


ASCR@40: Highlights and Impacts of ASCR's Programs



2020

A report compiled by the ASCAC Subcommittee on the 40-year history of ASCR for the U.S. Department of Energy's Office of Advanced Scientific Computing Research

ASCR@40: Highlights and Impacts of ASCR's Programs

A report compiled by the ASCAC Subcommittee on the 40-year history of ASCR for the U.S. Department of Energy's Office of Advanced Scientific Computing Research

Chairs, Bruce Hendrickson, Paul Messina

Committee members: Buddy Bland, Jackie Chen, Phil Colella, Jack Dongarra, Thom Dunning, Ian Foster, Richard Gerber, Wendy Huntoon, Bill Johnston, John Sarrao, Jeff Vetter

Editors: Jon Bashor, Tiffani Conner

Contents

1.0 INTRODUCTION.....	1
Document Overview	3
2.1 THE THIRD PILLAR OF SCIENCE: DELIVERING ON THE PROMISE OF COMPUTATIONAL SCIENCE.....	5
Background.....	6
Infrastructure for Computational Science & Engineering	7
Computational Resources	7
Advances in Applied Mathematics	9
Advances in Computer Science	11
High Performance Computing Initiatives and Programs	12
2.2 APPLIED MATHEMATICS: LAYING THE FOUNDATION	18
Differential Equations.....	19
Partial Differential Equations	20
Solvers	25
Optimization	30
2.3 COMPUTER SCIENCE: UNLEASHING SCALE AND SPEED	33
Massively Parallel Processing Computing.....	33
Large-scale data analysis and visualization	38
Deciding who to trust in a world of strangers.....	42
DOE Science Grid, ANL, LBNL, ORNL, PNNL.....	45
Laying the foundations for scientific AI.....	46
2.4 COMPUTER ARCHITECTURE	48
Context.....	48
Paradigm shifts	48
Computer architecture testbeds.....	50

Government and Industry Partnerships.....	52
2.5 ASCR COMPUTING AND NETWORKING FACILITIES IN SEVENTH DECADE OF DRIVING SCIENTIFIC DISCOVERY	54
Early History.....	54
Bringing Computation to Open Science	55
Working in Parallel.....	56
Big Changes in the 1990s	58
ASCR Cements DOE’s Role in HPC	61
Performance Modeling	65
Getting Big Science out of Big Data.....	67
Networking at the Speed of Science: The Energy Sciences Network	71
The First DOE Science-Oriented Network.....	72
Saving the Internet	73
Pushing Networking Technology	74
Evolving to 400 Gbps networks.....	75
Services for Advanced Networking Supporting Data-Intensive Science	76
perfSONAR Performance Monitoring Software	77
ESnet’s Sixth-generation network built on SDN expertise.....	77
ASCR’s Supercomputing Landscape.....	78
Conclusion	79
3.0 ASCR INVESTMENTS HELP DRIVE INDUSTRIAL DISCOVERY AND ECONOMIC COMPETITIVENESS.....	80
Software Libraries Boost Circulation of ASCR Expertise	80
ASCR Computing Facilities Help Drive Industrial Innovation.....	82
Helping Small Businesses Innovate.....	89
4.0 ASCR@40: BROADER ACHIEVEMENTS AND CONTRIBUTIONS	93
Strategic Planning.....	93
Community Standards.....	95
Open Source.....	95
Message Passing Libraries and MPI.....	97
TOP500	100
Research Modalities.....	102
Conclusion	104
5.0 LESSONS LEARNED AND CHALLENGES OF THE FUTURE	105
Lesson 1: A Compelling and Consistent Vision can Drive Scientific Revolutions.....	105
Lesson 2: Diverse Funding Models are Required for Diverse and Impactful Outcomes	105
Lesson 3: Workforce Investments have been Critical	106

Lesson 4: Partnerships are Essential	106
Lesson 5: Testbeds and Platform Access Funding Models are Important	107
Challenge 1: Technology Disruptions	107
Challenge 2: Funding Balance	108
Challenge 3: Software Stewardship	108
Challenge 4: Broader Partnerships.....	108
Challenge 5: A Sought-after Workforce	109
Challenge 6: New Roles for Computing to Advance Science	109
APPENDIX 1: CHARGE TO ASCAC	110
APPENDIX 2: CONTRIBUTORS TO THIS REPORT	112
APPENDIX 3: FOR FURTHER READING	114
APPENDIX 4: ACRONYMS.....	117

1.0 INTRODUCTION

The Office of Advanced Scientific Computing Research (ASCR) sits within the Office of Science in the Department of Energy (DOE). Per their web pages, “the mission of the ASCR program is to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to the DOE.” This succinct statement encompasses a wide range of responsibilities for computing and networking facilities; for procuring, deploying, and operating high performance computing, networking, and storage resources; for basic research in mathematics and computer science; for developing and sustaining a large body of software; and for partnering with organizations across the Office of Science and beyond.

While its mission statement may seem very contemporary, the roots of ASCR are quite deep—long predating the creation of DOE. Applied mathematics and advanced computing were both elements of the Theoretical Division of the Manhattan Project. In the early 1950s, the Manhattan Project scientist and mathematician John von Neumann, then a commissioner for the AEC (Atomic Energy Commission), advocated for the creation of a Mathematics program to support the continued development and applications of digital computing. Los Alamos National Laboratory (LANL) scientist John Pasta created such a program to fund researchers at universities and AEC laboratories. Under several organizational and name changes, this program has persisted ever since, and would eventually grow to become ASCR.

In 1975, the AEC split into the NRC (Nuclear Regulatory Commission) and the ERDA (Energy Research and Development Administration), and the research activities, including Mathematics, became part of ERDA. In 1977 ERDA was combined with the Federal Energy Administration to become the Department of Energy. With the creation of the DOE, the Mathematics program became part of the Office of Basic Energy Sciences. At that time the program had two elements, Applied Mathematical Sciences Supercomputing Research supporting core research in applied math and computer science, and the Energy Sciences Advanced Computation element managing access to hardware and computational technologies at the National Magnetic Fusion Energy Computing Center (NMFEEC). The NMFEEC was created in 1974 at LLNL with support from the Office of Fusion Energy Sciences (FES). NMFEEC was the original DOE supercomputer user facility and was the prototype for others that were to follow. NMFEEC needed connectivity to support users at LANL and Princeton Plasma Physics Laboratory. In 1976 FES supported the creation of the Magnetic Fusion Energy Network, the precursor of today’s ESnet (Energy Sciences Network). In 1979, the Mathematics program was renamed the Applied Mathematical Sciences program to reflect its growing scope of research.

In 1983, DOE created the Scientific Computing Staff with four elements: Analytical and Numerical Methods, Information Analysis Techniques, Advanced Computing Concepts, and Energy Sciences Advanced Computing. This organization was part of the Office of Basic Energy Sciences and was subsequently renamed the Division of Mathematical, Information, and Computational Sciences (MICS).

In 1990, NMFECC was rechristened NERSC (National Energy Research Supercomputing Center) with the broader purpose of supporting the advanced computing needs for all DOE science activities. Additional high performance computing facilities were deployed in 1992 at Oak Ridge and Argonne national laboratories. In 1996, NERSC (its name adjusted to National Energy Research Scientific Computing Center) and ESnet moved to LBNL (Lawrence Berkeley National Laboratory). In 1997, MICS was combined with NERSC and ESnet to form the Office of Computational and Technology Research. Two years later, in 1999, this Office was renamed the Office of Advanced Scientific Computing Research, with a mission to support the DOE Office of Science's most challenging scientific needs.

In its first year, ASCR created the hugely impactful program SciDAC (Scientific Discovery through Advanced Computing) which continues to this day. SciDAC accelerated the adoption of computing across the Office of Science and built strong partnerships between ASCR and the other Office of Science program offices. In recent years, ASCR has partnered with the Advanced Simulation and Computing program within the National Nuclear Security Administration to fund the Exascale Computing Project—the largest computing project in U.S. history.

Throughout this long history, ASCR and its predecessor organizations have adhered to a consistent vision—to create and deploy computational technologies that serve the nation's energy and national security needs. The DOE is widely acknowledged as a world leader in advanced computing technologies and applications, and this is largely due to the advances that ASCR has enabled. Many of these contributions are detailed elsewhere in this document, but some of the highest-level impacts are:

- ASCR's sustained investment in the creation and maturation of the discipline of computational science—now widely viewed as the third pillar of science, on par with theory and experimentation. Computational science has transformed many fields of science and engineering and has had enormous impacts on industrial competitiveness. This impact has permeated DOE through deep partnerships between ASCR and other offices in the Office of Science and beyond.
- The ASCR Applied Mathematics Research Program is second to none and has supported the creation of mathematical approaches, core algorithms, and software tools that undergird computational science worldwide.
- ASCR's origins date back to the earliest days of digital computing and DOE's missions have always demanded the most advanced computing capabilities. ASCR has helped drive the supercomputing field through its research into, and procurements of, state-of-the-art platforms, and through its foundational research in computer science. Research into parallel computing by ASCR's predecessor organizations in the 1980s paved the way for the broader transition to parallel computing in the 1990s. Similarly, ASCR investments helped drive the development and adoption of high-speed wide area networks for distributed computing and collaboration, and more recently for the revolution in data-centric science.
- ASCR's ESnet supports collaboration, distributed computing, data management, and data analysis for scientific communities around the world. In the late 1980s, ASCR researchers created the congestion management protocols that allowed the Internet to thrive and scale.

- ASCR’s computing user facilities are among the best in the world and have been used as models by the National Science Foundation and others. These facilities provide access to computing capabilities that have enabled extraordinary scientific progress and also helped U.S. industry to innovate and compete.
- ASCR has also played a key role in growing and enabling a STEM workforce for the needs of the DOE and more broadly for the U.S. economy. Scientific applications have driven ASCR priorities, but every advance has come from people—world class researchers along with visionary program leaders.

As impressive as they are, these specific contributions miss the overall, integrated impact of decades of investment and leadership from ASCR and its predecessor organizations. Much of the software that drives scientific discovery today has been prototyped, developed, and/or supported by ASCR-funded researchers. The range of these software products includes operating systems; network protocols and libraries; data analysis, management, visualization, and collaboration tools; mathematical libraries; programming models; performance analysis and prediction tools, machine learning frameworks, and many genres of modeling and simulation suites. These software products often build upon each other and interoperate to create a rich ecosystem for scientific exploration and insight.

In a similar fashion, the set of fields and activities that ASCR and its predecessors have combined create a whole that is greater than the sum of its parts. Fundamental mathematical advances are instantiated into software libraries that enable unprecedented simulation capabilities. Advances in programming models underpin message passing libraries and portable performance abstractions to ensure that applications make the very best use of state-of-the-art computers at ASCR facilities. Enormous simulation and experimental data sets are exchanged via data services and collaboration frameworks, managed by input/output tools, and analyzed by visualization and analysis algorithms and tools, all of which were developed and deployed with ASCR funding. ASCR played a pivotal role in envisioning, delivering, and sustaining this rich ecosystem for innovation and discovery.

Document Overview

Through the decades ASCR’s vision and sustained investments have produced extraordinary scientific and technological impacts. The purpose of this document is to record this history and to convey some of the most significant outcomes. Naturally, in a document of this length it is not remotely possible to touch on all the deserving accomplishments, and the authors make no pretense that the result is comprehensive. Rather, the goal is to describe some of the highlights and, through these stories, explain how ASCR has transformed the scientific landscape.

In gathering this information, the authors have reviewed many hundreds of documents and communicated with hundreds of individuals who have lived these stories. The content herein is thanks to these many individuals, and any errors or misrepresentations are entirely the fault of the authors. In Appendix 2 we provide a list of those who helped inform this document, and we apologize for any names we unintentionally overlooked.

This document began with a charge to the Advanced Scientific Computing Advisory Committee which can be found in Appendix 1. The team of authors made both targeted and open requests to the broad community to contribute content using a variety of communication mechanisms. Of particular note was a panel discussion at SC18, the large international conference on advanced computing and networking. The team then worked for many months on the difficult task of selecting a subset of stories that, in our view, best convey the breadth and depth of ASCR's vast impact. This distillation was inevitably somewhat subjective, and only a small fraction of the raw material found its way into this final document, but the rest is also of significance and importance to various communities. We believe it would be useful to provide a searchable, permanent archive for all of the materials we uncovered.

Section 2 of this document is focused on ASCR's most significant technical impacts and is broken into five subsections: computational science, mathematics, computer science, computer architecture, and computing and networking facilities. Section 3 describes some of the impacts on U.S. industry, while Section 4 reviews some broader impacts that do not fit cleanly in other sections. Finally, Section 5 tries to distill some key lessons from ASCR's history that have relevance for thinking about its future. Note that throughout this document, when referring to impacts over the span of decades, we sometimes use "ASCR" as shorthand for "ASCR and its predecessor organizations."

2.1 THE THIRD PILLAR OF SCIENCE: DELIVERING ON THE PROMISE OF COMPUTATIONAL SCIENCE

Computational modeling and simulation are among the most significant scientific and engineering achievements in the 20th Century and are transforming the nature of scientific and engineering research and practice. Computer-based modeling and simulation provides a means of solving the most challenging mathematical equations representing the fundamental physical laws that govern our world. Computational science and engineering can predict the behavior of natural and engineered systems that otherwise could only be described empirically. In fact, many science and engineering problems can only be addressed by harnessing the most advanced computing technologies, e.g., mapping the human genome or predicting the future climate of the earth. All fields of science and engineering have been impacted by the rise of computational science and engineering, from modeling catalytic processes to produce sustainable fuels to modeling the human heart to develop more effective treatments for cardiac problems.

Rapid advances in computing technologies during the past three decades have led to a major revolution in computational modeling and simulation. However, the advances in computing power have come at the expense of computational complexity—the architectures of today’s high performance computers pose a daunting challenge for scientists and engineers developing new computational modeling capabilities. To meet this challenge, the research programs in the office of Advanced Scientific Computing Research (ASCR) evolved to meet the needs of the computational science and engineering community, enabling scientists and engineers to take full advantage of these advancements in computing technologies. Since the 1990s, when these advances began accelerating, ASCR and its predecessors made investments to build the foundation upon which scientific applications for simulating physical, chemical, and biological phenomena can be developed, from implementing efficient and effective algorithms for the mathematical equations solved in these codes to developing the software technologies to take full advantage of increasingly powerful and increasingly complex computing systems. In addition, ASCR established computational facilities to provide the computing and technical support resources needed to address the scientific and engineering challenges that face the U.S. Department of Energy (DOE).

