beyond Moore	18
• Superfacility	18
Learning Initiative	21
1. Applications	23
2. Methods	25
3. Deployment	27
4. Empowerment	29
Beyond Moore Initiative	31
1. Heterogeneous Accelerators	33
2. Algorithms and Software for Accelerators	35
3. Multiscale Modeling of Post-CMOS Devices and Systems	38
4. Quantum Computing Algorithms and Software Tools	40
5. Quantum Computing Testbeds	42
Superfacility Initiative	45
1. User Engagement	49
2. Data Lifecycle	51
3. Automated Resource Allocation	53
4. Computing at the Edge	55
Summary	57

A Vision for Scientific Discovery

In today's research ecosystem, computation and mathematics are essential to scientific discovery. Over the next five years, their roles will continue to grow as the scientific process becomes increasingly automated, distributed and reproducible.



Scientists in the future will increasingly pose high-level questions using domain-relevant terminology, and automated systems will respond by running a set of experiments or simulations. Computational methods will sift through the resulting data — large, noisy and complex — to extract features and construct models to interpret the data. The boundary between theoretical and experimental science will blur as scientists combine observational and simulation data and use a variety of first-principles and learned models to make predictions. The most powerful instruments will be accessed remotely, and online collaboration tools will support dynamically configured, distributed science teams. Finally, science will be more reproducible. Scientific data will be open, available and searchable, and complex workflows captured in electronic notebooks will adapt in robust ways to new inflows of data and new software and hardware.

This vision requires advances in mathematical and statistical methods, as well as in programming systems, libraries and tools for analytics, simulation and learning (**Figure 1**). It will require more powerful networking and computational platforms, along with novel approaches to address the challenges in computing performance as device technologies reach the end of traditional scaling. Machine learning methods tailored to scientific data will provide powerful tools for analysis of complex data sets and the control of experiments, infrastructure and environmental processes. The use of learning for scientific inquiry will require improved methods that are consistent with physical laws, robust to gaps or anomalies in data and interpretable in ways that are meaningful and defensible to the scientific community.

A promising lever to meet future performance needs is processor specialization, which will usher in a new era of co-design where hardware is tailored to a given algorithm or application. In order to satisfy broad science needs, mixtures of specialized processors will be deployed in heterogeneous high performance computing (HPC) systems, and computation will be embedded near experiments and throughout the network. Longer term, quantum technology will offer new capabilities for computation, communication and sensing, complementing traditional digital devices and offering unique capabilities for certain science questions. Experimental facilities and computing facilities will need to be tightly integrated with high-speed networks, supporting real-time analytics from experiments and with large community data sets co-located with computing and data services, turning the network of facilities into a kind of *superfacility*.




Figure 1: In the near future, algorithms and systems will respond to science inquiries with automated experiments and analysis on seamlessly integrated facilities.

Progress on the Computing Sciences 2015 Vision


In 2015, the Computing Sciences Area drafted a strategic vision outlining three critical, overarching areas: Exascale computing to satisfy performance demands up to 2025; Mathematics to address the ever-increasing complexity of scientific questions; and Data to meet the needs of experimental and observational science. Over the last four years we have made considerable progress in all three areas.

Exascale

DOE formed the Exascale Computing Project (ECP) in 2016 and began work to deliver exascale systems by the early 2020s. With a focus on scientific impact, Berkeley Lab defined a lab-wide priority in “Breakthrough Science at the Exascale” designed to combine the best mathematical methods with innovative software technology to deliver new science capabilities for future exascale systems. Berkeley Lab was chosen to lead or co-lead 11 of the 25 applications projects under the ECP and other key projects across the Laboratory. These cover all major scientific domains at the Lab: Physical Sciences, Earth and Environmental Sciences, Energy Sciences, Biosciences and Energy Technologies. The application areas cover cosmology, astrophysics, accelerator design, chemistry, subsurface flows, genomics, urban systems, carbon capture and earthquake simulations.

Additionally, the Lab is playing a lead role in three ECP co-design centers, in which algorithms and applied math are developed and tailored for ECP applications and future exascale hardware. For example, the adaptive mesh refinement (AMR) methods developed by the Lab have been adapted for exascale use and, through the AMReX co-design center, have been incorporated into at least five of the ECP applications (**Figure 2**). Staff at the National Energy Research Scientific Computing Center (NERSC) launched exascale programs in simulation, data and learning applications to help its thousands of users prepare for NERSC’s series of pre-exascale systems in 2016 and 2020 that are deploying

Figure 2: This image shows a simulation of an accelerator experiment using the WarpX software developed as part of ECP. WarpX is using AMR software developed through the ECP codesign center AMReX at Berkeley Lab.



key technologies on the path to exascale. Combined with ECP software and co-design efforts from the Computational Research Division (CRD) and high-speed networking from the Energy Sciences Network (ESnet), these activities will broaden the impact of exascale computing to the science community beyond the 25 ECP applications.

Math

Our focus on math in 2015 was designed to further extend the scientific impact of the Lab's Applied Mathematics program. Berkeley Lab's program is known for making fundamental advances in the development of new models and algorithms and for providing usable math tools for scientists. In addition to its traditional application areas, the applied math expertise at the Lab now comes to bear on a new class of problems through CAMERA, the Center for Advanced Mathematics for Energy Research Applications. CAMERA brings together applied mathematicians, image and signal processing experts, data scientists and experimental scientists to attack algorithmic and data challenges at DOE Basic Energy Sciences experimental facilities (including the light sources, neutron scattering sources and nanoscience centers) and to accelerate the transfer of new mathematical ideas that can significantly improve the analysis and understanding of experimental data.





Figure 3: The CAMERA team developed the M-TIP algorithm to show reconstructions of a rice dwarf virus (a-b) and a PR772 bacteriophage (c-d) from experimental correlation data, with b and d showing asymmetries in the internal genetic material for each virus reconstruction.

As one example, CAMERA researchers contributed key algorithms that helped an international team achieve a goal first proposed more than 40 years ago: using angular correlations of X-ray snapshots from non-crystalline molecules to determine the 3D structure of important biological systems. The breakthrough resulted from a single-particle diffraction experiment conducted at the DOE's Linac Coherent Light Source (LCLS) by the Single-Particle Initiative organized by the SLAC National Accelerator Laboratory. The team used CAMERA's multi-tiered iterative phasing (M-TIP) algorithm to perform the first successful 3D virus reconstructions from experimental data (**Figure 3**).

Data

Sometimes described as a tsunami, the onslaught of experimental and observational data has disrupted long-standing ideas about the design of computing centers which for years have focused primarily on producing and storing results from simulations. The ability to collect and understand these massive data sets, often generated at remote experimental facilities, requires new approaches to advanced computing and networking. NERSC and ESnet have been used in production analysis for data from the Large Hadron Collider (LHC) and the Joint Genome Institute (JGI); in supernova detection,

Figure 4: CRD is home to the Center for Computational Cosmology, which produces scalable data science tools for cosmological analysis of phenomena such as the cosmic microwave background (pictured).



studies of the cosmic microwave background and more; and, most recently, have created real-time job scheduling for light sources and electron microscopes to support analytics during experiments. CRD has built tools to manage data pipelines, as well as analytics software for both experiments and simulations. These tools include custom software for projects in cosmology, particle physics, environmental science, materials science, imaging and biology, as well as general-purpose analysis and visualization software. Some of these tools help scientists perform live analyses on their data, enabling them to quickly determine the quality of the data and ascertain whether to repeat an experiment or move on to the next one.

Ongoing upgrades to ESnet and NERSC also support the growing vision of interconnected facilities for science. In 2017, ESnet launched ESnet6, a major project to upgrade to terascale networking and provide more programmability and automation within the network. In 2018, NERSC announced the procurement of its next-generation supercomputer, a pre-exascale machine slated to be delivered in 2020. Named “Perlmutter” in honor of Berkeley Lab’s Nobel Prize winning astrophysicist Saul Perlmutter, it is the first NERSC system specifically designed to meet the needs of data streams from experimental and observational facilities and machine learning applications, as well as large-scale simulations.

The Computational Cosmology Center in CRD has worked closely with NERSC and ESnet to develop high performance tools to stream and analyze supernova images and cosmic microwave background (CMB) data for national and international science collaborations (**Figure 4**). With a leadership role in data analysis for the next major CMB experiment (CMB-S4), the team has generalized the problem of CMB map-making to the reduction of any pointed time-domain data with TOAST (Time Ordered Astrophysics Scalable Tools), a tool that is optimized to run at scale on NERSC’s pre-exascale systems.

2019: Moving Forward

While the progress in these three areas has led to new capabilities for science, work remains to be done to meet the goals of automation and reproducibility and to address the new challenges facing computing performance. This reality led us to reassess our position and our areas of future emphasis. The 2019 Strategic Plan details our objectives and motivations for three strategic initiatives to address the greatest needs and provide the greatest impact for the DOE research community over the next five years: Learning for Science, Beyond Moore and the Superfacility model. These interconnected initiatives are supported by our recognized expertise in applied mathematics, HPC, data science and networking, as well as existing investments in the physical infrastructure for computing and networking.

Learning for Science: Machine learning has emerged as a powerful set of methods for learning predictive models, and it is rapidly moving into the analysis of scientific data and control of scientific experiments and systems. High performance computing has proven to be invaluable in scaling some of the largest learning methods, but methods need to be adapted for the complexity of scientific data, given that standards for understanding and explaining learned models is much higher in science than in commercial or social applications. We need much more robust and interpretable machine learning methods, and they need to scale with the data, models and machines being used. Most important, we can't have methods that just give us a correlation or answer without any interpretation behind it.

Beyond Moore: DOE mission requirements will continue to demand increased computing performance even while the traditional performance gains we have come to expect from lithography improvements taper off as we approach atomic scale. We cannot keep scaling microelectronics performance without co-design that spans all layers — from atomic-scale materials to large-scale complex systems — to meet the needs of emerging mission scientific applications. We need multiple approaches to future hardware, including exploration of new transistor technologies, heterogeneous accelerators and fundamentally new approaches such as quantum computing.

The Superfacility Model: Increasingly, users of DOE's national user facilities access HPC systems remotely, capturing the data and then moving it, analyzing it and repeating the whole process multiple times. But what if these steps could be woven into a seamless progression of phases? The superfacility concept is a blueprint for seamlessly integrating experimental, computational and networking resources to support reproducible science. This initiative is designed to go beyond the successful point demonstrations of the past and define an automated architectural model to support streaming data from experimental facilities, a single interface for users, improved resilience and integrated tools for sharing, searching and analyzing data for more productive, reproducible science.

These three initiatives will help Computing Sciences harness the expertise and resources of our three divisions — NERSC, ESnet and CRD — to deliver on our vision for science.

Overview of Berkeley Lab and Computing Sciences



Berkeley Lab is known for its scientific excellence, with 13 associated Nobel prizes, 15 awardees of the National Medal of Science and 91 memberships in the National Academies of Science, Engineering and Medicine. The Lab vision, “Bringing Science Solutions to the World,” goes beyond fundamental science to develop transformational impacts on society, particularly for energy and environmental challenges. Interdisciplinary, team-based science has been at the core of the Lab’s success for more than 85 years, emanating from founder Ernest Orlando Lawrence, who is often referred to as the “father of big science.”