Foundations of Scientific Simulation

Research supported by ASCR program has long provided the foundation upon which scientific and engineering simulation applications are built. Research in applied mathematics and the development of efficient algorithms for fundamental mathematical operations, e.g., solution of linear equations and diagonalization of matrices, for HPC systems has enabled computational scientists and engineers to access state-of-the-art implementations for the applications they develop.

As the transition to parallel computing began in the 1980s, research on computer architectures for scientific simulation identified promising pathways for future development. As the transition picked up speed in the 1990s, ASCR-supported researchers led the push for standardization of the inter-processor communication protocols, which is at the very heart of parallel computing, resulting in the widespread availability of the MPI.


More recently, research supported by ASCR has focused on software technologies that facilitate the development of scientific simulation software for the heterogeneous architectures of petascale and exascale computers as well as handling and analyzing the massive amount of data that these simulations produce. Examples of these development R&D activities include an exascale software stack, an extreme-scale scientific software development kit, software to facilitate moving scientific applications among computing systems, and many more.

This section describes the role ASCR has played in the advancement of the disciplines of computational science and engineering. Beginning with the Manhattan Project and the AEC (Atomic Energy Commission) computing and computational science and engineering have played a prominent role in DOE's research and development program. In 1983, ASCR led the nation in recognizing the potential of parallel computers by establishing the Advanced Computing Research Facility at Argonne National Laboratory (ANL). Once parallel computers began to be more widely deployed, ASCR focused on developing the software technologies needed to create scientific and engineering applications for parallel computers. To ensure that computational scientists and engineers could take advantage of the latest developments in computer science and engineering and applied mathematics, ASCR formed partnerships with other programs in the Office of Science, and elsewhere in DOE, that brought together computational scientists, computer scientists, engineers, and applied mathematicians into teams whose synergy was essential for developing a new generation of computational modeling and simulation applications for the most advanced computing systems. These close partnerships continue today as ASCR's Exascale Computing Project is developing a suite of software technologies and computational science and engineering applications to take full advantage of the extraordinary computing capabilities provided by the exascale computing systems.

Background

Recognition of the enormous potential of computing technologies for advancing science and engineering has been a hallmark of DOE since its birth during the Manhattan Project in the 1940s as the AEC (Atomic Energy Commission). Despite the limitations of the computers available at that time, computational modeling enabled designers to understand the nature of the shock waves needed to detonate the Fat Man bomb. The role of computational modeling grew in importance as scientists began developing the hydrogen bomb and, shortly thereafter, creating a device based on the harnessing of fusion processes in the hydrogen bomb for production of energy.

To support these activities the Mathematics Program was established in the early 1950s at the suggestion of John von Neumann, one of the early pioneers in computing. He recognized that the use of computers to solve problems faced by the AEC required a detailed understanding of their mathematical underpinnings. Although the Mathematics program evolved over the years, renamed the AMS (Applied Mathematical Sciences) program in 1979 and becoming ASCR in 1999, it maintained its focus on enabling the use of the latest high performance computers to solve problems critical to the mission of the AEC and its successors, ERDA (Energy Research and Development Administration), and DOE. What has changed is the nature of those high performance computers and the associated fields of computational science and engineering.



Alvin Trivelpiece

In the early 1980s, Dr. Alvin Trivelpiece recognized that a focused effort was needed in scientific computing and created the Scientific Computing Staff (SCS), subsuming the AMS program. Recognizing the coming revolution in computing technologies, the SCS started the Energy Sciences Advanced Computing program. This program provided advanced computing resources to researchers across the DOE complex as well as those in selected

universities. What began as the Controlled Thermonuclear Research Computer Center at LLNL (Lawrence Livermore National Laboratory) in 1974 (established by Trivelpiece as part of the Controlled Thermonuclear Research Division) eventually became the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL). Today, two other major supercomputing centers, the Argonne Leadership Computing Facility (ALCF) and the Oak Ridge Leadership Computing Facility (OLCF), have been established and are available to scientists and engineers in DOE laboratories, universities, and even the private sector. Together these three centers, have been invaluable in advancing the state-of-the-art in computational science and engineering, providing expertise in computing technologies as well as computing resources.

Building on this early pioneering work in applied mathematics, computer science, and computational science and engineering, the research programs in DOE's Office of Science—Basic Energy Sciences (BES), Biological and Environmental Research (BER), Fusion Energy Sciences (FES), High-Energy Physics (HEP), and Nuclear Physics (NP)—initiated research efforts in computational modeling and simulation designed to address the major scientific challenges in each of their programs. Today, the research supported by the Office of Science is a balanced portfolio of experimental and computational science.

Infrastructure for Computational Science & Engineering

As a result of the Manhattan Project, the AEC, ERDA, and DOE recognized that to meet its mission goals it must:

- Make the most advanced computational resources available to the scientists and engineers involved in its research and development programs, and
- Actively support research in applied mathematics and computer science and engineering essential to continuing the advancement of computational science and engineering.

Computational Resources

Since the Manhattan Project scientific computing has been considered key to achieving the goals of the AEC and its successors with the national security laboratories, LLNL, LANL (Los Alamos National Laboratory), and SNL (Sandia National Laboratories), often pushing these technologies to their limits. Not only were the technologies important, the AEC recognized that an entire computational ecosystem was needed to enable it to address the most challenging problems—computing system software, software tools and technologies, mathematical algorithms and libraries, and science and engineering applications.

The AEC had a history of working closely with the computer industry to develop and deploy the most advanced computing systems available for science and engineering research, beginning with the IBM “Stretch” computer (delivered to LANL in 1961). In 1976 Cray Research shipped the first Cray-1 to LANL. This system provided a peak speed of 160 million arithmetic operations per second (megaflop/s) and, by the end of the next decade the Cray-2 computers provided two gigaflops of

computing power. This increase in computing power enabled computational scientists and engineers to dramatically escalate the fidelity of their simulations. However, as the 1980s drew to a close, it became clear that there were limits to the speed of the technologies upon which these computers were built.

These limitations were recognized in the early 1980s by the SCS who began supporting research into alternate computing technologies. The most promising computing technology involved harnessing the rapidly evolving capabilities of microprocessor technologies. Building computers would now involve hundreds and eventually tens-of-thousands of these processors to work in parallel to perform scientific calculations. Today, ASCR is planning the deployment of computers with speeds in excess of one exaflop/s in the early 2020s—a factor of 10^{10} faster than the heady days of the Cray vector computers. These parallel supercomputers required a substantial change in the design and implementation of scientific codes and the algorithms on which they were based. As the computing power of these systems increased, they became more and more complex with deep memory hierarchies, communications and data technologies, and imbalances in the speeds of the microprocessor technologies. To continue to advance their research and address new challenges, computational scientists and engineers needed new software tools and technologies.

There is a rich history of computing in the national laboratories supported by the Office of Science, from the early efforts at ANL and Oak Ridge National Laboratory (ORNL) on hand-built computers in the 1950s, to the megaflop computing systems built and installed in these laboratories by the Control Data Corporation in the 1960s. In the early 1970s, Dr. Alvin Trivelpiece recognized that a national supercomputing center was needed to advance research in fusion energy. This led to the creation of the National Magnetic Fusion Energy Computer Center (NMFEECC) at LLNL in 1974 (originally called the Controlled Thermonuclear Research Computer Center). In 1983, NMFEECC opened its computer system, a Cray-1, to the other programs in the Office of Energy Research (the predecessor to the Office of Science). By 1990, NMFEECC was allocating computer time to such a wide range of projects from all of the Office of Energy Research programs that the name was changed to National Energy Research Supercomputing Center and in 1996 the center was moved to LBNL and renamed the National Energy Research Scientific Computing Center.

Today, ASCR, the inheritor of the computational vision established by the AEC and nurtured by the subsequent AMS, SCS, and Mathematical, Information, and Computational Sciences (MICS) programs, supports thousands of researchers at national laboratories and universities as well as industrial partners with three major supercomputing user facilities: ALCF, OLCF, and NERSC. Each of these facilities is at the forefront of scientific computing, providing the most advanced supercomputers available as well as the expertise needed to ensure that these resources have the in-house expertise needed to continue the advance of computational science and engineering. In the last 15 years the combined computing power at the three centers has increased by 100,000-fold, which has led to extraordinary advances in physics, chemistry, materials science, climate prediction, and biology. Even medicine has been impacted.

Currently, the ASCR computing facilities at ANL, ORNL, and LBNL are part of a community effort focused on creating an exascale computing ecosystem—computer systems, an operating system, software tools and technologies underlying computational science and engineering, and science and engineering applications. The value of exascale computing is primarily in the new or expanded science and engineering applications it will enable. Exascale computing will also stimulate the emerging integration of supercomputing, big data analytics, and machine learning across a wide variety of science and engineering domains and disciplines.

As the computing facilities were playing an increasingly important role in science and engineering research, ASCR pioneered the development of a high-speed communications network that linked scientists and engineers to its computing centers. The Energy Sciences Network (ESnet) was launched in 1986 and provides access to ever-larger data sets produced in experiments and computational simulations. ESnet delivers connections between scientists and engineers to support collaborative research—an increasing trend as the problems addressed by scientists and engineers become more and more complex. Today, the Office of Science operates the world’s fastest network dedicated to scientific and engineering research, making ESnet another national user facility.

Advances in Applied Mathematics

Mathematical equations describing the physical world are extraordinarily difficult to solve; until the invention of computers their solution was limited to the simplest systems. As the power of computers has grown, so has the ability of scientists and engineers to solve complex equations associated with real-world systems. Research in applied mathematics fostered by ASCR and its predecessors has been instrumental in the development of new approaches for solving these mathematical equations, in the development of mathematical software and libraries for advanced computers, and in solving the broad class of problems known as “optimization”.

Much of the progress in computational science and engineering has involved the design of efficient and accurate numerical methods for solving the differential and partial differential equations involved in computational modeling and simulation. During the 1960s, AEC researchers were interested in modeling the evolution of chemical reactions involved in a number of chemical phenomena such as the combustion of fuel in a gasoline engine. The solution of these equations became unstable unless prohibitively small time-steps were used. Motivated by these problems, applied mathematicians at ANL developed an alternate approach, which eliminated the requirement for uniformly small time-steps. In the 1970s, the GEAR software package, based on this new approach, was developed at LLNL—the first of a series of widely used ordinary differential equation solver packages developed under DOE funding.

By the 1980s, computers had become powerful enough to solve time-dependent multidimensional problems. However, it was highly inefficient to use uniform grids to solve problems that had localized small-scale features, particularly in shock dynamics and nuclear fusion. Investigators at LBNL and NYU (New York University) developed an adaptive mesh refinement technique that enabled the local resolution of the mesh to change automatically as finer and finer details needed to be resolved. The

approach enabled computational modelers to use fine grids where needed and coarser grids elsewhere, reducing the computational effort. In the 1990s-2000s, efforts at LANL, LBNL, LLNL, and NYU extended this approach in multiple directions, enabling it to be applied to a diverse set of science problems.

Since the mid-1970s, ASCR and its predecessors have been involved in the development of mathematical libraries that have had an unprecedented impact on scientific computing. Significant work on these libraries was performed at ANL. Sequentially, AMS, SCS, MICS, and ASCR supported the growth of the mathematical packages that led the way for development of scientific libraries—a large and growing resource for high-quality, reusable software components upon which scientific and engineering applications can be rapidly constructed. These packages provided computational scientists and engineers with improved robustness, portability, and sustainability. The development of these packages—LINPACK and EISPACK, BLAS and ATLAS, LAPACK and ScaLAPACK, to name a few—continues to evolve as computer architectures change, providing a solid foundation upon which computational science and engineering applications can be based.

Finally, research supported by ASCR in applied mathematics has led to the development of algorithms and software for science and engineering applications where the matrices of interest have only a few non-zero elements (there may be billions of such matrix elements). Solving these equations in a straightforward fashion, using dense linear algebra solvers, leads to an explosion in the amount of storage required for the data and in the computational work to solve the equations. The SuperLU package was initially developed at LBNL in the late 1990s for non-symmetric sparse systems on distributed-memory parallel systems. ASCR also funded the development of domain-specific preconditioners to accelerate the convergence of Krylov methods for solving these linear systems of equations. Researchers supported by ASCR have developed robust, general-purpose iterative solvers on parallel computing systems: PETSc (ANL), which initially focused on Krylov-type methods and domain decomposition preconditioners; and HYPRE (LLNL) which focused on multigrid algorithms. These packages are now core elements of the software technologies that support DOE's efforts in computational modeling and simulation.



NWCHEM

HIGH-PERFORMANCE COMPUTATIONAL
CHEMISTRY SOFTWARE

Ushering in a New Era in Computational Molecular Science

In the late 1980s discussions between computational chemists, computer scientists, and applied mathematicians at Argonne National Laboratory led to an effort to develop a new computational molecular science code to take advantage of what eventually became the high-performance computers of today—massively parallel computers. Because these computers require the work to be broken down into chunks and coordinated over the whole computer, this required a completely new design for the code.

The resulting code, called NWChem, was developed in the 1990s with support from both the Biological and Environment Research and Advanced Scientific Computing Research programs. The NWChem team consisted of computational chemists, computer scientists, and applied mathematicians, who worked side-by-side in the development of NWChem to ensure that the latest advances in computer science and mathematical algorithms for parallel computing, most of which was still in development at the time, were incorporated

NWChem has dramatically advanced chemists' ability to predict the properties of molecules. It has also had a major impact on parallel computing—the parallel computing technologies developed by the team have been incorporated in other computational modeling packages. NWChem has been downloaded more than 70,000 times since it was introduced in 1997 and has provided the modeling capabilities needed for more than 3,000 research projects reported in the scientific literature.