Overview of Berkeley Lab

Berkeley Lab is a member of the DOE's national laboratory complex, and its scientific preeminence is cemented by its proximity to the University of California at Berkeley, which has the largest number of top ten graduate programs of any university in the U.S. and is known around the world for excellence. We are a laboratory committed entirely to open, unclassified research and are the most collaborative of any laboratory, with nearly 10,000 publications jointly written with universities, industry and other national laboratories. The national laboratories all excel in designing, building and operating large, one-of-a-kind instruments for the science community, but in this aspect Berkeley Lab stands out. Not only do we develop some of the technology used worldwide in detectors, accelerators and other devices, but the five national user facilities at Berkeley Lab serve nearly 10,000 users each year, more than any other national laboratory.

Maintaining a global reputation for excellence in breakthrough science requires more than excellence in research. As an institution managed by the University of California for the benefit of the nation, Berkeley Lab is committed to the highest level of scientific integrity and dedicated to ensuring that all staff demonstrate impeccable ethical conduct in all aspects of their employment.

Underpinning all of Berkeley Lab's research efforts is an ongoing commitment to the health and safety of employees through continuous training and awareness programs and regular efforts to remind staff that we are all responsible for our collective well-being. And just as in science, where differences in perspective and knowledge often meld and lead to bigger breakthroughs, Berkeley Lab is committed to building an inclusive, collaborative and open environment, honoring diversity in people and in ideas. Berkeley Lab is determined to increase the diversity of its workforce and to promote diversity in STEM for the nation's next generation of engineers and scientists. Toward this end, Berkeley Lab was the first national lab to post its diversity statistics online and led the effort for all 17 DOE national labs to post their collective diversity statistics. The Lab also continues to build and expand its local and national outreach efforts in K-12, higher education and diverse professional associations, targeting underrepresented communities in STEM.

Overview of Computing Sciences at Berkeley Lab

The Computing Sciences Area mission is to achieve transformational, breakthrough impacts in scientific domains through the discovery and use of advanced computational and mathematical methods and systems that are accessible to the broad science community.

This mission embodies our commitment to influence basic and applied science areas, using computation to solve societal and environmental problems, address national and international priorities and aid in the pursuit of fundamental scientific discovery. We are equally committed to advancing the field of computing writ large, including underlying

mathematics, statistics and algorithms as well as computational devices, programming methods and systems. In keeping with Berkeley Lab's tradition of broad impact, these capabilities are delivered to the community through world-leading facilities and widely used software that solves specific scientific challenges or is generic across a range of applications.

The Computing Sciences program at Berkeley Lab is anchored by two major user facilities — NERSC and ESnet — and CRD, the strongest computing research program across the national laboratory complex, with major activities in applied math, computer science, data science and domain-specific partnerships in computational modeling and data analysis. The facilities serve a large cross-section of the DOE science community, with NERSC supporting several thousand researchers from universities, national laboratories and industry working on DOE mission problems, and ESnet connecting the DOE community to each other and the rest of the Internet for distributed team science.

The strength of the research program and facilities rests on the ongoing excellence of the staff, who bring outstanding intellect, expertise and leadership to address a broad set of research and operational challenges. The Area has a number of activities to attract new talent, train them in the collaborative science model, instill a culture of inclusion and support their professional growth and careers. It has a rich set of weekly seminars to keep people informed of the latest work inside and outside the Area. In addition to leadership development programs run by DOE and the Lab, the Area has a biennial internal mentoring program that is open to all staff wishing to be a protege or a mentor, and this successful model has been adopted by other areas at the Lab. To expand the workforce vital to the DOE's mission areas in advanced computing, Computing Sciences partners with the Sustainable Horizons Institute to bring more than 100 faculty and students from institutions historically under represented in the research community to Berkeley Lab. These faculty and student teams collaborate with Computing Sciences researchers in paid internships each summer, often returning for two or even three years to continue their work.



Figure 5: The Berkeley Lab Computing Sciences area has over 400 staff across three divisions: NERSC, ESnet and CRD.



Figure 6: NERSC staff in front of Shyh Wang Hall at Berkeley Lab.

The National Energy Research Scientific Computing Center (NERSC) is the mission high performance computing facility for the DOE's Office of Science. NERSC supports more than 800 projects covering a diverse workload that represents every scientific discipline within the Office of Science. The scientific record of NERSC is exceptional, with users each year reporting more than 2,000 publications and about 20 cover stories stemming from their work at NERSC. In the past 15 years, six Nobel prizes have been awarded to scientists or projects affiliated with NERSC.

In addition to delivering outstanding computational and data services, NERSC drives the industry toward innovation of high performance computing solutions, with first-of-a-kind systems designed to run a rich workload of applications, including those that require the full system capability. The NERSC Exascale Science Applications Program (NESAP) devotes personnel, computing resources and training programs to work closely with the user community to move their applications toward exascale architectures.

NERSC procures supercomputing systems every few years and often has two in production to ensure smooth transitions for users. The NERSC-8 system, named after Nobel Laureate Gertrude Cori, arrived in 2016 using manycore compute nodes and a flash-based burst buffer to accelerate data-intensive applications. NERSC recently announced NERSC-9, a system named Perlmutter to be delivered in 2020 that will integrate both CPU-only and GPU-accelerated nodes with a Cray network designed to support the rapid ingest of data from external facilities and a system architecture that can be extended to include other types of accelerators. Perlmutter will have an all-flash file system to accelerate heavy I/O applications and, like NERSC-8, will include a robust application readiness program.




The Energy Sciences Network (ESnet) fuels collaborations around the globe by connecting the DOE national laboratories to each other and to the rest of the national and international scientific world. ESnet exists to provide the specialized networking infrastructure and services required by the national laboratories, large science collaborations and the DOE research community. ESnet is recognized nationally and internationally as one of the premier science networks. It has a long track record of innovation in network design, performance and service delivery, highlighted by major contributions that are utilized by other research networks around the world. It offers novel services, such as bandwidth reservations to accelerate the transfer of very large data sets, and its widely adopted “Science DMZ” model ensures that benefits of high-speed networking and big data movement extend into the private networks of universities and laboratories worldwide.

ESnet transferred close to an exabyte of data in 2018, with over two-thirds of the traffic composed of large data science data flows. ESnet is a national and international resource for collaborative science, with 80 percent of the data moving outside the DOE laboratories. Initiated in December 2014, and upgraded in 2017, ESnet put four trans-Atlantic links into production, giving U.S. and European researchers 400Gbps of dedicated bandwidth, ensuring the fastest connections between scientists in the U.S. and collaborators and facilities in Europe, such as the Large Hadron Collider. Late in 2016, DOE approved the start of the ESnet6 project, which is the next major upgrade to the network. The ESnet6 project will achieve a transformational change in network capacity, resiliency and flexibility that will pay tangible benefits to the DOE mission and U.S. competitiveness.

ESnet serves as a vital “circulatory system” for all DOE research facilities and major projects and for every mission space within the Office of Science. ESnet’s strong history of innovation in network architecture, software and science-driven services support breakthrough science with large, remote and distributed data.



Figure 7: ESnet staff at a group team-building meeting in January 2019.



The Computational Research Division (CRD) is the Computing Science Area's research arm. CRD's research program is internationally recognized for excellence in applied mathematics, computer science, computational and data science.

CRD's Applied Mathematics program rivals that of any university in scientific excellence, with six National Academies memberships, one National Medal of Science awardee, and seven SIAM Fellows. Leveraging the UC Berkeley connection, the program has consistently achieved fundamental breakthroughs in science-driven modeling, adaptive mesh refinement, interface methods and scalable algorithms. Just as mathematics is the universal language for science, the math program is the linchpin for collaborative research, with the latest methods used for everything from simulating the Lyman-alpha forest in cosmology to the formation of bubbles used in industrial foams. Berkeley Lab has shown strong leadership in modeling and simulation for DOE's applied mathematics

program and currently leads the DOE SciDAC (Scientific Discovery through Advanced Computing) FASTMath Institute. CRD has also led the community in establishing the mathematical foundations of scientific data analysis for the Office of Science. Building on Lab funding, DOE established CAMERA, the Center for Applied Mathematics for Energy Research Applications at Berkeley Lab, which includes work with the Advanced Light Source and other DOE Basic Energy Sciences scientific user facilities to enhance the analysis of data from beamline experiments using a rich set of advanced mathematical methods.




Figure 8a: Center for Applied Mathematics for Energy Research Applications

Computational and data science research within CRD has strong connections with other Berkeley Lab divisions, leading to computational and data activities in computational cosmology, biology, chemistry, materials, high energy physics and climate modeling. The boundaries between math, computer science, data science and domain science are fluid, with project teams drawing from expertise across the Lab. The research program and NERSC have a long history of collaboratively developing tools for major experiments as well as modeling and simulation software. The science engagement activities in both divisions draw upon user expertise to define their future directions and collaborate in deploying state-of-the-art computational systems and networking services. Several of the cross-disciplinary activities address the end-to-end science problem, delivering complete solutions to science communities.

CRD's computer science research activities are well-recognized across the Lab complex, with four ACM Distinguished Scientists or Fellows and best paper awards at all of the major conferences in high performance computing. Berkeley Lab computer scientists led two of the four DOE SciDAC Institutes in 2011-2016. The Lab conducts research across the core areas of computer systems research, including computer architectures, operating systems, programming models, parallel algorithms, performance modeling and optimization, as well as data science-related areas of visualization, graph analytics, machine learning, data management and collaboration tools.

Figure 8b: Computing Sciences Area and Computational Research Division Management



Helen Cademartori, CSA/CRD Deputy for Operations
David Brown, CRD Division Director
Peter Nugent, CRD Computational Science Department Head
Kathy Yelick, CSA Associate Laboratory Director
Jonathan Carter, CSA Deputy for Science
John Shalf, CRD Computer Science Department Head
John Bell, CRD Chief Scientist
Deb Agarwal, CRD Data Science and Technology Department Head
Edmond Ng, Division Deputy (not shown)
James Sethian, Head of CAMERA (not shown)

Computational Science

Data Science & Technology

Computer Science


Applied Mathematics



Computational Cosmology



Data Analytics & Visualization




Performance & Algorithms




Applied Numerical Algorithms



Computational Biosciences



Scientific Data Management




Computer Architecture




Center for Computational Sciences & Engineering



Physics & X-ray Science Computing



Integrated Data Frameworks



Computer Languages & Systems Software




Mathematics



Computational Chemistry, Materials & Climate




Usable Software Systems



Scalable Solvers

Computing Sciences Strategic Initiatives


- Learning
- Beyond Moore
- Superfacility





The 2015 Computing Sciences initiatives are now ongoing activities in support of our vision for scientific discovery. In particular, the Exascale Computing Project and CAMERA are well-established DOE-funded activities, and there are multiple data-intensive science projects using ESnet, NERSC and CRD for experimental and observational science.