Advances in Computer Science

Although research in computer science has been supported by ASCR and its predecessors for decades, its importance increased dramatically in the 1980s and beyond as parallel computing was recognized as the future of scientific computing. Parallel computers posed a number of challenges for computational scientists and engineers that could only be addressed by computer scientists. The first area in which SCS-supported research in computer science led to progress was in the programmability of parallel computers. To efficiently organize a computation on a parallel computer, the computational work must be distributed across all of the processors, meaning the processors must be able to communicate with each other. In early parallel computing, each computer company created its own communication protocol. ASCR-supported researchers overcame the communication barrier by developing two widely used message-passing libraries and then convening and driving a standardization process that defined a standard interface allowing the processors to efficiently pass messages to each other. This message-passing interface was adopted by all of the major computing companies, allowing computational scientists and engineers to develop applications that ran efficiently across all of the computer systems from these companies.

The majority of scientific and engineering codes use mathematical libraries to perform a wide range of standard mathematical operations, e.g., diagonalizing a matrix or solving a set of linear equations. These libraries were developed by applied mathematicians and optimized for each type of computer system to ensure the operations were both efficient and accurate. This problem became acute in the transition to parallel computing systems. Using these libraries, computational scientists and engineers could organize their programs by making calls to library functions that create and manipulate parallel data structures distributed over many processors, greatly simplifying the development of science and engineering codes for parallel computers.

Most scientific applications use or generate blocks of data that cannot be stored in memory. This data must be stored on hard disks or, more recently, solid state disks or on a combination of both. As parallel systems grew in power, so too did the volumes of data to be manipulated. Thus, the speed at which data was written to and read from disks became a potential bottleneck in the simulation. Like other elements of parallel computing systems, the disk systems had to incorporate more and more parallelism as the power of parallel computers advanced. Parallel I/O (Input/Output) systems posed major system design, performance, and usability challenges. In response, the computer science community developed new methods and tools that improved both the usability and performance of the I/O systems. ASCR investments played an important role in the development of both object-based parallel file systems and parallel high-level I/O libraries.

As the volume of data produced by the scientific and engineering simulations increased, the analysis of this data also became a bottleneck. How can a scientist extract insights from gigabytes, terabytes, or petabytes of data? This concern motivated ASCR-funded work on large-scale data analysis (to transform data into results), visualization (to create visual representations of data amenable to human visual inspection), and specialized display devices (to increase the number of pixels delivered to the scientist).

High Performance Computing Initiatives and Programs

Because of the investments that ASCR's predecessors made in the 1980s on the evaluation of future computing architectures, ASCR realized that parallel computers, based on rapidly advancing microprocessor technologies, would dominate the future of computational science and engineering. However, the state of the art in parallel computing was still in its infancy in the late 1980s and early 1990s. This required a change in strategy for SCS. Fortunately, this problem was recognized more broadly by the federal government and the 1991 High Performance Computing & Communications Initiative was funded by the U.S. Congress.

High Performance Computing & Communications Initiative

The U.S. government passed the High Performance Computing Act of 1991 (Public Law 102-194). As a partner in this initiative, ASCR decided to take a new tack in its research programs—to actively support collaborations between applied mathematicians, computer scientists, and computational scientists and engineers. SCS reasoned that such collaborations were essential to enable the computational science and engineering community to take advantage of the rapidly developing capabilities of parallel computers. The success of this program provided convincing evidence of the value of this approach and led to formalization in the Scientific Discovery through Advanced Computing (SciDAC) program.

Scientific Discovery through Advanced Computing

*In March 2000, ASCR, and the other programs in the Office of Science, submitted a report to Congress outlining a major new research program, SciDAC. This Office of Science-wide initiative had three major thrusts: science and engineering applications, software technologies and computing systems software, and computational mathematics algorithms and software. The structure of this new initiative was based on that piloted in the multi-agency HPCC (High Performance Computing & Communications) program.

The goal of SciDAC was to dramatically accelerate progress in computational science and engineering throughout the Office of Science in climate science, fusion energy, high energy physics, nuclear physics, astrophysics, materials science, chemistry, particle accelerators, biology, and subsurface science. SciDAC is now widely recognized as a leading force in advancing computational science and engineering through the use of high performance computing (HPC). SciDAC funding brought together teams of theoretical and computational scientists and engineers, computer scientists and engineers, and applied

SciDAC is the gold standard nationally and internationally for fostering interactions between disciplinary scientists and HPC experts.

2014 Committee of Visitors
Advanced Scientific Computing Research



SciDAC
Scientific Discovery through
Advanced Computing

mathematicians—the triad of disciplines critical to exploiting advances in computing technologies working closely together to achieve breakthroughs in science and engineering.

SciDAC has been embraced by all of the programs in the Office of Science as well as DOE’s National Nuclear Security Administration (NNSA) and the Office of Nuclear Energy, as a model for advancing scientific discovery and engineering practice. Although SciDAC is a modest part of each Office’s research portfolio, its increasing adoption by other offices in DOE, as well as other federal agencies, indicates the growing importance of the SciDAC model. In addition, the success of SciDAC has led to its adoption of a similar approach in many other countries as well.


Through SciDAC partnerships ASCR has supported research by providing funding for computer scientists and applied mathematicians to participate in the computational science and engineering research programs across the Office of Science. In BES, ASCR strengthened six new Computational and Theoretical Chemistry projects by attaching mathematical and computer scientists to the projects. These scientists provided mathematical support for seed computational chemistry initiatives. In particular, the SciDAC team of Car and Klein (2017) joined with the BES-supported density-functional group of Perdew et al. to demonstrate unprecedented accuracy in the quantum-mechanical simulation of water.

ASCR and BER have a long and successful history of jointly supporting numerical aspects of climate and earth system model development. In addition, BER funded work at ORNL to develop a molecular dynamics simulation to improve and test the performance of methods of biomass pretreatment for conversion to fuels and other valuable products.

Advancements in Understanding and Simulation of Earth’s Climate

The ASCR program office, in all its instantiations over the past four decades, has played an important role in the Department of Energy’s commitment to the advancement of climate science. In addition to DOE’s substantial scientific investments in research and analytical efforts on carbon dioxide and global climate change research in the 1980s, DOE computational facilities supported the ever-increasing role of computational simulations in climate modeling beginning at the National Magnetic Fusion Energy Computing Center (NMFEC) in the 1980s, which later became NERSC in 1990. At the same time, MICS partnered with the Atmospheric and Climate Research Division to initiate the Computer Hardware, Advanced Mathematics, Model Physics (CHAMMP) Program, an effort that advanced climate science by fully utilizing the hardware and software capabilities of a new generation of distributed-memory supercomputer systems. These investments led to a decade-long effort in scalable mathematical, algorithmic, and software developments. This work helped transform climate modeling and simulation capabilities and, eventually, led to formal ASCR partnerships with DOE scientific program offices through the SciDAC program. Outcomes of these investments included the development and adoption of advanced mathematics, computational, and software methodologies, along with the development of a rich supporting environment in the form of data storage, visualization, analytical, and dissemination infrastructures (e.g., Earth System Grid Federation (ESGF)). These developments led to the Department of Energy’s central role in the Intergovernmental Panel on Climate Change’s (IPCC) scientific assessment activities which was acknowledged in the 2007 Nobel Peace Prize Award to the IPCC and former Vice President Albert Gore, Jr. These capability-building efforts also proved foundational to the development and recent deployment of DOE’s state-of-the-art global climate modeling system, Energy Exascale Earth System Model (E3SM).

Figure: The high-resolution Energy Exascale Earth System Model (E3SM) captures hurricanes with surface winds exceeding 150 mph. In this image, a large hurricane is visible as a bullseye of white (high water vapor) in the mid-Atlantic, leaving a trail of cold (green)



Taming Fire Extreme Scale Combustion Simulation

High fidelity simulations of combustion allow scientists and engineers to take a peek inside of an engine to examine in detail the nuances of turbulence and its coupling with chemical reactions that convert chemical energy to thermal and mechanical energy. The energy density of liquid hydrocarbon fuels is 100 times greater than a lithium-ion battery and, hence, the world will continue to rely on combustion augmented by other sources of energy for several decades. The holy grail is to design engines that have high efficiencies while also minimizing emissions harmful to humans and the environment.

To maximize efficiency in an internal combustion engine requires exquisite control over the combustion processes. High-fidelity simulations that resolve all turbulence scales and incorporate complex chemical mechanisms require all of the horsepower that high performance computing can provide. Through close collaboration with computational fluid dynamics engineers, computer scientists and applied mathematicians supported by the ASCR's SciDAC, ECI and ECP programs, simulation software for modeling these combustion processes has been developed to take advantage of evolving computer architectures over several decades. The simulations have, over time, encompassed ranges of scale and increased multi-physics complexity, including e.g. dispersed spray and soot formation. At the exascale, these simulations will rival experiments in the level of detail they provide, filling in gaps of knowledge unattainable solely by experimentation. With the insights encapsulated in these predictive models, industry has the ability to optimize engine design, increasing the efficiency, decreasing the cost and shortening the time -to market.




Figure: Volume rendering of HO₂, an ignition marker, and OH, a high temperature marker, in a turbulent hydrogen/air flame.

In FES, the Center for Edge Plasma Simulation (EPSI) SciDAC partnership encouraged, contributed to, and enabled the development of particle resampling algorithms for the XGC gyrokinetic codes, novel sensitivity analysis regimes for XGC runs, accuracy improvements in magnetic flux- and field-following discretizations, and a novel hybridizable numerical method to solve heat diffusion and thermal fluxes in fusion reactors. Other FES advances, enabled by ASCR technologies, include improvements in full-wave electromagnetic field simulations, high harmonic fast wave, and lower hybrid range of frequencies coupled with the computation of non-thermal ion and electron particle distributions in the core plasma in devices such as Alcator C-Mod and NSTX.

Partnerships between HEP and ASCR include making Lattice QCD software available for use across the community and facilitating world-class calculations outside of the big collaborations. SciDAC support of computational cosmology includes the development of astrophysical codes for Type IA supernova explosions and the development of emulators to expand the reach of computational cosmology simulations as well as advancing large-scale parallel

computation, particle methods, and statistical methods for cosmic frontier experiments. In accelerator modeling, SciDAC projects have improved the performance of existing accelerators and have contributed to the design of future HEP accelerators including development of specific codes like Synergia, Ace 3P, and WARP, as well as HEP facilities like the Tevatron, BELLA, and FACET.

A quintessential example of the cross-cutting benefit of ASCR-NP synergy in SciDAC is the utilization of ASCR resources by the Universal Nuclear Energy Density Function (UNEDF) and the Nuclear Computational Low Energy Initiative (NUCLEI) collaborations. Examples of the ASCR/NP SciDAC teams include UNEDF's calculation of all possible isotopes within the limits of nuclear existence and

a new paradigm for performing computational nuclear physics research, and NUCLEI's use of quantum Monte Carlo calculations to constrain the value of symmetry energy and density dependence.

Exascale Computing Initiative

The Exascale Computing Initiative (ECI) is a joint ASCR/NNSA partnership that lays the foundation for development of an exascale computing plan to enable DOE science and engineering applications to make effective use of future heterogeneous exascale computers. One of the primary challenges to achieving exascale computing is designing new computer architectures and new computational science and engineering applications that work under the enormous power and cost constraints established for a capable exascale computer. In close partnership with the HPC industry and academia, the DOE laboratories are conducting multi-disciplinary research to design all aspects of the simulations from numerical algorithms to programming environments to hardware features necessary to facilitate effective physics-based simulations on future exascale architectures—an approach referred to as codesign. In the codesign process scientific problem requirements influence computer architecture design and technology paths, and constraints on these architectures and technology paths inform formulation and design of algorithms and software.

The ECI builds on ASCR's history of successful partnerships between computer scientists, applied mathematicians, and computational scientists. These partnerships are exemplified by the many scientific successes that grew out of the SciDAC era and early vendor/laboratory codesign and co-development of hardware, systems software and programming models, and applications running on cutting-edge computers. LANL and IBM codesigned and built Roadrunner, the NNSA machine that broke the petaflops barrier in 2008, ushering in the era of heterogeneous supercomputers. Soon after ORNL partnered with Cray to codesign ASCR's first heterogeneous system, Titan, which was delivered in 2012 and immediately took first place on the TOP500 list of the world's fastest computers. A set of DOE applications spanning a broad range of computational motifs were refactored to perform effectively on Titan through early-user application readiness programs hosted at the laboratories, another downstream impact of SciDAC.

To ensure that a broad range of science applications could run efficiently on the emerging heterogeneous architectures, DOE recognized that computational science research and development centers must be formally engaged in the hardware, software, and algorithms codesign process. Application-driven codesign required the combined expertise of vendors, hardware architects, system software developers, domain scientists, computer scientists, and applied mathematicians working together synergistically to make informed decisions about features and tradeoffs in the design of the hardware, software, and underlying algorithms. To bring this expertise together, ASCR supported three centers focused on applications codesign beginning in 2011. These centers comprised multi-laboratory, multi-disciplinary teams focused on developing exascale-capable integrated simulation software to design materials in extreme mechanical and radiation environments, ExMatEx; predict performance and safety of a new generation of advanced nuclear reactor designs, CESAR; and control combustion processes with evolving fuel streams to design energy-efficient, low-emission internal combustion engines, ExaCT. In all three centers the core solver(s) were only one component of the required

methodology. Disparity in growth rates of I/O systems and storage relative to compute throughput necessitated a full exascale workflow capability that supported a broad range of in-situ analysis and uncertainty quantification methodologies.

There were a host of complex issues and nuanced details related to efficient and effective use of upcoming exascale computing platforms. While a standard solution would be to address these challenges via a set of narrowly focused efforts, the cross-cutting, multi-disciplinary nature of application driven codesign in the ECI was extremely valuable and significantly more productive at addressing these challenges. By tackling the full scope and interplay of the challenges, ranging from the overall workflow, software development environment, and choice of methods and algorithms, to the nuances of architecture characteristics, the multidisciplinary teams were in continual communication. As a result, the ECI codesign process was invaluable as an initial step toward the development of a productive exascale infrastructure.

Exascale Computing Project

The foundations created by SciDAC and ECI are currently being leveraged by the Exascale Computing Project (ECP). The ECP, another joint effort between the Office of Science and NNSA, is the largest computational science project ever undertaken, with a budget of approximately \$2B spanning FY16-FY23. The ECP was formally established to support President Obama's National Strategic Computing Initiative, an executive order that was issued to ensure continued U.S. leadership in HPC for the benefit of the American economy and scientific discovery. In support of the ECP, ASCR/NNSA established an aggressive goal of developing an entire suite of scientific and engineering applications that can take full advantage of the computing technologies upon which the next generation of high performance computers—exascale computers—will be built. Although the history of the ECP remains to be written, its approach builds on the legacy established over the years in ASCR—building teams of computational scientists, applied mathematicians, computer scientists, and engineers to focus on development of next-generation scientific and engineering applications closely coupled to efforts devoted to constructing software technologies upon which these applications will be built.

ECP's goal is to accelerate the delivery of a capable exascale computing ecosystem and provide breakthrough modeling and simulation solutions to address critical challenges in scientific discovery, energy assurance, economic competitiveness, and national security. To realize this goal the scope of the ECP spans computational application development, computing systems software, hardware technologies, and computer architectures, along with critical workforce development.

Currently, the ECP's project scope spans three focus areas: application development, software technology development, and hardware integration. Application development will deliver a diverse suite of computational applications that effectively utilize exascale platforms to advance the frontiers of simulation science in areas ranging from chemistry and materials, energy, earth and space science, and national security while also emphasizing data analytics, optimization, and codesign. The approximately 25 application development teams build on prior SciDAC and ECI efforts and teams. Whereas earlier ASCR efforts were focused primarily on Office of Science and NNSA science and




Figure: Calcite. This model of calcite pores was created to investigate underground carbon capture. The iridescent grains mimic crushed material geoscientists extract from saline aquifers deep underground to study with microscopes. Researchers want to model what happens to the crystals' geochemistry when the greenhouse gas carbon dioxide is injected underground for sequestration.

mission drivers, ECP has extended its portfolio to include DOE's applied energy areas, e.g., wind turbine technologies and multi-phase reactors for energy conversion, and even non-DOE areas, e.g., deep learning-enabled precision medicine for cancer sponsored by the National Cancer Institute. Software technology efforts will provide a comprehensive and coordinated software stack enabling applications to run productively, efficiently, and portably on diverse exascale architectures. Consistent with ASCR's history of, and commitment to, codesign, close integration of application development and software technology is a key priority of ECP.

When ECP successfully concludes in the next few years, it will have transformed not only the discipline of computational science and engineering but also the many technical domains that are the target of ECP applications. Credit for these successes will be due to the hard work of the ECP teams that delivered them. However, ASCR's long and deep focus on research in computer science and engineering and applied mathematics, coupled with innovative approaches such as SciDAC and codesign, will have played an essential role in this success.

2.2 APPLIED MATHEMATICS: LAYING THE FOUNDATION

The U.S. Department of Energy’s (DOE) Office of Advanced Scientific Computing Research’s (ASCR) Applied Mathematics Research program has developed the mathematical and scientific computing foundations that enabled computational science to impact DOE scientific discoveries. In its nearly 70-year history, the program has made significant advances in applied mathematics essential to modern computational science. The core methodologies and mathematical software libraries are the basis for many of the insights and discoveries made from using HPC. Funding from this program has supported generations of graduate and postdoctoral students who have continued to contribute productively to research organizations in industry, academia, and national laboratories. The combination of these mathematical advances and the resulting software has enabled high performance computers to be used in ways that could only be imagined.

Early in its history the Mathematics program focused on foundational questions related to the properties of models and the numerical accuracy and performance of numerical methods. As computers became more widely available with increasing performance, the emphasis of the math program continually broadened and evolved. Computations that were heroic in one decade became commonplace in the next, opening up new scientific opportunities with fresh needs for mathematical insights. For instance, the growing size of computer memories allowed for larger and larger problems to be simulated, which drove the need for solvers with near-linear scaling properties. As accurate simulations became affordable, new questions like optimization and uncertainty quantification became accessible. All of these advances raised new mathematical and algorithmic questions and this dynamic continues today with the growing utility of machine learning for scientific problems.

In order to successfully undertake scientific simulations, a number of difficult mathematical problems must be addressed. The mathematical model for representing a scientific problem must be well-posed, and specifically stable to the small perturbations as they are introduced by numerical approximations. Similarly, the discretized form of the model must be stable to small perturbations, as well as to perturbations introduced by finite-precision arithmetic. Finally, simulation software should be built using robust mathematical software components whose behavior is well-characterized and implemented for the highest possible performance.

ASCR investigators have developed a body of work to address these issues. It is built on a collection of detailed analysis, going back to the 19th century, understanding the fundamental properties of mathematical problems arising in modeling the physical sciences, along with more recent efforts in understanding the mathematics of computational methods. In order to apply these ideas to the problems arising in modern computer simulation, new tools are constantly being developed—more precise formulations of the mathematical theory as the basis for new computational methods, various combinations of asymptotic analysis, and model problems, and computational experiments to bridge the gap between rigorous mathematical theory and complex scientific problems. Note that there can be a significant time lag between the creation of fundamental mathematical advances and their full impacts on applications. For this reason, in the discussion of foundational work in this section we focus mostly on activities that were conducted long enough ago that their impact is now evident. This

does not mean that more recent mathematical developments are less impressive, just that their impact is less clear due to insufficient time for percolation and adoption.

Much of the development of this set of methodologies for designing and implementing scientific simulation on HPC systems was funded by the applied mathematics program at ASCR and its predecessors going back to the 1950s. At that time, the Atomic Energy Commission (AEC) laboratories built up mathematics programs for simulation and modeling in support of their mission areas. The AEC also funded a mathematics and computing center at NYU (New York University) and these institutions provided, and continue to provide, the leadership in this field.

Differential Equations

As of the 1950s, there had been over a century of development of mathematical theory aimed at understanding the properties of differential equations describing physical problems that exposed a rich and detailed structure to such equations. Much of the progress of computer modeling and simulation in the physical sciences has involved understanding the relationship between the mathematical structure of specific classes of differential equations and the design of efficient and accurate numerical methods. The methods described below constitute a major component of the modern methodology for solving differential equations for scientific applications in which DOE investigators, funded by the applied mathematics research programs, have played a leadership role.

Time discretization methods

During the 1960s, one of the emerging difficulties for AEC researchers wanting to solve complex chemical reaction equations was the appearance of stiffness, i.e., subprocesses on fast time scales that rapidly equilibrate, but whose presence would require the use of a prohibitively small time-step if one used classical explicit methods for integrating the system. Motivated by these problems, William Gear, while working at Argonne National Laboratory (ANL), used discretizations based on backward differentiation formulas (BDF) as a starting point for ordinary differential equations (ODE) algorithms and software that were able to automatically control the time-step for accuracy, while eliminating the requirement for uniformly small time-steps for stability. In the 1970s, the GEAR software package, developed at Lawrence Livermore National Laboratory (LLNL) based on this approach, was the first of a series of widely-used ODE solver packages.

The BDF-based approach was observed to be robust for stiff ODEs, and it was generalized in a straightforward fashion to differential-algebraic equations (DAE), i.e., systems of evolution equations and nonlinear constraints. In the late 1970s, investigators at Sandia National Laboratories (SNL) observed that the existing DAE methods failed to converge on systems arising from fluid dynamics simulations of solar energy problems. The resolution of this problem led to the development of a new theoretical analysis for DAEs, provably convergent methods for a large class of problems, and the development of DASSL (Differential/ Algebraic System Solver), a software package for DAEs. Further developments continued at LLNL through the 1980s and beyond, including support for matrix-

free iterative methods for solving the sparse linear systems arising from large DAE problems, and integration into the LLNL ODE software development efforts.

Spectral Deferred Corrections (SDC) for ordinary differential equations were first introduced in work funded in part by ASCR at NYU in 2000. In this approach, a solution to a system of ODEs is expressed in terms of an equivalent Picard integral equation in time, where the integral is approximated using numerical quadrature. The integral equation is put in a residual-correction form that permits the use of an iterative method for solving the correction equation for the values of the solution at the quadrature points using a sequence of forward-Euler or backward-Euler methods. The systematic use of the integral formulation throughout, combined with the careful choice of quadrature methods, leads to a family of methods that are stable for very high orders of accuracy (up to 20th-order), and stiffly stable for the backward-Euler based methods. This approach has been the starting point for a broad range of new developments in methods for ODEs carried out by ASCR-funded investigators. These include the development of multiply-implicit methods, in which the solution method for the correction terms is split into multiple parts, each of which can be treated implicitly or explicitly; combination with multigrid and multiresolution discretization methods for partial differential equations problems, for more efficient calculation of the correction terms; development of parallel-in-time methods; and application to a variety of stiff problems, such as combustion.

Partial Differential Equations

Hyperbolic conservation laws

Beginning in World War II, there was an intense interest in the calculation of shock waves and other phenomena arising from the dynamics of explosions. This led to the following fundamental questions: what does it mean to have discontinuous solutions to partial differential equations (PDE), and how does one compute them? The then-emerging theory for designing numerical methods for PDEs required the existence of derivatives, which are not defined at shocks and other discontinuities. Starting in the 1950s, Peter Lax and his collaborators at NYU and at the Los Alamos National Laboratory (LANL) provided fundamental answers to both questions, developing the theory for weak solutions for hyperbolic conservation laws and conservation-form discretization methods for such problems. In the 1970s and 1980s, many of the ideas from the mathematical theory were applied to the development of numerical methods such as total-variation diminishing methods and Glimm's method. Of particular note are the second-order extensions of Godunov's method developed at Lawrence Berkeley National Laboratory (LBNL) and LLNL. These methods obtained high-resolution simulations in one or more space dimensions of complex combinations of shocks and smooth flows by combining the use of solutions to the Riemann problem as analyzed by Lax in the 1950s, conservation-form discretization, and limiters and other nonlinear dissipation methods to optimally stabilize the treatment of discontinuities.

Incompressible and low-Mach number fluid dynamics

Velocity-based formulations

The difficulty with the incompressible flow equations for the velocity variables is that they are not pure evolution equations, but a combination of evolution equations and constraints. In the late 1960s, Alexandre Chorin of the University of California at Berkeley and LBNL introduced two new methods for grid-based discretizations of the incompressible flow equations on grids. In the artificial compressibility approach, the steady-state equations are embedded in a well-posed pseudo-compressible system for which the steady equations are a stationary solution. The steady-state solution is obtained by iterating in time, with the compressibility chosen to optimize the convergence properties. For time-dependent problems, the Hodge/Helmholtz projection operator is used to eliminate the constraints, leading to a well-posed, first-order-in-time system of evolution equations. The 1980s and 1990s saw the development of higher-order projection methods at LLNL and LBNL based on the use of high-order Godunov methods to treat advective terms.

Vortex methods

In the early 1970s, Chorin introduced particle methods for representing vorticity transport based on regularization of the singularity of the Biot-Savart kernel and stochastic representation of viscous effects (including no-slip boundary conditions). One of the main advantages to vortex methods is that they admit mathematically rigorous convergence theory, much of which was developed in the 1970s and 1980s at LBNL.

Low Mach-number flows

In fluid dynamics if the Mach number—the ratio of the fluid velocity to the speed of sound—is small, then acoustic waves are of low amplitude, relax rapidly, and thus contribute little to the dynamics. Classically, the zero-Mach number limit leads to the incompressible flow equations.

However, it is possible to have such low-Mach number flows that are not incompressible: for example, in engineering combustion applications, the Mach number is small, but the fluid undergoes large expansion due to heat release. Simulation of such problems as fully compressible flows leads to impractically small time-steps if treated using an explicit method, and conditioning problems if treated implicitly. This problem was addressed in work at LBNL beginning in the 1980s, using the idea that bulk compressibility effects can be represented using equations similar to the incompressible flow

Turning a computer into a numerical microscope for science

When a computer is used to model a scientific problem, it usually breaks the model down into a mesh of uniformly sized cells. In the early 1980s, however, Marsha Berger of NYU and Phil Collela of LLNL began to apply Berger's thesis of AMR to specific research problems in physics, such as determining the shape of a shock wave bouncing off a hard surface.

"This was a problem that had been studied since the 1800s, but with AMR we were finally able to do the right kind of careful calculations and get definitive answers we couldn't get before," said Colella, now at LBNL. "What made this work is we had a scientifically interesting problem we wanted to solve."

Instead of modeling a problem using a uniform mesh, AMR algorithms create cells of different sizes and apply the smallest ones to the most interesting parts of the problem, which allows scientists to study those aspects in greater detail. Collela has likened this to a "numerical microscope." For example, in a combustion simulation, the critical aspect is the flame front as it moves across a combustion chamber so AMR focuses on that portion, adapting and applying the finer resolution to the moving front as it burns.

Berger and Collela described their AMR solution in a 1989 paper "Local adaptive mesh refinement for shock hydrodynamics" that was published in the *Journal of Computational Physics*. Today, AMR is used worldwide and in a wide array of modeling applications, from combustion to Antarctic ice sheets to supernovae. Collela credits the success of AMR to strong collaboration between the math and computer science communities and the long-term collaborative support of ASCR and the Office of Basic Energy Research over the long term.

equations, but with an inhomogeneous divergence constraint in regions when there is bulk expansion due to thermal sources. This approach was put on a firm analytic footing with a systematic derivation using asymptotics of the zero-Mach number limit for reacting fluid flow in open or closed vessels. Further work at LBNL and LLNL in the 1990s and 2000s included the extension of the high-order projection methods to these equations, as well as their extension to more general settings in astrophysics and gravitationally stratified flows.

Adaptive mesh refinement

By the 1980s, computers had become powerful enough to solve time-dependent multidimensional problems. The basic approach is to create a discrete representation of the spatial domain of interest using a grid (also referred to as a mesh). Derivatives in the PDE are computed using a combination of adjacent values of the approximate solution at discrete locations on the mesh. A structured grid is one where the grid-points lie in a regular array; an unstructured grid allows for essentially arbitrary connectivity of the grid-points.