For 2019, we have identified three new initiatives that deserve an area-wide focus, with the expectation that they will grow over the next five years into self-sustaining projects that will deliver scientific impacts for many years to follow: Learning, Beyond Moore and Superfacility.



Learning Initiative

In scientific discovery, learning has become essential to identifying features of massive collections of images, genome sequences, text and simulation output, and it will increasingly be used to control complex energy and environmental systems. Predictions based on modeling and simulation of physical processes using first-principles applied mathematics (e.g., partial differential equations) will continue to play a critical role in science, and they will be augmented with models inferred from data, especially in scientific areas where first-principles are unknown or inaccessible. Applications of advanced computing will go beyond traditional foci on the physical sciences to problems that involve infrastructure management, human behavior, ecosystems and biological processes.

In the CS Learning Initiative — which incorporates deep learning, traditional machine learning and statistical methods — our vision for the future includes scientific instruments controlled by intelligent, high-precision robotics; data analytics methods that learn models from noisy, complex, multi-modal and multi-scale data sets; high-throughput simulation campaigns managed by algorithms; and computational and networking facilities that automatically adapt resources to user demand and respond to anomalous behavior.

Learning methods will empower scientists by allowing computers to identify emergent or unexpected patterns in heterogeneous, complex datasets and find hidden signals recalcitrant to classical interrogation. Our approach is driven by the needs of science, so rather than limiting attention to one class of methods, our Learning Initiative very broadly includes deep learning, traditional machine learning, statistical methods and scalable analytics. We will use interdisciplinary application teams to solve open scientific questions using real data sets, addressing data paucity along with abundance, systematic errors along with noise and constraints due to laws of physics along with other known structures in the data.

But answering fundamental questions on the foundations of machine learning is essential to their use and acceptance in scientific discovery. Deep learning methods result in highly accurate predictions, outperforming more traditional techniques, but they remain almost entirely “black boxes” — we can see that they work, but they





Figure 9: This image shows a developing jet diffusion flame simulation run at NERSC using Berkeley Lab’s PeleM code developed as part of the ECP’s combustion application project and with adaptive mesh refinement methods and software from the ECP AMReX Co-Design Center. In fluid dynamics problems more broadly, applied math researchers have used optimization, statistical methods and deep learning to aid in parameter estimation and fast surrogates and to derive highly accurate models from data.

provide little or no insight to the systems they model. Further, they require substantial trial and error to use, the theoretical foundations are poorly understood and the resulting models are complex and lack natural interpretations. This may be acceptable for placing advertisements or finding cat videos, but it is a serious concern in scientific discovery. Our research program will therefore address these and other foundational questions, including possibilities for bias, lack of stability and open questions in explaining how, when and why particular methods work — all with the goal of enabling scientific discovery in regimes and at scales that have not previously been accessible to inference.

For some of the largest and most challenging problems, the power of HPC systems will be needed, providing a unique capability in DOE to accelerate model training and inference. Berkeley Lab is already a leader in the use of HPC for machine learning in science and has a broad set of activities adapting and applying machine learning methods for scientific discovery. Our work on scalable linear algebra libraries, parallel algorithms and HPC systems software will enable new and existing learning methods to take advantage of the massive memory and compute power of systems at NERSC and other facilities.

Berkeley Lab and the surrounding community are a powerhouse of machine learning talent and expertise, including UC Berkeley's academic programs, local startups and large corporations, non-profits and other universities and laboratories. To catalyze a learning effort for scientific discovery, we will offer a set of seminars, training material, public data sets and benchmarks to empower scientists to use the best learning methods, software and systems.

Figure 10: Researchers in Berkeley Lab's CAMERA center have developed a mixed-scale dense convolution neural network (MS-D) that requires far fewer parameters than traditional machine-learning methods, converges quickly and has the ability to learn from small training sets. This image shows a slice of mouse lymphoblastoid cells, with raw data (a), corresponding manual segmentation (b) and output from an MS-D network with 100 layers.



Our Learning Initiative forms the computational core of a broader Lab-wide initiative called ML4Sci that involves Laboratory Directed Research and Development (LDRD) funded application projects in learning from across the Lab. Playing to the strengths and ongoing activities in Computing Sciences, the four high-level goal categories for our initiative are:

1. Adapt and apply learning techniques to address the complexity of fundamental and applied science applications through cross-disciplinary collaborations.
2. Develop robust and interpretable learning methods.
3. Develop and deploy premier hardware, software and services for learning at NERSC.
4. Empower the science community with learning collaboration activities, training and seminars.

1. Applications

Adapt and apply learning techniques to address the complexity of fundamental and applied science applications through cross-disciplinary collaborations.

In nearly every domain of science, researchers are considering machine learning algorithms for analyzing data, solving inverse problems, approximating time-consuming simulations or experiments, controlling complex systems and for design and engineering. Indeed, multiple DOE workshops and reports highlight the urgent need to apply and adapt learning techniques to solve our most pressing challenges. In recent years, significant progress has been made across CRD, NERSC and ESnet for some application areas. Examples include enabling interpretation of multispectral images from light sources; predicting real-time network behavior and avoiding data traffic congestion; and feature detection, classification, segmentation and generation for neutrinos, extreme climate events, supernovae, microbial genes and more, generally for data from experiments, observations, and simulations. Ongoing Lab efforts reach into new areas, including the use of machine learning to automatically label science data with metadata information for improved search, the development of surrogate models for managing groundwater and methods to analyze multi-modal sensor data from highly instrumented farmland for sustainable agriculture.

Expanding this effort, we will embark on in-depth engagements with key scientific partners to build learning algorithms directly into modeling and simulation applications, experimental and observation data (EOD) science pipelines and the operation of facilities and infrastructure. We will capture some of the best practices and common pitfalls in a lessons-learned report on learning in science. This report will include areas such as generative networks, weakly supervised deep learning, graph neural networks and large-scale learning, and will compare traditional analytics, deep learning and other forms of machine learning to provide a perspective from science. The report will go beyond summarizing existing efforts to identify infrastructure requirements, limitations faced with current methods, common computational and data access patterns and how the data characteristics affect the choice of techniques. This applications effort will take advantage of the lab-wide LDRD initiative in ML4Sci, the Exascale Computing Project's (ECP's) ExaLearn project and NERSC's learning applications in the NESAP partnership program.


2-Year Milestones

- Produce a website describing Berkeley Lab learning experience in science applications.
 - Demonstrate new learning applications across multiple EOD projects and simulation applications.
-

5-Year Milestones

- Develop an on-demand hub of pre-trained models and application-level services.
- Demonstrate new learning applications for engineering design and control of complex systems.

Figure 11: Machine learning is being used to reconstruct the tracks of particles in high energy physics experiments at facilities such as the Large Hadron Collider at CERN. CRD and NERSC staff collaborated on the TrackML Challenge, a competition designed to inspire new algorithmic developments that can quickly reconstruct particle tracks from three-dimensional coordinates left in silicon particle detectors following millions of particle collisions.



For EOD pipelines, learning will be incorporated into all stages, including filtering, compression and reduction *in situ* at instrument; smart scheduling of facilities; advanced analytics; and statistics/machine learning on HPC machines. For modeling and simulation applications, we will build on our existing strength and ownership of these applications to use learning where data-driven models are needed, to analyze outputs or to manage ensemble calculations. Longer term, we will develop methods for facilities and infrastructure management, optimizing resource allocation and detecting problems such as component failures and cyber incidents. We will demonstrate these within our own computing and networking facilities operations and in infrastructure, such as the control of transportation or energy systems, and through the use of reduced order models for real-time prediction of water systems management.

Guided by domain experts and this experience, we will implement common methods and processes in new libraries on top of popular frameworks and develop model hosting platforms (“hubs”) for publication of models, weights and data that can be reused for transfer learning across related problems with restricted or limited data. Our implementations will be optimized for heterogeneous hardware architectures, will use a variety of parallelism techniques for HPC scalability (e.g., model, data and domain parallelism) and will be available to the community through well-engineered and well-supported software suites.

Ultimately, our goal is to ensure that all science projects from Lab researchers, NERSC users and the wider DOE community have access to state-of-the art learning approaches and necessary expertise that can be easily applied to their application at-scale throughout their pipelines.

2. Methods

Develop robust and interpretable learning methods for science.

Mathematically, all learning methods can be characterized as identifying and representing structure in the data. Geometrically we might be looking for low-dimensional spaces or sparse representations in some basis to characterize the data. Optimization is used to take advantage of the special structure in data, both for efficiency and quality of results.

Deep learning and most learning techniques available today are able to identify complex and subtle patterns in data. However, those patterns remain locked away in the derived model — we are able only to glimpse aspects of the fitted functions and cannot yet extract internal data representations. Our goal is to improve foundational understanding of machine learning methods and to develop methods with improved interpretability; such methods will capture realistic mechanisms and not simply correlations, are consistent with known constraints such as physical laws and reveal quantifiable uncertainties and identifiable biases (**Figure 12**).




Figure 12: Berkeley Lab researchers are using interpretable machine learning methods to analyze and understand the interrelationships between soil chemistry, microbes, plant types, water and other factors in precision agriculture.

We aim to go beyond correlation, or statistical association, to causation and mechanism. This research will focus on two paradigms: building methods that facilitate interpretation up front (e.g., produce models that have few parameters and satisfy known properties of the domain); and designing interpretation techniques for the interrogation of

2-Year Milestones


- Develop new strategies for building interpretability and explainability into machine learning methods (including deep learning modalities) with minimal impacts on performance.
- Develop new techniques for extracting the internal data representations learned by “black box” algorithms.
- Develop frameworks for the quantification of uncertainty in data representations extracted from learning machines, including deep neural networks.

5-Year Milestones

- Develop the mathematics to extract human-interpretable surrogate models that are robust to outliers and with quantifiable uncertainty.
- Apply new surrogate modeling techniques to enable adaptive optimization and control of complex systems, such as energy systems and watersheds.

learned models, regardless of how they are constructed. An example of the first would be incorporating physical or mathematical knowledge to the learned models to capture “cause and effect” by design. On the other hand, black-box architectures, such as convolutional neural networks (CNN), will require extraction techniques where the underlying (hidden) model is exposed in a human-navigable surrogate amenable to visualization and inferential procedures. **Figure 13** illustrates this type of interrogation scenario.

Figure 13: Our vision for machine learning methods addresses data from a variety of instruments, simulations and feedback to control experiments. Methods need to produce models with natural interpretations or techniques for interrogating those models.



Interpretation is especially challenging for high-dimensional data where the underlying structure is not easily derived. Techniques such as topological analysis can help discover this structure and determine its inherent complexity. One might use such a structure as a guide to find better embeddings of the data (ones that amplify the low-dimensional structure) or to translate input data into a different representation on which learning would be effective. To advance our mathematical understanding and capabilities in learning, we will look for ways to leverage internal expertise and develop additional expertise in fields of importance to learning, such as measure theory and analysis, black-box and non-black-box optimization, differential geometry, algebraic topology, optimal transport (for comparing probability distributions), Hamiltonian Monte Carlo and linear algebra.