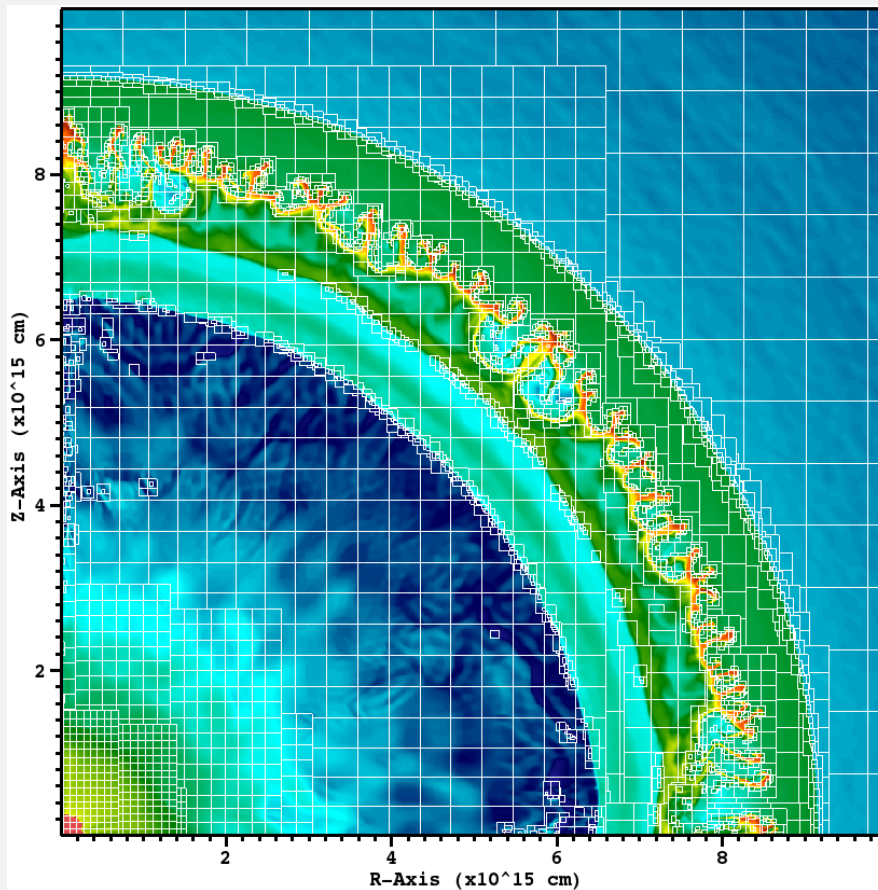


Figure X: This 2D simulation of a supernova illustrates the use of adaptive mesh refinement to focus computing power on areas of particular interest, such as the Rayleigh-Taylor instabilities in the outer shell of the star, by using finer mesh resolution. The simulation was created using Castro, a compressible hydrodynamics code for astrophysics (then based on BoxLib, now on AMReX).

As scientists began to study more challenging problems, it became clear that using uniform grids could be highly inefficient in situations where localized small-scale features appear in multiple dimensions, particularly in shock dynamics and astrophysics. Investigators at LBNL and NYU developed adaptive mesh refinement (AMR) for locally changing the resolution of the mesh automatically as a function of space and time. The approach was based on using nested rectangular patches, which permits efficient implementation and minimizes computational overhead due to irregular data placement. In the 1990s and 2000s, efforts at

LANL, LBNL, LLNL, and NYU extended this approach in multiple directions, including the application to elliptic PDEs, incompressible flows, low-Mach number models for reacting fluid flows, and Cartesian-grid methods for complex geometries, with applications to a diverse set of science problems. Software frameworks (BoxLib, Chombo) for high-performance scalable implementations on distributed-memory multiprocessors to facilitate reuse across applications was funded by ASCR through the High Performance Computing and Communications (HPCC) and Scientific Discovery through Advanced Computing (SciDAC) programs.

Complex geometries and fronts

In the 1980s, as multidimensional simulations became feasible, it was necessary to address the challenge of representing complex boundary geometries and propagating fronts. Both the spatial discretization of PDEs in irregular geometries, and the representation of the geometry and motion of boundaries themselves, were addressed.

Discretization of PDEs on unstructured grids

Methods based on unstructured grids were originally developed for structural mechanics in the 1950s. Starting in the 1990s, ASCR-funded investigators at ANL developed methods for time-dependent fluid-dynamics problems based on high-order-accurate finite element-type discretizations (“spectral-element methods”), with a particular focus on scalable and high performance implementations on distributed-memory multiprocessor systems.

Composite overlapping grid methods for PDEs

This approach, developed simultaneously by both by NASA investigators and ASCR-funded researchers at LANL and LLNL, uses a set of structured curvilinear grids that overlap in space, and together define a complex geometry. Finite difference or finite volume methods are used to represent the PDEs, and information is communicated between the overlapping grids using interpolation. Overlapping grids have the advantage that boundaries are represented accurately, high-order derivative approximations are straight-forward to compute, and problems with moving boundaries, for example those describing the interaction of fluids and solids, can be computed very effectively. Further efficiency is provided by combining this approach with AMR. Applications have included the study of detonation waves, problems involving deforming bodies, and aerospace applications.

Cartesian-grid finite-volume methods for PDEs

This is an alternative approach for representing complex geometry that uses Cartesian grids instead of curvilinear grids, thus considerably reducing the algebraic complexity of the difference schemes from general curvilinear grids. Instead, using an idea first suggested in the 1960s, the computational domain is represented by intersecting the irregular domain boundaries with a Cartesian grid. The complexity of generating arbitrary curvilinear grids is thus replaced with a lower dimensional problem of computing intersections of the boundary with the Cartesian grid. This method was combined with AMR in the 1980s by ASCR investigators at LBNL, LLNL, and NYU to turn it into a broadly applicable approach. While conceptually simple, the method requires complex automatic calculations of intersections between control volumes and arbitrary boundaries, as well as sophisticated approaches

for maintaining stability of the method. These challenges have been successfully addressed, as well as the extension of the method to problems with moving boundaries. Applications include a variety of problems in geophysical fluid dynamics and aerodynamics. The latter application includes the Cart3D code developed by NYU and NASA Ames Laboratory, which is a heavily used production code for NASA and the Department of Defense aerodynamics applications.

Level-set methods for propagating fronts

The initial problems motivating this work were the propagation of premixed flames in reacting fluid flows, and the growth of crystals, both represented as a discontinuous interface propagating at a scalar speed in the direction normal to the front. Through work funded by ASCR at LBNL in the 1980s it was recognized that the theory and numerical methods for hyperbolic conservation laws could be applied to the Hamilton-Jacobi equations for a continuous function in space whose zero level set specifies the location of the front. The level-set approach is the leading method for computing propagating fronts in a broad range of problems, including fluid dynamics, materials science, semiconductor processing, image analysis, grid generation, and computational geometry, with contributions made at LBNL in these and other areas. LBNL has also made ongoing contributions to the mathematical and numerical methodology for level set methods, including the development of more efficient narrow-band methods and of a fast solver for the steady-state form of the equations (eikonal equations).

Fast integral-equation-based methods

A broad range of problems in physics include the discrete or continuous convolution of a charge distribution with a long-range kernel, e.g., the computation of electrostatic or gravitational forces on a collection of particles, or the solution of Poisson's equation using integral equation methods. The naive approach to this computation is to evaluate the field by a direct computation of the convolution, leading to methods that scale like the product of the number of quadrature points in the charge and the number of field evaluation points for the potential. The Fast Multipole Method (FMM) exploits the smoothness of the long-range coupling to construct systematic approximations that reduce the cost to linear or log-linear in the number of quadrature points / field evaluations. Beginning in the 1990s, ASCR-funded investigators at NYU, led by one of the original developers of FMM, developed a variety of extensions and applications, including an optimal version (in terms of accuracy and efficiency) of the algorithm for the Newtonian potential in 3D; extensions to other integral kernels for constant-coefficient PDEs, such as the heat equation, Stokes' equation, and the Helmholtz equation; and applications of this approach to construct fast, scalable, low-communications solvers for constant-coefficient elliptic PDEs on rectangular and AMR grids.

Compatible discretization methods for PDEs

Discretization transforms a PDE living in an infinite dimensional function space into a finite dimensional algebraic problem that can be solved on a digital computer. It is inevitable that some information about the governing equations and their structure is "lost in translation" during the discretization. Compatible, or mimetic, discretizations aim to ensure that the discretization process preserves key structural properties of the underlying PDEs. In so doing, compatible and mimetic

discretizations provide a rigorous basis for building physically consistent, stable, and accurate discrete models.

The importance of structure preservation was recognized early on by many scientists and engineers working on specific application problems. Perhaps the best-known examples are the Yee scheme for the Maxwell's equations developed by Kane Yee in 1966 at the Lawrence Radiation Lab and the MAC scheme of Harlow and Welch developed at LANL in the 1960s. These schemes revolutionized simulations of electromagnetic and fluid phenomena by providing physically meaningful solutions free of spurious numerical artifacts.

However, it was not until the late 1980s and early 1990s that scientists and engineers began to recognize that many seemingly different methods across multiple disciplines share a common provenance that is responsible for their exceptional properties. At the forefront of this work was a group of scientists at LANL who pioneered the mimetic finite difference method and established one of the first general mathematical frameworks for compatible discretizations.

With ASCR support in the late 90s, research in compatible discretizations intensified, culminating in a workshop sponsored by SNL, LANL, and LLNL that took place in 2004 at the Institute for Mathematics and its Applications. This workshop marked the emergence of compatible discretizations on unstructured grids, including polygonal and polytopial grids, as a key enabling technology for DOE.

In the following years ASCR funded a comprehensive research portfolio in compatible discretizations leading to their adoption across multiple mission areas. Some examples are the Amanzi Advanced Terrestrial Simulator based on mimetic finite differences on polygons, and the Model for Prediction Across Scales (MPAS) which is developing models for the DOE's E3SM Earth System Model based on compatible discretizations on unstructured Voronoi meshes. In 2007-2012 ASCR-funded efforts at SNL led to the development of modern software tools supporting high-order compatible discretizations on a wide range of unstructured grid, such as the Intrepid package in Trilinos, which is currently being used in simulations of complex multiphysics problems at SNL and elsewhere.

Solvers

Since the mid-1970s, ASCR has been involved in the development of mathematical software which has had a tremendous impact on scientific computing. ASCR invested heavily in the progression from the physics to the mathematical model to the algorithmic decision and then to the software implementation. ASCR-funded mathematical software that was developed into packages led to the development of scientific libraries that provide a large and growing resource for high-quality, reusable software components upon which applications can be rapidly constructed—with improved robustness, portability, and sustainability. These activities of developing robust packages continue to evolve as computer architectures have changed and provide the basic foundation on which scientific computing is performed.

Direct solvers for dense linear systems

LINPACK and EISPACK

LINPACK and EISPACK are Fortran subroutine packages for matrix computation developed in the 1970s at ANL. Both packages are a follow-up to the Algol 60 procedures described in the Handbook for Automatic Computation, Volume II, Linear Algebra, referred to now as The Handbook. Between 1965 and 1970, J. Wilkinson and 18 of his colleagues published a series of papers that offered procedures for solving different types of linear equations and eigenvalue problems. These research papers presented results about numerical stability, subtle details of implementation, and, in some cases, new methods. The selected algorithms represented the best available in terms of their generality, elegance, accuracy, speed, and economy of storage. The Handbook is a landmark in the development of numerical software and organizing library algorithms for the eigenvalue problem for dense, stored matrices. A team led by ANL rewrote the Algol 60 version of EISPACK in Fortran, being careful to retain the structure and the numerical properties. Wilkinson visited ANL for a few weeks every summer and provided the ANL team with advice about numerical behavior and testing procedures.

In 1976, as EISPACK was nearing completion, a second project was funded by the DOE and the National Science Foundation as a research program to investigate methods for the development of mathematical software. The resulting package, called LINPACK, was designed to solve linear equations and linear least-squares problems. One of LINPACK's most distinctive characteristics is its efficiency, which is achieved through two features: the column orientation of the algorithm, and the use of the Level 1 BLAS. When LINPACK was designed in the late 1970s, the state of the art in scientific computers was pipelined scalar processors, like the CDC 7600 and the IBM 360/195. On scalar computers, LINPACK also gains efficiency from the use of the Level 1 BLAS when large matrices are involved. This is because doubly subscripted array references in the inner loop of the algorithm are replaced by singly subscripted array references in the appropriate BLAS. The effect can be seen in matrices of quite small order, and, for large orders, the savings in terms of computational time are significant. Finally, improved efficiency can be achieved by coding a set of BLAS to take advantage of the special features of the computers on which LINPACK is being run—either by producing machine language versions or by taking advantage of features like vector operations. On vector machines, like the Cray-1, the loop was a single vector instruction. The final version of LINPACK was generated by the TAMPR (Transformation Assisted Multiple Program Realization) system. TAMPR was a powerful software analysis tool, developed at ANL, that manipulated and formatted Fortran programs to clarify their structure. The single (S), double (D), and complex (C) versions of LINPACK were produced by TAMPR.

BLAS and autotuning

The original BLAS, now called Level 1 BLAS, defined by a group at NASA's Jet Propulsion Lab, perform low-level operations like dot-product calculations and adding the multiple of one vector to another. The BLAS promote efficiency by identifying frequently occurring operations of linear algebra that can be optimized on various computers (e.g., by coding them in assembly language or otherwise taking advantage of special machine properties). In the 1970s, using these optimized operations could yield dramatic reductions in computation time on some computers. ASCR supported the development

of the released version of the Level 1 BLAS software. As computer architectures became more sophisticated in the 1980s, it became clear that Level 1 BLAS operations were not going to provide the kind of performance that the architecture was capable of achieving. Architectural features, such as cache, require a coarser level of granularity in the algorithms to take full advantage of the highest levels of performance. With support from ASCR, ANL led an international effort to improve the efficiency of BLAS in light of improvements in computer memory systems. As a result, Level 1 BLAS operations were augmented from 1984 to 1986 with Level 2 kernel operations for vector-matrix operations. Memory hierarchy was also recognized as something to exploit. Today, computers have cache memory that is much faster than main memory, and keeping matrix manipulations localized allows for better use of the cache. In 1987 and 1988, the Level 3 BLAS were identified for doing matrix-matrix operations. The Level 3 BLAS encouraged block-partitioned algorithms.

DOE's applied math program provided the research funding for ATLAS, a program generator that focuses on the BLAS operations. ATLAS is often recommended to automatically generate an optimized BLAS library. While its performance often trails that of specialized libraries written for one specific hardware platform, it is often the first, or even the only, optimized BLAS implementation available on new systems and is a large improvement over the generic Fortran BLAS available at Netlib. For this reason, ATLAS, which was one of the first widely used program generators, is sometimes used as a performance baseline for comparison with other products.

LAPACK and ScaLAPACK

LAPACK is written in Fortran 90 and provides routines for solving systems of simultaneous linear equations, least-squares solutions of linear systems of equations, eigenvalue problems, and singular value problems. The associated matrix factorizations (LU, Cholesky, QR, SVD, Schur, generalized Schur) are also provided, as are related computations such as reordering of the Schur factorizations and estimating condition numbers. LAPACK supports dense and banded matrices with real or complex numbers in both single and double precision.

The original goal of the LAPACK project, funded in part by ASCR, was to make the EISPACK and LINPACK libraries run efficiently on shared-memory vector and parallel processors. On these machines, LINPACK and EISPACK are inefficient because their memory access patterns disregard the multi-layered memory hierarchies of the machines, thereby spending too much time moving data instead of doing useful floating-point operations. LAPACK addresses this problem by reorganizing the algorithms to use block matrix operations, such as matrix multiplication, in the innermost loops. These block operations can be optimized for each architecture to account for the memory hierarchy and so provide a transportable way to achieve high efficiency on diverse modern machines. LAPACK routines are written so that as much of the computation as possible is performed by calls to the BLAS, particularly Level 3 BLAS.

The ScaLAPACK library, also funded in part by ASCR, includes a subset of LAPACK routines redesigned for distributed memory computers using explicit message passing with message passing interface for inter-processor communication. LAPACK is one of the most successful and influential software packages produced under the ASCR program. LAPACK is used by hardware vendors in their

software libraries (IBM, Intel, Cray, NVIDIA, AMD, Apple, ARM, Atos (Bull), Fujitsu, Hitachi, Huawei, Sugon, and Sunway) and by numerous software companies (MathWorks, Mathematica (Wolfram), ANSYS, NASTRAN, MSC Software, Siemens PLM Software, ActiveState, Enthought, Numfocus, SIMULIA, ABAQUS, Synopsys, Julia Computing, Anaconda (R), NAG, Red Hat, ProNova, and Facebook).

Algorithms and software for sparse systems of equations

Many applications, including most PDEs, lead to large linear systems that are sparse, that is, most of the matrix elements are zero. A tremendous amount of work and memory can be saved by exploiting this fact and avoiding computing with or storing zero values. Dense factorization methods like those in LAPACK are often inefficient on these problems since they require memory that scales as the square of the number of unknowns, and computational work that scales as the cube of the number of unknowns. As computer memories have grown vastly bigger in recent decades, problem sizes have grown commensurately. This has led to a compelling need for solution methods whose memory and computational complexity grows much slower than dense factorizations.

There are two approaches to dealing with this problem. One is to use factorization-based methods, like LAPACK, while paying careful attention to avoiding storing or operating on the many zero values in the matrix. The other is to approximate the solution by the result of an iterative process that involves only successive applications of a matrix-vector multiplication. In both approaches, the storage and the computational work scale more favorably with the number of unknowns, with a tradeoff between generality, numerical stability, and the ability of exploiting the problem-specific properties of the system to improve efficiency.

Direct methods for sparse systems

Early development of direct, factorization-based methods for sparse linear systems were driven by the requirements for finite element discretizations of self-adjoint elliptic PDEs, which lead naturally to symmetric positive-definite (SPD) matrices. The earliest ASCR contributions in this area were at ORNL (Oak Ridge National Laboratory) beginning in the 1980s, with the development of new sequential and parallel algorithms and software for SPD problems. Beginning in the 1990s, the SuperLU package was developed at LBNL for non-symmetric sparse systems on distributed-memory parallel systems, with ASCR support coming under SciDAC. The design tradeoffs for the non-symmetric case are more difficult than in the SPD case due to the conflicting requirements of data locality (both for parallel efficiency and reducing the number of nonzeros in the LU decomposition) and of minimizing the effects of roundoff. SuperLU has emerged as a widely-used standard for sparse matrix calculations. First released in 1997, SuperLU is downloaded more than 30,000 times a year and has been included in many vendors' mathematical libraries, including Cray LibSci, FEMLAB, HP's MathLib, IMSL, NAG, OptimaNumerics, Python (SciPy), and Interactive Supercomputing. SuperLU has been adopted in industrial mathematical libraries and simulation software, including AMD (circuit simulation), Boeing (aircraft design), Chevron, ExxonMobil (geology), and Walt Disney Animation. Sparse direct methods like SuperLU scale much better than dense methods, but the storage and computing requirements still grow much more than linearly with the problem size.

Iterative methods

An alternative class of solutions methods does away with factorizations altogether and instead iterates toward a solution by repeated applications of matrix-vector multiplication. These so-called iterative methods have the potential to be more efficient for large problems, but they are less robust and more general than factorization-based approaches. ASCR researchers in academia and across the DOE labs have played a key role spanning several decades in developing and maturing iterative methods. Some of the early work focused on the foundational methods and their convergence properties. A key factor in the successful application of these methods is the development of a problem-specific preconditioner, a method to solve the system very quickly, but only approximately. ASCR researchers have been trailblazers in developing high performance preconditioners for a range of physical phenomena in complex geometries. ASCR has funded the development of domain-specific preconditioners at ORNL, ANL, and SNL. These techniques have allowed iterative methods to be successfully applied to science and engineering problems in electromagnetics, complex fluid flows, and other applications of critical importance to the DOE. Importantly, these preconditioners have demonstrated good parallel performance on very large applications running on the biggest DOE computers.

A very important class of methods that can be used as solvers or as preconditioners is known as multigrid methods. The core idea in multigrid is to approximate a problem by a series of simpler and simpler representations (multiple grids), and then to carefully orchestrate solutions of simpler problems to drive corrections to an approximate solution for the full problem. It was known in the 1970s that multigrid could achieve optimal, linear scaling in both memory and computational complexity for very simple problems. Researchers at the University of Colorado, LLNL, and SNL developed a series of algorithmic and software advances in the 1990s and 2000s that dramatically increased the practical robustness and applicability of multigrid methods. Thanks to ASCR-funded advances, multigrid is now the preconditioner of choice for many real-world problems in structural mechanics, electromagnetics, fluid mechanics, and more.

These algorithmic advances in solver technologies have been encapsulated in state-of-the-art software libraries that are used around the world. Beginning in the 1990s, many of the labs have developed parallel iterative solver libraries and used them to conduct research in preconditioning. ANL is a leader in software, particularly with the PETSc (Portable, Extensible Toolkit for Scientific Computation) project, an ambitious effort to build an open-source infrastructure to support scientific simulation. PETSc is widely-used in DOE and in academia and industry. In the 2000s, SNL began developing an alternative framework called Trilinos that provides solvers but also support for optimization and derivative computations. In the 1990s and 2000s, LLNL and SNL began developing advanced multigrid libraries known as HYPRE and ML, respectively, under ASCR math funding. In the 2000s and 2010s, a series of SciDAC Institute projects built common interfaces and interoperability amongst the various sparse direct, iterative, and multigrid solvers within the DOE labs. At the beginning of SciDAC in 2001, applications scientists mainly implemented their own iterative solvers; by the end of the decade, they almost entirely used one or more of the community tools.

Eigenvalue problems

Sparse eigenvalue problems are an important computational component of DOE scientific simulation, such as in the areas of R-F accelerator cavity design, electronic structure calculations, and nuclear physics. In the 2000s, the SciDAC program funded mathematicians at LBNL to develop high performance domain-specific eigenvalue solvers for distributed-memory HPC systems. The starting point for this work was the existing software for computing a small number of eigenvector/eigenvalue pairs using Arnoldi's method as implemented in SNL's ASCR-funded ARPACK. In order to meet the requirements of the applications, a number of ideas and software components—domain decomposition, the use of shift-inverse Lanczos, and high performance direct solvers such as SuperLU—were combined in highly innovative ways to exploit specific features of the applications. In addition, some problems, such as electronic structure, are actually nonlinear eigenvalue problems, leading to mathematically sophisticated integration of the nonlinear dependence into the algorithms. The sparse eigenvalue calculation is one of the major application impacts of computational mathematics from the SciDAC program.

Optimization

A number of DOE applications involve finding a set of inputs that result in the best possible output. The field of applied mathematics that addresses this wide class of problems is known as “optimization.” General optimization problems are quite hard, but many problems have a special structure that allows for effective algorithms. ASCR has invested in optimization mathematics, algorithms, and software since the 1970s. As computers have become vastly more capable through the decades, the applications of optimization and related methods have grown enormously.

Optimization algorithms and software

With the growing availability of computers in the 1970s, optimization was a rapidly advancing field and ASCR supported groundbreaking research at ANL. This early work was focused on applications in which the function being optimized could be written as a simple computer code. Developed in the 1980s at ANL under funding from one of ASCR's processor organizations, trust region approaches solve general nonconvex optimization problems as a series of local quadratic approximations. Each quadratic optimization step can be solved efficiently, and the overall algorithm adaptively grows and shrinks the size of the region in which the approximation is accurate, hence the name “trust region.” Researchers at ANL developed the theoretical background and prototype code for trust region methods. This proved to be a major breakthrough in the field with the ability to efficiently solve problems that had previously been intractable. Trust region methods and concepts remain in use today.

ANL has also been a pioneer in developing and deploying general-purpose optimization software, beginning with MINPACK, a software package for solving systems of nonlinear equations and least-squares minimization problems, first developed in the late 1970s. Starting in the 1990s, ANL led the development of TAO (Toolkit for Advanced Optimization) software package for solving large-scale optimization problems on distributed-memory HPC systems, an effort funded under ASCR's applied math base program and SciDAC. TAO allowed for much more complex functions to be optimized, enabling simulations to start to be used as design tools, permitting computers to find the best (cheapest,

strongest, fastest, etc.) design for a physical or engineering system. ANL also deployed the Network-Enabled Optimization System (NEOS) in the mid-1990s. NEOS allowed a user to submit a function in the form of a snippet of computer code over the web, and NEOS' servers and algorithms would find values that optimized the function. NEOS was one of the very first examples of software-as-a-service, particularly in the scientific community.

“Outer-loop” capabilities

As computers became faster and simulation methods improved, single simulations no longer required heroic, multi-hour runs on the most advanced computers. Instead, simulations could be viewed as function calls, albeit fairly expensive ones. Starting in the 1990s ASCR-funded researchers at SNL, ANL, University of Texas, and elsewhere began exploring “outer-loop” methods for design optimization and analysis. This work built upon the capabilities created by ASCR's foundational investments in mathematics, computer science, and computational science. The goal was to explore the landscape of simulation outputs to support optimal design, failure margins, uncertainties, and more. As simulations are increasingly used to support critical decision making, it is essential to understand margins and uncertainties. In the early 2000s, ASCR's AMR program initiated the Limits of Predictability which focused on quantitative understanding of the accuracy of simulations. This program resulted in mathematical insights into numerical methods, but it also helped drive the maturation of outer-loop algorithms and libraries. SNL's DAKOTA (Design and Analysis toolKit for Optimization and Terascale Applications) toolkit was initially built to support design optimization and has broadened to be a widely used tool for function space exploration.

When possible, optimization algorithms are greatly enhanced by knowing the slope of the landscape. For simple closed-form functions it is often easy to provide this derivative information, but it is vastly more challenging when the function is a complex simulation. Building upon their long history of optimization research ANL has been a pioneer in the field of algorithmic differentiation. The goal of this work is to take a computer code that computes a function and produce a companion piece of code that computes the function's derivatives. For complex, high-dimensional functions, it is critical to find ways to compute derivatives efficiently. Starting in the late 1980s, researchers at ANL and Rice University developed a series of software tools for algorithmic differentiation. These initially targeted Fortran codes, but have been extended to C and C++. This work has generated many algorithmic challenges that have contributed to DOE's growing interest in discrete optimization. Starting in the 2000s, SNL developed an alternative vision for algorithmic differentiation software in which derivatives are propagated automatically across layers in the simulation software stack. This idea is central to the power of Trilinos and is supported by DAKOTA.

Discrete optimization

Discrete optimization emerged in the 1990s and 2000s as an important topic for supporting decisions in resource management, cyber security, and other applications. In discrete optimization problems, some or all of the values in a function are limited to discrete possibilities, not real numbers. For example, a switch is on or off, the number of facilities in a complex is an integer, a pipe in a network is included or it is not. In these problems, derivatives are not available, so different optimization

methods have to be employed. ASCR-funded researchers at SNL and ANL developed different methods to explore a space of inputs using heuristics like genetic algorithms, or comprehensive search methods like branch-and-bound. SNL developed parallel versions of these methods to make use of advanced computers for the most challenging problems. SNL also developed Pyomo, a new, open-source modeling language to facilitate optimization.

In the 1990s, SNL began working on the closely related topic of graph algorithms. The initial driver behind this work was the need for algorithms and tools to facilitate the effective use of parallel computers. An important early result was the idea of multi-level graph partitioning and the Chaco tool for dividing work amongst the processors of a parallel machine. In the 2000s, ANL also took an interest in graph algorithms as a means for devising efficient differentiation algorithms. ASCR-funded researchers have also applied graph algorithms to problems in mesh generation, statistical physics, cyber security, critical infrastructure vulnerability, and more.