Developing stronger foundational insights into why and how predictive architectures function will enable us to enhance the efficiency of machine learning design and training and encapsulate user-defined hyperparameters in fitting routines. Future generations of learners will need to “know when they don’t know” — to adapt to, and correctly flag, unexpected or adversarial inputs. These capabilities, including the capacity to register “surprise,” are essential for the process of scientific discovery. We aim to improve uncertainty quantification, enable sensitivity analysis and, ultimately, statistical inference.

3. Deployment

Develop and deploy premier hardware, software and services for scalable learning.

Powerful platforms designed for data management and analytics enable the DOE open science community to address problems at scale. This requires efficient use of existing and soon-to-be deployed supercomputing hardware. In collaboration with industry, and in coordination with the science user community, we will develop scalable learning libraries on the Cori and Perlmutter platforms. This will build on our work in scaling and optimization of dense and sparse linear algebra, graph algorithms, randomized algorithms, communication-avoiding techniques and data analytics frameworks. It also overlaps with our Beyond Moore initiative in the use of algorithms and software for next-generation hardware. Usability is a key consideration, and we will develop productive interfaces (e.g., notebooks) for running jobs at scale for distributed training, inference and hyper-parameter optimization.




Figure 14: Trained deep learning neural networks can identify weather fronts, tropical cyclones and long narrow air flows that transport water vapor from the tropics, called “atmospheric rivers.”

To aid in vendor partnerships and inform future procurements, we will develop a suite of learning benchmarks to evaluate emerging computing hardware and serve as use cases. New and existing vendor relationships will ensure that software works well on HPC systems and for science use cases, leveraging collaborations and lessons learned from application areas. The NERSC exascale partnership program

2-Year Milestones

- Deploy a rich learning stack at NERSC in partnership with software and hardware vendors.
- Develop a comprehensive learning benchmark suite for future procurements to shape evaluation of emerging hardware and NERSC-10 system design.

5-Year Milestones

- Design, build and adapt learning algorithms and software for scalability on next-generation HPC platforms.
- Deliver a NERSC-10 system with appropriate, highly optimized, energy-efficient technology for learning applications.
- Deploy an expanded learning software stack on NERSC-10 that includes scalable interpretable methods.

(NESAP) will include learning applications, which will further shape application benchmarks for configuring the NERSC-9 system and designing NERSC-10. Documentation and training events (linking with empowerment and application areas) will help ensure optimum use of each system. Finally, we will ensure that infrastructure hooks exist for enabling smarter facilities and computational steering.

The five-year time frame will see evaluation and incorporation of machine/deep learning-focused hardware into systems, evaluated using the benchmarks and applications developed in the previous milestone. These may be deployed within NERSC-9 (Perlmutter) or as standalone testbeds and built into the procurement plans for NERSC-10. In addition, we will see full productization of the learning software stack, libraries and hubs developed in the application area for use by a large fraction of the NERSC user base.

Systems beyond NERSC-10 will be compelling platforms for learning, built through the enhanced vendor engagements and benchmarking efforts in the next few years. EOD and simulation pipelines will by then be able to seamlessly plug in to NERSC for their mission computing, gaining automated access to appropriate state-of-the-art models served on-demand. Creative applications of new methods will continue to be a focus, with close engagement with strategic science partners, but these will be made immediately available for use by the entire DOE community via productive interactive interfaces.

Figure 15: Horst Simon, Berkeley Lab Deputy Laboratory Director, gives opening remarks at the first ML4Science workshop in September 2018. This popular event brought together luminaries in machine learning with scientists across multiple domains.



4. Empowerment

Empower the science community and DOE Office of Science to use learning methods by establishing collaboration activities, training and seminars.

To realize the potential of our work for scientific discovery, we will enable and drive adoption in the scientific community. We will create collaboration activities, training materials and sessions, seminars and symposia to make our learning algorithms easy to deploy and scale across HPC lab facilities. We will develop training materials and activities to help new users gain access to our interpretable methods, application hubs and high-performance software stack. We envision building a powerful community of learning experts and applications across the Lab that also leverages activities on the Berkeley campus, as well as the industry and research ecosystem physically and virtually surrounding Berkeley Lab.

Bringing together the DOE community across the Laboratory complex, and leveraging the expertise in local academic and industry leaders in machine learning, we will establish an annual symposium to highlight the latest results and trends in machine learning for science. This will include highlights of recent hardware and software platforms, new foundational work and methods relevant to science and scientific impact across a broad range of areas. It will also underscore limitations of current approaches and highlight remaining open questions, as well as lessons learned across facilities.


We will focus on providing workshop and tutorial content appropriate to the conferences attended; examples of conferences we will target include Supercomputing (SC), International Conference on Machine Learning (ICML), Knowledge Discovery and Data Mining (KDD), ISC High Performance and Cray Users Group (CUG). We will also expand content to develop institutional forums for capabilities within the annual Machine Learning for Science workshop at the Lab. This work will include a number of outreach efforts, including a weekly seminar series to coordinate machine learning efforts with ECP projects (e.g., ExaLearn), other research projects, the machine-learning-focused summer program and NERSC tools (e.g., partnering with software carpentry or the LabIT team) to establish learning training within and outside DOE science domains.

2-Year Milestone

- Establish a Berkeley Lab-led annual symposium and workshop in learning for science.

5-Year Milestone

- Develop training materials that become cornerstone resources for the integration of learning algorithms and theory through the DOE science space.



Beyond Moore Initiative

Moore's Law is a techno-economic model that has enabled the information technology industry to nearly double the performance and functionality of digital electronics roughly every two years with fixed cost, power and area. As photolithography approaches atomic scale and fabrication costs continue to rise, the classical technological driver that has underpinned Moore's Law for the past 50 years is already failing and is anticipated to flatten by 2025 (**Figure 16**).




Figure 16: The history of Moore's Law, from the 2017 Turing Award lectures of David Patterson and John Hennessy.


But DOE mission requirements will continue to demand increases in computing performance in the absence of traditional technology scaling. Thus, the overarching goal of our Beyond Moore Initiative is to create new technologies and computing systems that transcend the performance limitations caused by the tapering of lithographic scaling.

The Beyond Moore Initiative is aligned with two broader Lab-wide initiatives: one in digital technologies, which includes advances in materials, devices, lithography and manufacturing; and a second in quantum information science (QIS), which is developing and using quantum sensors, computing and communication devices for science problems in materials, chemistry, high-energy physics and nuclear physics. Computing sciences is a key player in both of these Lab-wide initiatives, providing modeling capabilities for new materials, hardware design and evaluation, as well as algorithms and software developments to enable improvements in performance, energy efficiency, integration and scaling. Algorithms and software work will serve to enable emerging science problems, inform hardware designers of computational requirements and options and adapt to future hardware features. The introduction of quantum devices will require even more aggressive algorithms and software work, in addition to classical hardware for control and use in hybrid algorithms.

These approaches are disruptive in varying degrees to applications and will therefore require cross-disciplinary research to develop algorithms and software solutions, while providing feedback on the performance, usability and deployment issues of specific hardware designs. We will attack the challenges to move beyond Moore’s Law through an integrated research, development and deployment program comprising the following five strategic goals:

1. Develop and deploy application-specialized heterogeneous computing systems to deliver scalable performance to meet the demands of future scientific challenges (Figure 17).
2. Advance algorithms, mathematical models and programming systems for future extreme heterogeneous-accelerated computers that will make the systems capable and usable by scientists.
3. Develop accurate multi-scale modeling and simulation of post-CMOS devices to rapidly evaluate microelectronic technology alternatives in the context of system and application performance.
4. Develop new quantum computing algorithms and software tools for the control and use of quantum systems for emerging scientific applications that are intractable using conventional computing technologies.
5. Architect and deploy quantum computing testbeds using state-of-the-art qubit technologies with a flexible software stack and user support for the science community to enable effective use of quantum technology for scientific discovery.

Figure 17: Programming support for heterogeneous architectures needs to support the variety of specialized hardware. This has been done for low-power devices in the past, as in this heterogeneous, data-centric programming model from the Qualcomm Snapdragon cell phone processor.



1. Heterogeneous Accelerators

Develop and deploy systems with heterogeneous processors that are specialized to accelerate emerging scientific applications.

Specialization is the most promising technique for continuing to provide the year-on-year performance increases that all users of computing systems have come to expect over the last four decades. This creates a particular need for DOE to focus on the unique aspects of scientific computing for both analysis and simulation. Recent communications with computing industry leaders suggest that post-exascale HPC platforms will become increasingly heterogeneous environments. Examples of this trend exist in smartphone technologies, which contain dozens of specialized accelerators co-located on the same chip; in hardware deployed in massive data centers, such as Google's Tensor Processing Unit that accelerates the Tensorflow programming framework for machine learning tasks; in-field programmable gate arrays (FPGAs) in the Microsoft cloud used for Bing search and other applications; and a vast array of other deep learning accelerators.

Heterogeneous processor accelerators — whether they are commercial designs (evolutions of GPU or CPU technologies), emerging reconfigurable hardware or bespoke architectures that are customized for specific science applications — optimize hardware and software for particular tasks or algorithms and enable performance and/or energy efficiency gains that would not be realized using general-purpose approaches. These long-term trends in the underlying hardware technology (driven by the physics) are creating daunting challenges for maintaining the productivity and continued performance scaling of HPC codes on future systems.

Our strategy will be to accelerate the assimilation of emerging heterogeneous specialized technologies that extend far beyond GPUs. We will use workload analysis to identify bottlenecks and opportunities for targeted acceleration, co-develop (with industry) effective accelerator technologies and deploy these technologies in production of effective heterogeneously accelerated platforms that are specialized for mission scientific applications. Our strategy is designed to maximize the impact of these trends on scientific computing.

2-Year Milestones

- Identify at least 3 application candidates by combining CRD core strengths with workload analysis for improved performance using heterogeneous accelerators (shared milestone with Algorithms/Software Thrust)
- Create a methodology to evaluate heterogeneous accelerators.
- Develop a heterogeneous accelerator testbed as part of NERSC-9 (Perlmutter) or as a standalone.
- Initiate 3 different efforts to evaluate the different accelerator options:
 - Lab-led: Accelerator circuit designs that can be implemented either in an FPGA or an ASIC.
 - Co-design with industry: Establish industry design collaborations.
 - Industry-led: Refactor core algorithms for diverse accelerators.

5-Year Milestones

- Evaluate the performance, programmability and efficiency of the 3 potential accelerator options (lab-led, co-design and industry-led) in the context of the selected applications.
- Determine the viability of which of these accelerator options (lab-led, industry co-design and industry-led) should be included in the NERSC acquisition strategy for future extreme heterogeneous-accelerated systems.
 - Initiate FPGA or ASIC development for lab-led accelerator designs *and/or*
 - Initiate NRE development for industry/lab co-developed design *and/or*
 - Initiate NESAP for application refactoring *and/or* algorithm redesign to target industry led accelerators

In the longer term, we expect our co-design methodology to be refined enough that domain scientists, application and software technology developers, hardware architects and industry partnerships will be established as U.S. Government best-practice. Berkeley Lab will drive the HPC ecosystem toward scientifically relevant heterogeneously accelerated solutions, and NERSC's transition to revised procurement methodology to co-develop and acquire such systems will be complete. The transition of at least 50% of the NERSC workload to modern programming models that enable effective use of heterogeneous accelerators will be complete.



Figure 18: In March 2018, Secretary of Energy Rick Perry visited Berkeley Lab to get a firsthand view of how the Lab combines team science with world-class facilities to develop solutions for the scientific, energy, and technological challenges facing the nation. During a tour of the NERSC computer room, he signed the center's newest supercomputer, Cori, which features a combination of traditional and energy-efficient manycore nodes. NERSC's next supercomputer, Perlmutter, will include GPU-accelerated nodes and allow other types of accelerators to be integrated into the HPC network.

2. Algorithms and Software for Accelerators

Develop and advance new mathematical algorithms and software implementations to take advantage of new heterogeneous systems for science.

New software implementations, and in many cases new mathematical models and algorithmic approaches, are necessary to advance the science that can be done with new architectures. This trend will not only continue but will intensify; the transition from multicore systems to hybrid systems has already caused many teams to re-factor and redesign their implementations. But the next step — to systems that exploit not just one type of accelerator but a full range of heterogeneous architectures — will require more fundamental and disruptive changes in algorithm and software approaches. This applies to the broad range of algorithms used in simulation, data analysis and learning.

New programming models or low-level software constructs that hide the details of the architecture from the implementation can make future programming less time-consuming, but they will not eliminate or in many cases even mitigate the need to redesign algorithms. Key elements of a path forward include:

1. Understanding the impact of proposed architectures on current mathematical kernels and algorithms and using this knowledge to steer the HPC hardware deployment choices through feedback in an iterative co-design process.
2. Re-designing current algorithms in response to proposed architectures; hardware choices should be based not only on current algorithms but on the potential performance of new algorithms and even new science use cases.
3. Developing advanced programming environments that ease the implementation of these new algorithms and numerical libraries and are able to generate code for these diverse, heterogeneous accelerators.

The NERSC facility will play a key role in the first thrust, in close collaboration with the mathematicians and scientists who design the algorithms and scientific applications that run on the machines. This will require performance analysis, modeling and optimization tools that allow application teams to map their algorithms and kernels to different potential components in a heterogeneous system, and an understanding of the opportunities and limits for optimization.

2-Year Milestones

- Identify at least 3 application candidates by combining CRD core strengths with workload analysis for improved performance using heterogeneous accelerators (shared milestone with Heterogeneous Accelerators Thrust).
- Select key computational motifs (patterns from key applications and algorithms) to use in assessing heterogeneous accelerator components.
 - Create the initial software proxy applications demonstrating for motifs to use in hands-on testing.
- Determine mathematical models, algorithm options and programming system requirements based on deep analysis of the selected motifs and opportunities/constraints of heterogeneous hardware accelerator components

5-Year Milestones

- Develop and demonstrate programming environment(s) concepts and library/framework interfaces that are used for proof-of-concept experiments to demonstrate superior performance and productivity advantages.
- Demonstrate mathematical / algorithm / hardware co-design targeted at a single, specialized accelerator based on selected motifs.
- Design and release prototype programming environment(s) that enables targeting of multiple heterogeneous accelerators.
 - Analyze performance and energy of combined hardware / software / algorithm for more than one emerging accelerator.

The CRD Applied Mathematics program is critical to the second thrust, in which one can identify two types of applications that will need to be redesigned to run effectively. In the first type, a single computational motif or kernel is paramount, such as stencil computations with fixed spatial patterns. In this case, there is likely to be a single best choice of hardware design. Most of the success stories regarding specialized architectures fall into this category. The advances in numerical methods can be encapsulated in numerical libraries (such as SuperLU and STRUMPACK) and application frameworks (such as AMReX, GraphBLAS and Chombo) to make these advances broadly available to the community.

The second, more complex type is that in which solving the science problem requires fundamentally heterogeneous operations. The heterogeneous operations can be staggered, as one might envision in a data pipeline; as the data moves through the pipeline, different operations are performed on it. In this scenario the data may also be moving physically in steps from source to destination, making the use of different architectures for different stages transparent and separable. Heterogeneous simulation algorithms place a different demand in that, unlike the data example, the flow is more fine-grained and tightly coupled. For example, in a simulation of a time-evolving state or any iterative solution procedure, each step may contain multiple heterogeneous substeps, with each step repeated multiple times, perhaps with different relative (i.e., dynamically changing) costs of the components. No single specialized architecture will be ideal for all stages, suggesting an architectural layout that allows a single code to exploit multiple specialized components. Existing hybrid CPU/GPU systems already allow this, and applications are being re-factored to use this capability; the current trend of offloading different algorithmic components to different specialized architectures will not only continue but become more important.

The Computer Science Department within CRD will play a lead role in the third thrust, which can make use of years of experience writing compilers and runtime systems for Global Address Space languages. This entails the co-design of new compiler technology and domain-specific languages (DSLs) designed around the requirements of the target computational motifs (the 13 motifs that extended Phil Colella's original 7 Dwarfs of algorithmic methods; see "The Landscape of Parallel Computing Research: A view from Berkeley," <https://www2.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.pdf>). The higher levels of abstraction and declarative semantics offered by DSLs enables more degrees of freedom to optimally map the algorithms onto diverse hardware than traditional imperative languages that over-prescribe the solution. Because this will drastically increase the complexity of the mapping problem, new mathematics for optimization will be developed, along with better performance introspection (both hardware and software mechanisms for online performance introspection) through extensions to the roofline model. Use of machine learning/AI technologies will be

essential to enable analysis and automation of dynamic optimizations.

New algorithms favoring less data movement or higher arithmetic intensity, such as communication-avoiding and high-order operators, are already being developed, and data-centric programming abstractions must be built into new PGAS programming systems in order to confer algorithmic information about data locality to the underlying software system. These capabilities are even more crucial for heterogeneous architectures where different accelerators have different memory/communication speeds. More complex algorithms increase the challenges of performance modeling, and tools such as the Roofline model need to be improved to take heterogeneity into account.

Although applied mathematicians must lead the effort to refactor core simulation and analysis algorithms, they should be working as part of collaborative teams containing algorithm, application, software, computer architecture and performance analysis expertise.

Looking ahead, we expect to demonstrate algorithmic re-design of simulation algorithms that target multiple specialized architectures and refine the software prototypes to the point that they can transition to production release and adoption on leading-edge facilities. The goal is to have transitioned at least 50% of the workload over to the new algorithmic methods, libraries and supporting




Figure 19: The ECP ChomboCrunch application is used to simulate subsurface flow (shown here). The team is developing advanced programming support, including DSLs for exa-scale systems, and will need even more aggressive code generation and optimization methods to address heterogeneous accelerators.


3. Multiscale Modeling of Post-CMOS Devices and Systems

Enable DOE to rapidly evaluate emerging CMOS replacement technologies in the context of system and application performance.

Typically, new electronic devices — such as new transistors or memory elements — are evaluated in isolation at a physical level, but this approach fails to capture the architectural-level impact of the device. It is essential to capture metrics that architects and system designers can use to reason about the impact of each to architectures, designs and their complex interactions with existing technologies. Existing hardware design tools do not account for the benefits, and limitations, of future devices. This creates an urgent and immediate need to efficiently and systematically explore the specialized architectural design space in combination with emerging device technologies to avoid stalling performance scaling while waiting for radical new technologies to mature.

The ability to guide development of future devices requires evaluation of their performance based on ultimate outcomes for target applications. The value of new and novel materials or device technologies is not currently understood in a system context. Performance and behaviors at a system context are not currently understood at a device or materials context. True co-design to advance future systems containing novel devices and materials requires feedback that spans all layers, from atomic-scale materials to large-scale complex systems, to meet the needs of emerging scientific applications.

Figure 20: Advanced mathematics is critical to the design and use of future hardware systems, as in this petascale simulation of a fusion tokamak using scalable numerical solvers. Higher order methods can improve computational intensity and make them more amenable to hardware acceleration, while common patterns such as stencils can be “encoded” in hardware.



Our vision is to develop a co-design framework that integrates the physical layers, logical layers and control. We must propagate the quantitative information to guide development of better solutions. The co-design framework would enable DOE to develop Unified Materials→Device→Circuit→System electronic design automation simulation tools to ensure resilience to variability and reduce the development timeline for mission-critical science. The long-term solution requires fundamental advances in our knowledge of materials and pathways to control and manipulate information elements at the limits of energy flow.

As we approach the longer term, we will require groundbreaking advances in device technology going beyond CMOS (arising from fundamentally new knowledge of control pathways), system architecture and programming models to allow the energy benefits of scaling to be realized. A complete workflow will be constructed, linking device models and materials to circuits and then evaluating these circuits through efficient generation of specialized hardware architectural models such that advances can be compared for their benefits to ultimate system performance. The architectural simulations that result from this work will yield better understanding of the performance impact of these emerging approaches on target applications and enable early exploration of new software systems that would make these new architectures useful and programmable.

In the longer term, we will expand the modeling framework to include non-traditional computing models and accelerators, such as neuro-inspired and quantum accelerators, as components in our simulation infrastructure. We will also develop the technology to automate aspects of the algorithm/architecture/software-environment system co-design process so developers can evaluate their ideas early in future hardware. Ultimately, we will close the feedback loop from the software all the way down to the device to make software an integrated part of this infrastructure.

2-Year Milestones

- Develop end-to-end co-design modeling/simulation framework that fills in modeling gaps of projecting from materials and device scale to circuits and system-scale performance.
 - Initially for chip-scale simulation and moving toward system scale.
- Validate co-design framework models end-to-end using existing characterized technologies internally and through industry collaborations.
- Develop hardware generators to enable agile creation and exploration of specialized architectures.

5-Year Milestones

- Use the co-design framework modeling infrastructure to model emerging devices developed by Berkeley Lab and provide feedback to device developers.
- Create the ability to flexibly adapt architectures to optimally use emerging microelectronic device technologies.
- Transfer technology to DOE advanced manufacturing office and industry to accelerate commercialization of promising technologies.

4. Quantum Computing Algorithms and Software Tools

Making quantum computing accessible and productive for scientific discovery.

To make quantum computing relevant for DOE-mission science problems in the near term, we will need to design and build complete applications — incorporating basic and expanded quantum computing algorithmic building blocks — that will leverage the unique characteristics of our scientific problem domains (e.g., symmetries or conservation properties). With this focus, we can develop scalable algorithms and error-mitigation techniques that improve performance while avoiding the large qubit resources needed for full quantum error correction. A key synergy is to work with and use the resources of Berkeley Lab’s soon-to-be-deployed Quantum Computing Testbed, described in the next section, to enable the exploration of algorithm and qubit design-space together and allow the lowest level access to hardware devices.