Since its earliest inception at von Neumann's urging, DOE's applied mathematics program has been the cornerstone of the research activities of ASCR and its predecessor organizations. The emphasis of the program has evolved and broadened through the decades, but it has consistently generated the key foundational and algorithmic advances that have underpinned computational science for DOE and the international scientific community.

2.3 COMPUTER SCIENCE: UNLEASHING SCALE AND SPEED

In 1980, scientists delighted at the 160-megaflop/s (floating point operations per second) speed of the world's fastest computer, the Cray-1, and dial-up lines that allowed connections at 9.6 kilobit/s. Today, supercomputers run at 100+ petaflop/s and are accessible over optical fibers at 100+ gigabit/s. Thus, in just 40 years, we have seen an astonishing nine orders of magnitude increase in computing speed and seven orders of magnitude increase in network speed. That these remarkable developments have yielded transformative advances in both the practice and pace of scientific discovery is due to work supported by the U.S. Department of Energy's (DOE) Office of Advanced Scientific Research (ASCR) and its predecessor organizations, under its programs in computer science, collaborations, and networking, and via initiatives such as HPCC (High Performance Computing and Communications), DOE2000, SciDAC (Scientific Discovery through Advanced Computing), and ECP (Exascale Computing Project).

ASCR's success over four decades owes much to the highly productive interplay that ASCR has long promoted among scientific applications, computing facilities, applied mathematics, and computer science. This section focuses on the computer science; the new methods and tools ASCR-funded computer scientists have developed to enable new approaches to solving scientific problems. These pioneers have adapted for science, and often shaped the development of, a series of revolutionary new technologies: massively parallel processing, ultra-high-speed optical networks, the Internet, grid and cloud computing, distributed collaboration tools, and other technologies. This work has enabled the creation of immensely powerful computing systems, the effective operation of global-scale scientific collaborations, new discoveries across many scientific disciplines, and the production of new technologies with pervasive impacts on industry and society.

Massively Parallel Processing Computing

The work of ASCR, and its predecessor units within DOE, in parallel computing dates back to the early 1980s when visionary scientists at the DOE laboratories recognized the dominant vector supercomputers at the time would eventually have to be replaced by a new breed of massively parallel processing (MPP) computers.

State of the art high performance computing (HPC) in the early 1980s involved specialized vector supercomputers. DOE scientists and others used these systems to create and run cutting-edge simulations in many scientific domains.

While the Cray-2 supercomputer delivered a then-remarkable two gigaflop/s, by 1985 DOE scientists recognized that much faster computers would be needed to solve next-generation problems in climate science, computational chemistry, and other fields. They argued that the key to achieving such speeds was to switch to a radically different technology based on interconnecting inexpensive microprocessors of the type being manufactured for use in workstations and personal computers. To believers in the new MPP technology its ultimate success was inevitable; to its detractors the proposed systems were unreliable, unproven, and difficult to program.

ASCR-supported visionaries were ultimately proven right as MPP systems became reliable, usable, and mainstream. By 2020, MPP systems are delivering computational performance nine orders of magnitude greater than that of the Cray-1. But for that to happen many difficult technical obstacles had to be overcome.

MPP Programming: Harnessing millions of processors

The first challenge was how to program these new MPP systems. Programming a conventional computer involves specifying “just one thing after another” (Arnold Toynbee). Programming languages like Fortran, C, Matlab, and Python have evolved to meet this need to specify sequences of actions concisely and correctly. But parallel computing, in which several processors run at the same time, involves many things happening at once and easily leads to confusion or stalemate if not carefully managed. MPP systems demanded new programming methods and tools that programmers could use to specify what different processors should do and how those different activities should be coordinated.

When the first MPP systems emerged in the mid-1980s, there was no shortage of often wild and wonderful ideas within the academic community about how these systems should be programmed. As in other areas of computer science, ASCR’s emphasis on enabling science led ASCR-supported researchers to identify and focus on ideas that worked. That focus, realized in programs such as HPCC and SciDAC, ultimately delivering the programmability, performance, and portability required for effective scientific use of MPP systems.

The first area in which ASCR-supported research led to progress in MPP programmability was in message-passing. In the early days of MPP systems, there was uncertainty as to whether processors should communicate by reading and writing shared memory locations (supposedly easier for programmers, but harder to implement efficiently and not scalable to the large systems needed to solve today’s problems) or by exchanging messages. Furthermore, each MPP system vendor (of which there were many) had its own libraries for writing parallel programs, each with unique features. ASCR-supported researchers overcame these barriers to portable and performant parallel programs by developing two widely used message-passing libraries, P4 and PVM. In 1993, ASCR-funded computer scientists convened the Message Passing Interface (MPI) Forum, which started from scratch to build an MPI library they hoped would become a standard across MPP systems. The result was MPICH and the group released a new version of MPICH every six weeks. Users running it on various systems provided feedback which was then used to improve the next release. Vendor representatives were among the 30 people who participated in the meetings from the start, realizing that having a standard like MPI could help expand the market for their computers; MPICH gave them a running start. The first MPI standard was published in May 1994. The MPICH team began a comprehensive rewrite of the implementation in 2000 and the result received an R&D100 Award in 2005.

Despite the passage of three decades and enormous changes in MPP system architecture, parallel programs of today are created with message-passing constructs largely identical to those developed in the 1990s. This is not to say that MPI implementations have not evolved greatly as MPP systems have scaled from tens to millions of cores, they have, often thanks to ASCR support. Indeed, ASCR research

on MPI implementations continues today within ECP. For example, addressing how message passing can co-exist with shared-memory programming models used in increasingly powerful processing nodes as defined by the OpenMP specification. But the underlying programming model has stayed the same.

While message passing has been, and continues to be, a tremendous success, it is a low-level programming model requiring programmers to spend time worrying about where data are located within a parallel computer. Both academic and commercial communities have continued to seek a programming model that could hide the complexity of parallelism while providing high performance across different machine architectures and scales. This ideal model would lower the barrier to entry for parallel computing, make parallel machines more attractive as commercial products, and help create a strong workforce for high-end scientific programming.

ASCR-supported research has long been at the forefront of this search for alternative parallel programming models. For example, in the 1980s, the SISAL (Streams and Iteration in a Single Assignment Language) project developed a compiler for a specialized high-level language that achieved performance comparable to Fortran compilers. However, SISAL ultimately had little impact on scientific programming due to its lack of compatibility with the programming languages used by

Global Arrays and Computational Chemistry

The Global Arrays library is an example of influential ASCR-funded innovations in parallel programming models. It implements a shared-memory programming model in which the programmer manages locality explicitly via calling functions that transfer data between a global address space (a distributed array) and local storage. This library was designed from the beginning to provide features necessary for convenient implementation of computational chemistry codes. It has become a de-facto programming model for mainstream scalable chemistry electronic structure codes. Some of the most widely used electronic structure codes make use of Global Arrays library: NWChem, Columbus, MOLPRO, MOLCAS, QChem and GAMESS-UK.

Additionally, the Global Arrays library has been used in many other applications that rely on large shared arrays such as the AMR-based Computational Physics applications NWPhys/NWGrid, the parallel version of the Visual Analytics application IN-SPIRE, and the bioinformatics application ScalaBLAST, a scalable implementation of the mostly widely used DNA sequence searching and matching program, BLAST. Other application areas that use the Global Arrays library include smooth particle hydrodynamics, image processing and analysis and groundwater modeling.

working scientists. Importantly, this experience revealed that any new programming language must function within a complex ecosystem involving commonly available hardware and existing software and programming tools; it must also garner sufficient algorithms and applications to draw a substantial user community.

Other ASCR work in parallel programming models has profited from this lesson. For example, ASCR's Center for Programming Models for Scalable Parallel Computing supported work on Global Arrays, a shared-memory programming model for highly distributed hardware that has no hardware-shared memory. Global Arrays is implemented as a library rather than a language, simplifying adoption by scientists; it has been widely adopted in computational chemistry. The same center conducted pioneering work on partitioned global

address space (PGAS) languages, which integrate a shared-memory programming model into conventional sequential programming languages: UPC, Coarray Fortran, and Titanium are extensions of C, Fortran-90, and Java, respectively.

ASCR support has also contributed to parallel programming methods in which parallelism is not embedded within scientific libraries and applications but defined by a separate coordination layer—the terms scripting and workflow are also used. This coordination layer may specify, for example, that 100 or 1,000,000 copies of a sequential program should be run, the results combined and analyzed, and then further applications called, depending on outcomes. As parallel applications become more complex, such application coordination systems are becoming increasingly important.

ASCR-funded work in this area dates back to the 1980s when Schedule (1987), Strand (1989), and Program Composition Notation (1992) were used by parallel programmers within and outside DOE laboratories. Later systems such as Swift, Kepler, Pegasus, and Parsl explored alternative approaches to describing coordination patterns, from scripting languages to graphical programming. The latter systems all continue to be used to develop parallel applications. Swift, for example, is used to implement hyperparameter search methods in the CANDLE (CANcer Distributed Learning Environment) deep learning project.

Parallel solver libraries: Fast paths to MPP

Parallel solver libraries represent another important thread in parallel programming models. Numerical solvers simplified parallel programming by enabling the use of SPMD (single program, multiple data) programming models (see section 2.2 for a mathematical perspective on numerical solvers). In SPMD the programmer organizes a series of calls to parallel library functions that create and manipulate parallel data structures that are distributed over many processors. This approach, as realized in PETSc (Portable Extensible Toolkit for Scientific Computation) from 1991 onward, as well as in other libraries such as TAO (Toolkit for Advanced Optimization) and Trilinos, has transformed how many scientists develop parallel applications. A PETSc user, for example, may call one PETSc function to create a multi-dimensional matrix and then another function to apply a solver to that matrix. Under the covers, MPI calls exchange information among processors, but this activity is invisible to the programmer. The programmer can use other functions to control how parallel matrices are distributed over processors, with distributions often effecting only performance, not the result computed.

Scalable input/output: Talking fast to data

Most scientific applications need to read and write data located on disk storage. As MPP systems accelerated from the 1980s onwards so too did the volumes of data to be manipulated and thus the speeds at which data needed to be read and written. As disk speeds increased more slowly than supercomputer performance, storage systems had to incorporate increasingly large amounts of parallelism so the data could be read from, and written to, many disks simultaneously—initially a handful, and ultimately many thousands. These developments posed major system design, performance, and usability challenges that were not being addressed by commercial file systems.

In the early 1990s, the computer science community responded by developing new methods and tools that improved both the usability and performance of HPC (high performance computing) I/O (input/output) systems. HPC applications are able to rapidly access these parallel file systems because

they are a coherent whole (disks, servers, and network links). Without parallel file systems, high-performance I/O for computational science would simply not be possible. I/O middleware implements critical optimizations that leverage knowledge about groups of processes accessing storage in a coordinated manner. And high-level I/O libraries provide convenient, portable abstractions for storing and accessing scientific data.

Three specific breakthroughs, enabled through ASCR funding, are in the areas of object-based parallel file systems, parallel middleware, and high-level I/O libraries.

The first of these breakthroughs concerns how parallel file systems organize data across many hardware devices. In a conventional block-based file system, every I/O node's kernel participates in a shared disk data structure comprising potentially trillions of sectors, introducing both code complexity and bottlenecks. In an object store, parallel files are constructed by striping across a relatively small number of variable-length byte streams, each realized independently as one object server. ASCR played a major role in the development and adoption of object storage principles in parallel file systems, a critical enabler for extreme-scale I/O, via its support for the pioneering PVFS2 file system from 2000 onwards, and for the realization of the Lustre parallel file system at extreme scales, delivering I/O performance of more than a terabyte per second.

The second and third breakthroughs concern how parallel programs access parallel file systems. This problem is mind boggling in its complexity. Data must be transferred between tens or even hundreds of thousands of computing nodes on the one hand, and storage systems comprising hundreds of storage servers on the other, all while achieving high performance for workloads that may vary widely in their data access patterns.

A first set of ASCR innovations in this area produced I/O middleware, ultimately incorporated within programming models such as MPI, that leverage knowledge about groups of processes that are accessing storage in a coordinated manner to implement optimizations such as gathering all data intended for a single I/O node to a single location before writing, rather than performing, hundreds or thousands of small independent write operations. I/O middleware delivered performance, but as they required application scientists to express I/O operations in terms of bytes, they were not intuitive to use.

Another set of ASCR innovations addressed this latter challenge by producing increasingly powerful high-level parallel I/O libraries that allow scientists to interact with file systems in terms they understand, such as multi-dimensional arrays, with variables, dimensions, and attributes. Sustained innovations realized in libraries such as parallel NetCDF, HDF5, and ADIOS (Adaptable I/O System) delivered portability of data formats, a natural interface for interacting with scientific data, and high performance for parallel applications. Other important work has addressed related challenges such as maximizing performance for highly concurrent updates, tagging data with provenance, moving data between systems, organizing data from many processes in a single file, and exploiting emerging, extremely fast non-volatile memory express (NVMe) solid-state storage devices.

Large-scale data analysis and visualization

The enormous data produced by massive simulations encounters another bottleneck when it comes to human perceptual and cognitive capabilities—extracting insights from giga-, tera-, and peta-bytes of data. This concern motivated ASCR-funded work on large-scale data analysis, visualization, and specialized display devices.

Data analysis is first and foremost a mathematical and algorithmic problem, in that it involves the development of methods for extracting significant information (means, extremes, trends) from data. However, as data volumes grow, the need arises to harness many processors and parallel I/O to enable timely analysis. Thus, large-scale data analysis also becomes an MPP problem. In response, ASCR-funded researchers developed open source scientific visualization software such as ParaView, developed at Los Alamos National Laboratory (LANL) and Sandia National Laboratories (SNL), and Kitware, and VisIt, developed at Lawrence Livermore National Laboratory (LLNL), Lawrence Berkeley National Laboratory (LBNL), Oak Ridge National Laboratory (ORNL), and the University of California at Davis, both of which have commoditized access to visualization capabilities in many fields of science. VisIt is a distributed, parallel visualization and graphical analysis tool for data defined on two- and three-dimensional meshes that can visualize data represented in more than 120 scientific data formats, while ParaView targets extremely large datasets using distributed memory computing resources. Both systems began development around 2000 with funding from the ASCI (Accelerated Strategic Computing Initiative) program and were later expanded with ASCR support to enable production-quality, petascale-capable visual data exploration and analysis on ASCR HPC platforms.