NOISY INTERMEDIATE-SCALE QUANTUM (NISQ) DEVICES

Quantum computing offers the promise of performance significantly greater than that realizable by classical computing for certain classes of problems. However, the fundamental




Figure 21: A quantum core with 8 superconducting transmon qubits arranged in a ring topology, developed in UC Berkeley’s Quantum Nanoelectronics Laboratory.

elements in quantum hardware — qubits — are currently a scarce and noisy resource. While the full scope of applicability of quantum computing is not completely understood, the areas of chemistry, materials, high-energy physics simulation, optimization and machine learning are some of the most fruitful targets for noisy intermediate-scale quantum (NISQ) devices. NISQ devices are quantum computers with 50-100 qubits that will be available in the next few years, where control over the qubits is imperfect and generates noise that limits what quantum devices can achieve. In the last few years, there have been many proof-of-principle

“quantum computations” in these domains that indicate quantum computing might offer a viable Beyond Moore strategy. Developing and deploying improved qubits is a vital part of the overall strategy, but developing algorithms and software tool-chains to program these devices efficiently is also urgently needed. The close collaboration of both these research thrusts will enable us to co-design hardware, software and algorithms.

Few quantum algorithms showing provable asymptotic speedup over classical approaches have been developed. Many of these algorithms currently function as simple building blocks described only at a general level, often with implementation details missing. At present, much of the research is focused on limiting computational (gate) complexity in the asymptotic regime, assuming both many more, and much higher fidelity, “idealized” qubits than are likely to be available in the next 5-10 years. To fill these gaps, Computing Sciences aims to create new quantum algorithms to address problems of DOE mission relevance, and broaden the applicability and usability of existing algorithms, to make NISQ hardware as productive as possible for the scientific community. New approaches need to be developed to effectively transform mathematical models that emerge from science domains into simple operations that can be executed on quantum hardware platforms. A major mathematical challenge is the choice of algorithm decomposition into fundamental operations while minimizing the accumulation of quantum hardware errors.

In order to ensure the success of quantum computing within DOE, we believe research is needed to design and develop high-level programming languages, compilation infrastructures and domain-science high-level infrastructures. In traditional HPC, a domain scientist can develop at least a basic simulation in a relatively short period of time without knowing many of the details of the underlying hardware. This is not yet the case for quantum computing, and program synthesizers are needed to make it usable to more than a handful of highly trained QIS experts. In the NISQ era, software toolchains need to expose and take into account many more hardware details than are necessary in conventional computing. In conventional computing stacks, layers of abstraction are commonplace and the small losses in performance are acceptable in exchange for a modular software development and user productivity. In contrast, on a NISQ device, a user may need to control in minute detail the exact qubit manipulations required to enable successful execution over failure.

Now and going forward, our efforts are focused on NISQ devices, which we anticipate will dominate quantum computing over the next 10 years. We will develop scalable synthesis algorithms that incorporate error mitigation (an optimal solution for midsize circuits with error mitigation). Longer term, efforts need to be concentrated on providing scalability of the software stack, and developing partial evaluation and manipulation techniques. Our goal is to demonstrate a quantum advantage for scientific simulations in the fields of chemistry, materials and/or high energy physics, with quantum computing delivering a new discovery in one or more scientific fields.

2-Year Milestones

- Develop algorithms and tools tailored to initial NISQ devices and optimized for small numbers of qubits, shallow circuits and targeted at hardware with noisy qubit operation and readout.
- Develop software tools for generating optimal gate sequences given broad information of the noise characteristics of the NISQ device (integration of error-mitigation methodologies).
- Demonstrate new synthesis techniques that can be used for quantum circuit generation when starting from a domain science problem formulation.
- Demonstrate scientifically relevant simulations in chemistry and high energy physics with error mitigation on quantum computers with up to 16 qubits.

5-Year Milestones

- Reevaluate algorithms and tools for hardware with increased capabilities.
- Develop robust program synthesis techniques that incorporate effective error mitigation.
- Develop algorithms that enable simulation of novel new scientific simulations on quantum systems containing 40-60 qubits, elucidate pathways to quantum supremacy and scientific discovery and lead to new approaches for quantum machine learning.

5. Quantum Computing Testbeds

Deploy state-of-the-art hardware targeted at science solutions.

Over the past few years, various implementations of NISQ devices have been built and used for small-scale simulations and benchmarking. A variety of quantum computing paradigms have emerged, including quantum annealing for optimization problems; analog quantum computing, in which a device has similar dynamics to the problem to be solved; and gate-based computing, in which a unitary operator that encodes a program is decomposed into a sequence of operations (or “gates”) acting on single or pairs of qubits.

Berkeley Lab’s Computing Sciences Area has elected to focus on gate-based quantum computing, being the most generally applicable to the modeling and simulation domain space. For gate-based devices, currently the two most promising qubit technologies are superconducting qubits and trapped ions, which have tradeoffs with respect to coherence time, gate execution time, qubit connectivity and scalable fabrication and construction.



Figure 22: Berkeley Lab’s Quantum Computing Testbed will include a dilution refrigerator that will house superconducting qubits. Cooling the qubits to a few milliKelvin is required for superconducting to be maintained in the material and to prevent thermal fluctuations from destroying coherence. Photo courtesy of BlueFors Cryogenics.

Our main effort will focus on superconducting technologies, owing to our strong Lab-wide and university collaborations' expertise in this field. While some commercial near-term NISQ devices are being offered via cloud access, the ability of users to access the hardware at a low-level or propose changes such as connectivity of qubits is severely restricted. The roadmap of these commercial devices will, by necessity, be directed toward either early emerging "killer applications" or general requirements for universal quantum computing, and this roadmap may not be optimal for scientific simulation.

We aim to design, fabricate and deploy best-of-class NISQ devices that are targeted at scientific computing problems, with the ability to employ different tradeoffs than would be made by those focused on a more general application space. The testbed will target users who expect to explore the co-design space for near-term algorithms and toolchains coupled with our hardware. With current design and fabrication techniques, there are many tradeoffs in the design of these early devices in terms of number and connectivity of qubits on a chip, coherence time, gate sets, gate fidelities, etc., and this exposes a rich design space to make devices more broadly useful for scientific computation. We expect to deploy numerous chips with different design points and to custom-fabricate chips in response to the needs of algorithm and toolchain developers. In addition, we will deploy extensible hardware (with designs and fabrication from projects in other laboratory science areas) and software to serve as the complex layer converting digital signals representing gate operations into analog microwave pulses, and analog signal processing technology for qubit readout. A key feature will be flexible interfaces that are required to tune the timing and execution of operations at the lowest level.

We anticipate a number of different categories of users that could make use of the testbed. For example, a user focused on tool development could implement improvements to a compiler based on the low-level noise characterizations measured at the testbed, while a user focused on designing a better interface to classical hardware could substitute in their device to evaluate performance. In addition, an algorithm developer could request a chip with a specific topology or gate set that would make successful execution much more likely.


Our longer-term goal is to enhance testbed capabilities to demonstrate quantum advantage for modeling and simulation application and deploy an extensible testbed infrastructure targeted at selected problem domains in science.

2-Year Milestones

- Deploy a superconducting 16-qubit quantum computing testbed with a low-level software stack enabling access via one or more open-source programming stacks to be used.
- Explore a variety of qubit types in collaboration with other laboratory divisions and university partners, designing and fabricating chips using 3D integration.
- Design and prototype customized hardware for quantum control (e.g., customized instruction-set architecture, digital-analog conversion, and packaging).
- Develop and deploy a low-level software stack allowing for tuning and optimizing testbed operation
- Provide multiple interfaces to the NISQ device, enabling fine-grained control by the algorithm writer.

5-Year Milestones

- Deploy increased capability with respect to qubit count and variety.
- Design and fabricate one or more superconducting qubit chips based on co-design with an algorithm developed by a collaborator in the testbed.
- Deploy customized hardware for quantum control.



Superfacility Initiative

Large-scale analysis of experimental data has revolutionized our understanding of the physical world, from the discovery of elementary particles and the accelerating expansion of the universe to new insights into the microbiome and the ability to better predict extreme weather events. The Computing Sciences Superfacility Initiative is a framework for further integrating experimental and observational instruments with computational and data facilities, bringing the power of exascale systems to the analysis of real-time data from light sources, microscopes, telescopes and other devices. Data from these experiments will stream to large computing facilities where it will be analyzed, archived, curated, combined with simulation data and served to the science user community via powerful computing, storage and networking systems. Tied together with high-speed programmable networking, this superfacility model is more than the sum of its parts, allowing for discoveries across data sets, institutions and domains and democratizing science by making data from one-of-a-kind facilities and experiments broadly accessible.




Figure 23: The superfacility vision involves data streaming from experiments and embedded sensor networks into supercomputing facilities like NERSC using the high-speed networking of ESnet, with a rich set of research questions around the data lifecycle.

The Superfacility Initiative creates an integrated framework to provide automated allocation of compute, storage and networking resources; access to efficient and effective edge computing devices; optimized compute pipelines; and easy-to-use data management tools. Our vision will make it routine to rapidly provision rich, productive and broadly accessible environments for data-intensive discovery, coupling remote instruments to data and computing facilities via advanced networking services. We will demonstrate this model using ESnet and NERSC facilities, while addressing key research and engineering challenges in automation, resilience, specialized edge devices, data lifecycle management and the development of end-to-end science pipelines built in collaboration with domain experts. This initiative also leverages the expertise and research in mathematics, scalable algorithms, data management, programming methods and high performance software that will provide the next-generation analysis capabilities that makes this model a truly

advanced capability for science. It ties together the novel computing technologies from the Beyond Moore Initiative with the methods and software from the Learning Initiative to deliver seamlessly integrated tools for scientists.

While much of DOE's past computing efforts have focused on modeling and simulation, Berkeley Lab Computing Sciences has a deep history of supporting experimental science as well. For example, ESnet has long supported high energy and nuclear physics experiments with bandwidth provisioning and customized services. In 2014, they built a transatlantic extension to Europe, in large part to support the large streams of data from the Large Hadron Collider, which now make up nearly a third of total ESnet traffic. Data transfer architectures designed by ESnet have seen widespread adoption across the ASCR supercomputing facilities (e.g., Petascale DTN project) and the university community through NSF's Campus Cyberinfrastructure program. ScienceDMZ architectures have been deployed across hundreds of universities, as well as throughout the Pacific National Research Platform, a partnership of more than 50 institutions, including the University of California, NSF and DOE.

Similarly, from its inception as a general-purpose supercomputing facility, NERSC has run systems specifically for physics experiments (ATLAS, Daya Bay, STAR, Planck, PTF and others), and recently transferred these pipelines onto HPC systems using containerized software for better efficiency, scalability, performance and economies of scale (**Figure 24**). NERSC also supports science gateways for serving data from cosmology (unWISE catalog from the Sloan Digital Sky Survey), climate observations (the 20th Century Reanalysis Project), material simulations (the Materials Project) and imaging (SPOT suite for ALS tomography beamline, CXIDB for coherent X-ray imaging), among others, and has added real-time queues to support data streaming from experiments.

CRD has built several customized workflow tools and data repositories for light source data, environmental monitoring, biology, cosmology and particle physics. The CAMERA project is an example of a highly successful collaboration to develop end-to-end application pipelines and is in use today at several labs around the world. Jointly funded by ASCR and BES, CAMERA partners with multiple labs to identify areas in experimental science that can be aided by new mathematical insights, develop the

Figure 24: Timeline of prototype pipelines integrating multiple facilities.



needed algorithmic tools and deliver them as user-friendly software to the experimental community. To tackle the data challenges at these DOE light sources and nanoscience centers, CAMERA has been building some of the required infrastructure for algorithm and data curation, as well as efficient user interfaces, real-time streaming, workflows and data-driven environments to meet emerging needs across a broad range of U.S. and international synchrotron and nanoscience facilities.

Thus, while progress has been made toward the vision of using combined experimental, networking and computing facilities, challenges — particularly around seamless use, automation and availability — will require a concerted effort and strategy to address. Today many of the data analysis and workflow pipelines from experiments have been developed for individual, custom applications that are domain-specific and cannot be reused or shared. Efforts to help a new experimental facility port their pipeline and analysis tools to a high performance computing facility like

NERSC remain labor-intensive, often resulting in one-off solutions. A more unified, seamless environment for experimental science — where hardware solutions, analysis software and data management tools are usable by multiple projects and datasets are easily searchable — will provide a profound opportunity to capture insights and accelerate scientific discovery.

Realizing the superfacility capabilities more fully will require significant research, strong partnerships with experimental facilities and innovations in networking, software stacks, compute architectures and high performance computing facilities policies. This section describes Computing Sciences' vision and strategy for creating a superfacility network between DOE facilities and experiments.




Figure 25: The STAR Experiment at Brookhaven National Laboratory is one of several physics experiments already employing the superfacility model to stream data into NERSC and thereby improve efficiency, scalability and performance.



2012



2015



MOLECULAR
FOUNDRY
NCEM

2020

BioEPIC

2021



2025



We see four main goals in achieving this vision:

1. **User Engagement:** Engage with experimental, observational and distributed sensor user communities to deploy and optimize data pipelines for large-scale systems.
2. **Data Management:** Manage the generation, processing, movement and analyses of data for scalability, efficiency, usability and reproducibility and enable data reuse and search to increase the impact of experimental, observational and simulation data.
3. **Automated Resource Allocation:** Deliver a framework for seamless resource allocation, calendaring and management of compute, storage and network assets end-to-end and across administrative boundaries.
4. **Computing at the Edge:** Design and deploy specialized edge devices required for real-time data handling and computation at experimental and computational facilities.

COORDINATION, PLANNING AND MANAGEMENT OF THE SUPERFACILITY INITIATIVE

In support of the Superfacility Initiative, we have created an internal Superfacility Project to coordinate cross-facility tasks and manage the diverse set of activities necessary to fulfill the goals of the Superfacility Initiative. While the Superfacility Project is focused on deploying infrastructure and creating software tools to support production workflows from experimental facilities, it also includes a research council to provide input and directions to the project. Approximately 30-40 staff members from ESnet, NERSC and CRD are involved in the Superfacility Project, which is currently managed by NERSC. The team meets bi-weekly, provides updates to senior management monthly and has created a set of lower-level milestones to track progress on the goals of the initiative.

1. User Engagement

Engage with experimental, observational and distributed sensor user communities to deploy and optimize data pipelines for large-scale systems.

The success of the Superfacility Initiative hinges on close and active engagement with experimental facilities and major science experiments to build tools and optimize pipelines that are of genuine use to the experimental and observational science community. This engagement will fold in the efficient use of high performance computing, data lifecycle tools and general or specialized computing at the edge. These engagements will, in turn, identify the commonalities in the needs of these diverse experiments, which will drive the requirements in the goals of automation, networking, data management and edge device design.

We plan to collaborate with a small number of partners (specified in the milestones) with the goal of enabling scientific discovery using the superfacility framework in such a way that the components developed can be used as the building blocks for other science projects. This will include developing a superfacility API for an application-independent services framework to enable cross-facility access to NERSC data and compute resources to meet challenges such as compute scheduling pattern for experimental facilities (e.g., advance planning, usage pattern during experimental running and resiliency planning), federated ID across facilities and new models for access to NERSC resources, including via an API and Jupyter.



Figure 26: The success of the Superfacility Initiative hinges on close and active engagement with experimental facilities such as the Linac Coherent Light Source at SLAC National Accelerator Laboratory.

2-Year Milestones

- Establish the base set of requirements to support current and future use cases in the superfacility framework by producing white papers describing the computing and data needs for multiple experimental facilities, including the mid-term partner facilities LSST-DESC and LCLS-II.
- Develop the first iteration of a superfacility API for an application-independent services framework to enable cross-facility access to NERSC data and compute resources.
- Demonstrate initial deployment of a prototype infrastructure for four experimental pipelines including federated identity management, scheduling across facilities and failover resilience.

5-Year Milestones

- Demonstrate three new large experiments (including LSST-DESC and LCLS-II) running their workflows seamlessly at NERSC with minimal human intervention, using the framework developed in the 2-year milestones.
- Extend the superfacility framework and infrastructure to meet the needs of multiple experiments across HEP, BER and BES.
- Optimize and demonstrate the framework of automated tools on ESnet6 and NERSC-9 (Perlmutter).

The time frames of the experiments in the DOE complex and at Berkeley Lab provide a natural timeline for the development of the infrastructure and software needed to support their application pipelines. The ALS and the JGI run production workloads at NERSC, and the Dark Energy Spectroscopic Instrument (DESI) and Lux-Zeplin (LZ) will be commencing operations at NERSC before the end of 2020. Some application teams, in collaboration with NERSC, ESnet and CRD, have already developed prototype pipelines integrating multiple facilities, and others, such as DESI and LZ, will be commencing operations at NERSC in 2020. Ensuring the success of these experiments will form the basis of our short-term milestones. The lessons learned from these experiences will inform mid-term engagements with the suite of facilities coming online in the near future, enabling us to build a comprehensive suite of pipeline tools that can be deployed for all new DOE experimental facilities.

In the future, we envision pipelines connecting local compute resources to edge computing to HPC centers, where experimental data is automatically moved and analyzed by tools that hide the location of computing and the types of hardware being used and end users will not require supercomputing expertise. Moving beyond dedicated experiments, the Superfacility Initiative will extend to applications involving sensor networks, including the planned experimental pods at BioEpic, a new facility to be constructed for Biosciences and Environmental Sciences collaborations.

JUPYTER NOTEBOOKS: A TRANSFORMATIVE TOOL

In 2018 the Project Jupyter team was honored with an Association of Computing Machinery Software System Award for Jupyter Notebook, an open-source web application that allows users to create and share documents that contain live code, equations, visualizations and narrative text. Project Jupyter evolved from IPython, an effort pioneered by Fernando Perez, an assistant professor of statistics at UC Berkeley and staff scientist in the Usable Software Systems Group in Berkeley Lab's Computational Research Division.

Today, more than 2 million Jupyter Notebooks are hosted on the popular GitHub service, covering technical documentation to course materials, books and academic publication. Jupyter has been transformative in scientific collaborations and reproducibility, as exemplified by its use at the LIGO observatory, whose discovery of gravitational waves was recognized with the 2017 Nobel Prize in Physics. The LIGO Open Science Center publishes Jupyter Notebooks that allow anyone to replicate their original analyses. Jupyter Notebooks also serves as a core infrastructure for research endeavors like the DOE-funded Kbase platform for predictive biology, the GenePattern Notebook project from the Broad Institute and UC San Diego and the European Union-funded OpenDreamKit project that is building virtual research environments for mathematics.

2. Data Lifecycle

Manage the generation, processing, movement and analysis of data for scalability, efficiency, usability and reproducibility. Enable data reuse and search to increase the impact of experimental, observational and simulation data.

The infrastructure at facilities will need to support management of the entire lifecycle of data from generation, processing, movement and dissemination to reuse. The expanding volume, variety, veracity and velocity of data drive the need to address specific research challenges in the data lifecycle. These challenges include:

- Programming the workflow pipelines
- Collecting provenance
- Managing the data for movement (i.e., filtering, compression and denoising)
- Scheduling and resource management of HPC resources
- I/O and object stores to manage the data on HPC systems
- Innovative analyses and search methods that take into consideration the nature of the application
- Characteristics of the underlying systems.

These challenges need to be met to ensure that workflow pipelines spanning multiple facilities can operate efficiently across facilities. Data produced at experimental facilities or by instruments need to be pre-processed or reduced at the data collection site, efficiently streamed or moved to computing facilities and appropriate resources need to be allocated at the computing facility. Additionally, at present, data is accessed and analyzed primarily by those who generate or produce the data, since it is difficult for others to search and find relevant data sets and confirm their provenance. Capturing provenance of the data and processes, extracting metadata and using indexing to enable search on data will accelerate scientific discoveries through virtual experiments, as well as multidisciplinary and multimodal data assimilation by enabling data reuse.

Recent cross-agency policies ensure that scientific data will be publicly available, but it will require infrastructure to organize, discover, transfer, merge and reanalyze data to make it scientifically useful. The changes


2-Year Milestones

- Develop innovative methods and libraries that allow users to easily process and analyze experimental data using superfacility resources.
- Develop innovative methods and libraries to capture, manage and share analytics pipelines.
- Develop tools and technologies to extract metadata and capture, preserve, share and reproduce the data, workflows and contextual information associated with an experimental setup.

5-Year Milestones

- Develop methods for streaming analytics engines and interactive notebooks that allow near-real-time and interactive workflows to be processed on HPC systems.
- Develop data access, replication and placement services to match hardware architectures with analysis tools to facilitate real-time processing and analytics, whether remote or *in situ*.
- Collect and curate metadata and provenance of the data and pipelines to allow users the ability to reproduce existing workflows.
- Demonstrate analytics and data management tools that run on top of the superfacility API developed under applications engagements to allow experimental scientists to run these tools seamlessly across multiple platforms.

Figure 27: Modern scientific experiments have complex data lifecycles, with data often being used, reused and analyzed at multiple points by multiple teams. This creates a number of challenges for research and software development.



to the underlying hardware at supercomputing facilities will require innovation in analysis methods, I/O management and workflow programming to achieve efficiency and scalability on future systems while maintaining usability; such methods could include *in situ* methods to overcome the I/O challenges, topological data analyses to identify salient methods, indexing, I/O library enhancements and object stores for faster data access, ensemble learning techniques for streaming analyses and interactive exploratory data analyses using Jupyter notebooks. The software stack at HPC centers will have to evolve to support the diverse workloads (i.e., batch and streaming, real-time) while balancing performance and utilization of HPC resources.

Looking ahead, a key aspect of our research and development efforts will be a focus on scale, which will address the data volumes and velocity while leveraging new methods and techniques. In addition, it will be possible for observational sensors to be auto-configured, allowing the real-time streaming data to be seamlessly processed at computing facilities. We ultimately envision a scenario where a scientist conducting an experiment at an experimental facility will be able to seamlessly and interactively search for other relevant simulation or experiment data, run real-time data analysis workflows on HPC machines to influence the experiment and reproduce existing pipelines with different parameters. Users will be able to search and access scientific data from one or more domains, reproduce a workflow, access the appropriate hardware resources and apply complex and advanced data analyses in a turnkey manner.

3. Automated Resource Allocation

Deliver a framework for seamless resource allocation, calendaring and management of compute, storage and network assets across administrative boundaries.

Today, domain science applications and workflow processes are forced to view the infrastructure as static distinct resources (i.e., local compute, remote HPC, storage, network) that require manual coordination to support a complex, multi-facility workflow. There is little ability for applications to interact with infrastructure to exchange information, negotiate performance parameters, discover expected performance requirements or receive status/troubleshooting information in real time. As a result, domain science applications frequently suffer poor performance or require significant manual support to reserve network and compute resources and tune the performance of the end-to-end infrastructure.

As part of its ESnet6 upgrade, ESnet will provide terascale networking necessary for streaming analytics across facilities. In addition, enabling end-to-end multi-domain resource provisioning requires three fundamental operations:

1. Resource exchange information, which includes developing a resource model that can appropriately characterize domain resources (e.g., type, capability, availability, status, etc).
2. Resource requisitioning through an agreed-upon programmatic interface (i.e., API) that is supported by the necessary domain-specific functions (e.g., AuthN/AuthZ, resource computation, provisioning, etc.) to enable inter-domain resource reservation.
3. Performance evaluation using a suite of benchmarks that defines key metrics for measuring the performance of a superfacility.

ESnet will support the first two operations with the deployment of software defined networking (SDN) for leveraging work in the DOE-funded SENSE project, in addition to the existing ESnet bandwidth reservation services (OSCARS) and the Network Services Interface connection protocol adopted by multiple networks supporting scientific research. This functionality will be interoperable with the superfacility API that is actively being developed as part of the applications goal.


2-Year Milestones

- Implement a programmatic framework to securely exchange resource information, and negotiate and orchestrate the end-to-end reservation of compute (e.g. local compute, remote HPC), storage and network resources.
- Deploy SDN for End-to-End Networking @ Exascale (SENSE) resource managers.

5-Year Milestones

- Develop techniques for resource management and optimization to decrease resource blocking, throttle workloads when necessary and provide resilient elastic compute, storage and network resources in real time.
- Adopt learning-based techniques that can detect anomalous resource behavior due to failures or cyber attacks.
- Develop resiliency strategies to prevent loss of data or minimize costly restart of jobs due to unforeseen failures.

Figure 28: A vision of superfacility automation. The ability for a science application to interact and negotiate with compute, storage and network infrastructure would be the hallmark of truly smart infrastructure and smart applications.



Looking ahead, to ensure optimized operations between multiple experimental facilities, ESnet and HPC facilities simultaneously, each resource or application will need to be able to understand the current state of both its own resources and others with which it may need to communicate. To scale the number of experiments and facilities deployed across facilities, it will be necessary to develop intelligent, automated processes that are not restricted by human-scale limitations. Such processes could include instrumenting the infrastructure for performance and availability monitoring, providing input into analytics and learning engines to predict failures and resource congestion and proactively continuously optimizing resource allocations. Security will be enhanced by leveraging learning algorithms to react and protect against byzantine behavior, such as rogue network switches or denial-of-service attacks.

To achieve these goals, we expect to build on a number of current research efforts, such as the ability of ESnet's resources and links to be used optimally in high-traffic and high-bandwidth demands and an ESnet and CRD collaboration to investigate dynamic data movements in situations when NERSC supercomputing resources are down. These research methods will focus on optimizing resource utilization in real time and will be demonstrated through simulated and real infrastructure deployments.

The ability for a science application to interact and negotiate with compute, storage and network infrastructure will be the hallmark of truly smart infrastructure and smart applications (**Figure 28**). Our goal is to implement such capabilities in the future ESnet and NERSC facilities, including ESnet6 and NERSC-9 (Perlmutter).

4. Computing at the Edge

Design and deploy specialized computing devices for real-time computing needs for streaming data.

Experimental and observational data is often collected at sites located a considerable distance from an available HPC center, and the volume and velocity of experimental data threatens to overwhelm both the wide area network and HPC center ingest rates, even with the upgraded ESnet6 capacity. For example, a single microscope under development at the National Center for Electron Microscopy (NCEM) is already pushing the limits of what can be transferred to NERSC over the network; even though a dedicated 400Gb/s network connection was installed directly to NERSC, the detector frame rate of the electron microscopes is still limited by the available network bandwidth.

Just as each experiment today frequently benefits from custom sensors and associated hardware to gather the data, the experiments of the future will benefit from custom hardware solutions to reduce, analyze and respond to collected experimental data in real time. Realizing the superfacility

Figure 29: A new detector is being built for NCEM's Transmission Electron Aberration-corrected Microscope 0.5 (TEAM 0.5) that allows researchers to access single-atom resolution for some samples. The new detector, unveiled in February 2019, will generate 4 terabytes of data per minute.



2-Year Milestones

- Deploy a framework to rapidly design, build and deliver specialized edge computing devices, increasing the impact and availability of custom computing.
- Work with end user facilities to complete an application and workflow analysis to understand the requirements for edge computing devices.
- Leverage a combination of existing high-level hardware description languages and generators with novel programming models to prototype a general framework for rapid creation of FPGA-based edge computing devices.
- Provide a proof-of-concept custom, on-site hardware solution for NCEM, providing data processing/reduction necessary to support higher data rates than are currently possible.

5-Year Milestones

- Enhance the computing capabilities of experimental facilities and DOE HPC centers through the integration of custom edge-computing devices for data reduction and analysis.
- Complete an initial version of a hardware and software generation toolchain that enables rapid design and deployment of at least two edge computing devices in two different facilities.

vision will require hardware deployed throughout the experimental pipeline to increase scientific throughput, automate control of experiments (no human in the loop) and reduce the burden on networks and HPC data centers. At the same time, these specialized analytics accelerators will support increases in the quality (resolution, sampling frequency, signal-to-noise, etc.) of data collected.

Deploying commodity clusters or servers adjacent to experiments is useful for some workflows, but is neither a complete nor scalable solution as commercial platforms may exceed power, space or portability constraints. In addition, deploying a full custom solution for each experiment is undesirable as the cost to design, deploy and support these full custom solutions quickly becomes unwieldy. Field deployable sensors optimized for science have different requirements than commercial devices in terms of reliability, resilience and accuracy and often contain highly specialized functionality. Finally, as sensor data rates continue to increase, it can become impossible to get all measured data off the physical sensor due to the physical constraints of the package – namely, I/O pin limitations. In this case, we must aggressively integrate data-reduction processing directly on die requiring a robust hardware generation toolchain.

To address the unique needs presented by the diversity of experiments, a generalized framework must be developed to allow rapid design, prototyping and deployment of edge computing devices. We will work with science teams deploying sensor networks so that DOE-designed edge computing devices will be integrated directly onto sensors, in the field, in the network and inside the data center. These devices can range from tiny distributed sensors deployed in the environment, such as Smart Dust, to powerful FPGA or ASIC accelerator cards performing data reduction and analysis for large-scale experiments. The creation of this framework creates the opportunity to use novel computing devices, including neuromorphic or other non-Von Neumann processors, to maximize the data reduction and analysis performed for minimal energy.

The use of edge computing devices will help Computing Sciences realize the superfacility vision of ubiquitous computing seamlessly integrated throughout complex scientific workflows. The creation of custom computing devices will be combined with commodity hardware and integrated into complex workflows as transparently as any software-based library, enabling non-experts to design, deploy and utilize these devices.

Summary

Our vision for the future of scientific discovery encompasses more powerful computing, data and networking systems and more tightly integrated facilities for experimental and observational science. But even more important will be the need for human talent to design new methods and algorithms, develop software solutions, partner in cross disciplinary science teams, deploy and operate hardware and software systems and engage with users to empower the broader scientific community.

The three initiatives in this strategy complement our traditional strengths in modeling and simulation, data-intensive science and HPC; moving aggressively into machine learning, digital electronics and quantum computing while having a focused effort in a superfacility model will enable more effective facilities and reproducible science. The combination of exploratory research, team science and engineering practice is well-suited to a national laboratory environment, and the model of driving toward exciting breakthrough science while enabling a broad and open science community is well-aligned with Berkeley Lab's history and culture.



Figure 30: Each year the Computing Sciences Summer Student Program welcomes college graduates and undergraduates from around the world who work with mentors from NERSC, ESnet and CRD over 12 weeks, culminating with a poster session where they present their projects to staff and peers.

Building on an excellent foundation of research and engineering expertise in Computing Sciences, this plan will evolve over several years and require a pipeline of talented personnel to lead and execute. We have a strong track record of hiring outstanding individuals, often as postdoctoral researchers or early career scientists, and mentoring them in the collaborative nature of our work and the value we place on high-impact research to serve the mission of the laboratory. To build the next generation of scientific leaders in computing, we have leveraged a large number of postdoctoral research initiatives, including our prestigious Alvarez Fellowship program, and we are a popular site for summer interns from DOE's Computer Science Graduate Research Fellowship program. Both DOE and the Laboratory management have prioritized the career development pipeline with Early Career research programs run by DOE and as part of the Lab's internal LDRD program.

The pipeline for future leaders starts with students, and each year the Computing Sciences Summer Student Program welcomes college graduates and undergraduates from around the world who work with mentors from NERSC, ESnet and CRD over 12 weeks gaining invaluable science and engineering research experience. Launched in 2010, the program typically draws more than 100 students annually; for many, the highlight is the opportunity to present a poster sharing the results of their summer research projects with their peers and CS staff at the Poster Session that concludes their summer at the Lab.

Team science benefits from diversity in perspectives and expertise, and Computing Sciences is in alignment with the rest of Berkeley Lab in its commitment to foster a diverse workforce — in experience, perspective, and background — along with a culture of equity and inclusion. In support of these goals, we have partnered with the Sustainable Horizons Institute on the Sustainable Pathways Program. This program recruits students and faculty from minority-serving institutions, junior colleges and other colleges supporting students from underrepresented or underprivileged backgrounds for summer research opportunities within the CS organization. As winners of the annual Sustainable Pathways Research Fellowships, these teams spend the summer collaborating with Berkeley Lab staff to further their own research.

Our vision cannot be realized in isolation; it will require partnerships with universities, industry and other laboratories. We believe our initiatives will draw support from across the research community, united by a shared commitment to the mission of DOE and goals of scientific excellence.

PRODUCTION CREDITS:

Project Managers: Helen Cademartori, Katie Antypas, Kathy Kincade

Initiative Leads: Deb Agarwal, Debbie Bard, John Shalf

Strategy Team Members: Ann Almgren, Wahid Bhimji, Ben Brown, Jonathan Carter, Bert De Jong, Doug Doerfler, David Donofrio, Chin Guok, Costin Iancu, Mariam Kiran, Sherry Li, Peter Nugent, Prabhat, Lavanya Ramakrishnan, John Shalf, Dilip Vasudevan, Nick Wright

Photography: Berkeley Lab Photographers

Layout and Design: Creative Services, Information Technology Division

cs.lbl.gov

Computing Sciences Area | Lawrence Berkeley National Laboratory | One Cyclotron Road
Wang Hall, MS59R3022B | Berkeley, CA 94720



U.S. DEPARTMENT OF
ENERGY

Office of Science