The use of specialized display technologies to facilitate understanding of large datasets is another area where ASCR researchers have played a leadership role, exploring from 1994 onward the use, for scientific purposes, of technologies such as the CAVE immersive virtual reality environment. Argonne National Laboratory (ANL) developed the first interactive supercomputer-in-the-loop CAVE application in 1994 and worked with academic research collaborators at the University of Illinois at Chicago for more than two decades to prototype, test, and build advanced visualization display and collaboration technologies that have gone on to be widely deployed for scientific use. The CAVE comm library for connecting virtual environments and supercomputers supported a wide range of applications in which scientists and engineers interacted with data produced by simulations. In one influential example, virtual reality was used to aid in the design of commercial boilers and incinerators.

Indexing, compression, and online analysis

Even massively parallel computing and storage systems encounter limits when processing large quantities of data. ASCR-supported research aided in the development of methods to reduce the amount of data that must be moved between processors and storage: in particular, indexing technologies for finding data in large data volumes, compression methods for reducing the overall footprint of scientific data, and online computing methods for avoiding I/O altogether by analyzing data when it is generated.

One example of ASCR-supported advanced indexing technologies is the FastBit package from LBNL, which enables rapid identification of data of interest in large scientific datasets, facilitating interactive data analysis on datasets of many terabytes. FastBit, which won a 2008 R&D100 Award, has been used to accelerate tasks such as analyses of trillion-particle datasets in laser wakefield particle accelerator studies and in commercial applications ranging from drug discovery to web applications. The related FastQuery system provides similar capabilities for data stored in structured formats such as HDF5. Indexing has also been applied to accelerate file systems. Projects such as Giga+ and IndexFS enabled much needed scalability of metadata and small files on parallel file systems, leveraging indexing techniques to partition metadata efficiently across many servers, to aggregate updates for higher insertion rates, in some cases up to two orders of magnitude faster than without these tools. DeltaFS applied similar methods to store data in ways that accelerate subsequent analysis, delivering a 5000 times speedup for a plasma-kinetic simulation code.

Compression methods seek to encode data in forms that require less storage space in order to reduce storage demands and accelerate I/O. Achieving useful levels of compression typically requires lossy methods that do not preserve all information in the original data. ASCR-supported lossy compression technologies developed at ANL, LLNL, ORNL, and elsewhere include the SZ, ZFP, and MGARD lossy compressors. These have proven capable of reducing data volumes by orders of magnitude in some settings, and the Z-Checker system for evaluating compression quality.

As computers accelerate to exascale performance, it can become impractical to write all interesting data produced by a computation out to disk for subsequent analysis or compression. In such cases, the need arises for what are variously termed “in-situ” or “online” data analysis methods, in which data is analyzed and/or compressed as it is produced. Influential systems resulting from ASCR support via SciDAC and other programs include ANL’s GLEAN in-situ analysis and data movement acceleration framework, ORNL’s ADIOS system for coupling application and analysis codes, and LBNL’s SENSEI (Sub-Electron-Noise Skipper Experimental Instrument) software for scalable and portable in-situ code coupling.

Collaborating at the speed of light

As one group of ASCR-supported researchers innovated in supercomputing others were pursuing equally dramatic transformations in wide area networking. By the early 1980s, the Internet was in use at most DOE laboratories and many universities. However, its 10 kilobit/s dial-up links were good for little more than exchanging email. Things changed dramatically in the 1990s as optical networks were deployed, providing 155 Mbps (megabits per second) connections to many laboratories. By 2018, ESnet (Energy Sciences Network), established in 1986, operated at 100 Gbps (gigabits per second) and had test networks running at 400 Gbps, a stunning eight orders of magnitude performance improvement in 35 years.

Networking technologies produced by the telecommunications industry provided the physical substrate for these developments. But, as with the MPP revolution, realizing the promise of these technologies for science required sustained innovation at many levels: new network protocols, new methods and

tools for data movement and collaboration, and new approaches to science based on large-scale data sharing, among many others. ASCR researchers were early proponents of the use of high-bandwidth, low-latency connections for scientific collaboration and contributed many important technological innovations. The DOE2000 program, launched in 1997, was an important source of support for this work.

ASCR innovations in networking and distributed computations were also spurred by the advent of science experiments like those at the LHC (Large Hadron Collider) that involve scientists, computing, and storage resources at dozens of institutions around the world. These experiments must routinely move terabytes of high-value data per day in order to accomplish their scientific goals. A distributed science analysis approach is essential in order to access and manage the resources required to provide the needed scale of computing and data management.

Network protocols that scale

Today's Internet emerged from early experiments on ARPANET (Advanced Research Projects Agency Network) in the late 1970s, which established the basic protocols that allow computers to exchange information on a global scale. In particular, the TCP (Transmission Control Protocol) ensures the reliable, in-order delivery of information between two communicating end systems. Since its introduction in 1974, along with the IP (Internet Protocol), TCP has become the ubiquitous transport protocol of the Internet, supporting myriad networked applications such as the web browser, email, file sharing, video and voice over IP, and file transfer. Yet the Internet's success was not guaranteed. In 1986, growing use led to a series of congestion collapses. Many concluded that the Internet could not scale. ASCR-supported researchers at LBNL diagnosed the problem (TCP's retransmission on failure caused a vicious cycle of congestion) and developed the congestion control algorithms that have allowed the Internet to scale in terms of number of nodes.

Over subsequent decades, ASCR-supported initiatives such as DOE2000 and the National Collaboratories Research Programs contributed to scaling the Internet along the performance dimension, to the extent that scientists today use essentially the same TCP to transfer data at speeds exceeding 100 Gbps. One important set of developments targeted the tuning of TCP software. By the mid-1990s, the emergence of 100 Mbps networking technologies began to expose a major limitation in the ability of the TCP software to efficiently scale its performance to match that of network hardware. Early attempts to address this limitation by manual tuning of TCP software suffered from three shortcomings: the manual tuning generally required a "network wizard" to complete; the manual tuning only created a static, per-connection optimization of the TCP software; and manual tuning often made inefficient use of local resources at the end system.

These shortcomings gave rise to ASCR-funded projects on TCP autotuning, with the goal of delivering improved network performance and end-system resource usage without users or application programmers having to modify their applications. At LANL, ASCR-supported researchers discovered a way to automatically characterize network dynamics at run time and enable the automatic tuning of TCP for high performance. This technique, dubbed DRS (Dynamic Right-Sizing), progressed from a

research prototype to becoming fully integrated into the Linux operating system (version 2.4.30) for the benefit of all applications and their associated end users. The ASCR-funded net100 project at LBNL and ORNL helped verify the effectiveness of these changes to TCP, as well as the efficacy of then-new TCP congestion control algorithms such as CUBIC and HTCP. Variants of these techniques have since appeared in the Microsoft Windows operating system, FreeBSD, and Mac OSX. In short, ASCR-funded projects in TCP autotuning resulted in powerful and self-adapting software that has been integrated into a multitude of operating systems for the greater good of all current and future end users.

Telepresence: Better than being there?

The emergence of high-speed networks also got researchers thinking about new applications that would use the Internet to transmit real-time streaming data: in particular, audio and video streams to support teleconferencing. ASCR researchers made major contributions in this area via the development at LBNL, in the early 1990s, of the vic and vat tools (“the Mbone tools”) for real-time, multi-party video and audio conferencing over the Internet. In 1994, the Rolling Stones used Mbone to broadcast a concert from Dallas, dubbed the “first major cyberspace multicast concert.” These tools leveraged the then-new RTP (Real-time Transport Protocol) and IP Multicast protocols developed within the Internet Engineering Task Force, packaging them into a form suitable for convenient use. Vic and vat were widely used within DOE laboratories and elsewhere and strongly influenced the development of Internet videoconferencing.



Figure X: Access Grid installations worldwide, in May 2009

Vic and vat were desktop clients, functioning much like commercial video conferencing tools of today. By 1999, researchers at ANL, with support of ASCR’s DOE2000 program on scientific collaboratories,

pushed the limits of network-based collaboration further via the development of Access Grid, a system that supported large-format audio- and video-based collaboration between groups of people in different locations (see Figure 1). An Access Grid node, of which there were at one time more than 500 worldwide, featured an ensemble of resources, including multimedia large-format displays, presentation and interactive environments, and interfaces with grid computing middleware and visualization tools. Access Grid also helped advance network engineering and collaboration technologies.

Deciding who to trust in a world of strangers

While bad actors were rare in the early Internet, trust emerged as a vital concern as the Internet grew and high-speed networks enabled new modes of working based on on-demand access to remote resources (scientific instruments, data, computers), federated use of resources at multiple sites, and new forms of interpersonal collaboration. By the mid-1990s, new technologies were clearly needed in order for these resource federation and collaboration methods to function securely, reliably, and conveniently. In particular, authentication had emerged as a major problem. Users at that time had to acquire and use a separate credential for every site they wished to access, a process that was often insecure and error prone, and also impractically complex when using multiple resources at once.

ASCR-funded computer scientists were leaders in the development of the concepts and technologies that addressed these difficulties, developing early single sign on (SSO) technologies and also pioneering concepts and methods to enable users to define, configure, and operate within distributed virtual organizations. As with other network technologies, progress proceeded through a combination of community standards development and high-quality software development. Important products developed at ANL and LBNL, with the support of ASCR programs such as DOE2000, included the pioneering Akenti authorization service, the GSI (Grid Security Infrastructure) for single sign on, and the ESnet DOE Grids Certificate Service for proving membership in a community of mutually trusted entities. GSI, in particular, made it possible for scientists and engineers to authenticate once and then access computers, storage systems, and other resources worldwide, all using a single credential. They could also delegate authority to programs to perform operations on their behalf. Resource owners could express and enforce policies concerning which credentials are accepted and which sets of users are allowed access. These capabilities were fundamental to the growth of distributed collaborative science, supporting for example breakthrough science within projects such as the LHC and LIGO (Laser Interferometer Gravitational-wave Observatory). They also influenced subsequent development, and adoption across DOE, of next-generation trust technologies in use today, such as Globus Auth.

Moving big data

Modern science and engineering must increasingly deal with enormous amounts of data. Petabytes (a billion megabytes) are commonplace and exabytes (a billion gigabytes) are on the horizon. As data is produced, shared, and analyzed, it must frequently be moved from place to place. Thus, the ability to move data rapidly across networks has become fundamental to many collaborative projects in both research and industry.

Excitement about the increasingly widespread deployment of high-speed networks from the early 1990s onward was tempered by the realization that the transfer speeds actually achieved by applications were embarrassingly low (often 1% or less of available network bandwidth). The effects on users could be devastating. For example, moving one terabyte over a one Gbps network could take many days rather than a few hours, essentially preventing the use of remote systems and distributed collaboration.

The rapid, reliable, and secure movement of data from one file system to another is a challenging end-to-end problem. A single transfer may involve many disks, switches, firewalls, routers, and networks, as well as file systems, network protocols, and authentication protocols. High performance may often require the exploitation of parallelism at the storage system and network level. Application-level software may want to move a single large file, stream data, or transfer many small files. Individual hardware and software components may vary greatly in their performance and reliability. Failures may occur at the hardware and software levels.

ASCR-funded researchers have developed the technologies required to overcome these challenges. For example, pioneering work at LBNL in the early 1990s produced the Distributed Parallel Storage System, which demonstrated how multiple networked data servers could be used to fill a high-speed network. Subsequent work over two decades within ASCR's National Collaboratories Research Program and related efforts has produced technical standards and software that are used internationally by tens of thousands to move exabytes in support of a wide range of science projects. As in MPP, this work engages both extensive use of parallelism to overcome performance limitations of individual devices, and abstraction and automation to address usability and combined research, software development, and standards development.

Much of this work builds on GridFTP, a set of extensions to the popular file transfer protocol (FTP) to enable secure, reliable, and high-performance transfers. These extensions range from messages for negotiating how data are striped over multiple data movers and network connections, to restart markers that permit an interrupted transfer to be resumed. ASCR funding to ANL scientists led to the codification of GridFTP extensions within the Open Grid Forum, an international standards organization. They also produced high-performance implementations of the protocol that use specialized methods to achieve extremely high performance in a variety of settings. For example, the Globus GridFTP implementation uses striping across multiple servers to permit files in a parallel file system to be moved across wide area networks at speeds exceeding 100 Gbps.

Experimental and simulation scientists worldwide use GridFTP; LHC experiments use it to transfer close to a petabyte per day. DOE supercomputer centers use Globus to move data in and out of their mass storage systems, achieving performance 20 or more times what was achieved previously; one single transfer in 2019 totaled nine petabytes. DOE light sources use Globus to allow U.S. and international users to access experimental data, again at vastly greater speeds than before.

Dynamic network circuit reservation: Dedicated lanes for science

Another area of network protocol innovation concerns mechanisms for providing service guarantees, as needed, for example, when moving data from instruments that operate in real time to remote computing and storage systems. Advanced applications increasingly require massive data movement at a given rate and in a specific time period. Such service guarantees are not possible with IP alone due to its best-effort nature. Further, network operators increasingly need the ability to manage network traffic on a per-flow basis (that is, at the level of individual application-to-application communications), when such site-to-site flows frequently account for a significant fraction of the total network traffic.

ASCR research in this space dates back to 1998 when scientists at LBNL and ANL successfully demonstrated differentiated services by sending two video streams over the Internet between the two laboratories. In a test, sharing an intentionally congested path, the priority-marked stream moved at eight frames per second, while the standard version transmitted just one frame per second. This result helped pave the way for more reliable and constant connectivity via priority bandwidth on the Internet.

Subsequently, ESnet developed OSCARS (On-demand Secure Circuits and Advance Reservation System) to provide guaranteed bandwidth based on advanced reservations, user-selectable backup path strategies for high reliability, data connections as either IP or layer 2 (Ethernet) paths, and a web services user interface.

ESnet developed this service as part of a successful international collaboration of research and education networks working to ensure the virtual circuit service is available from end-to-end, since most connections have one end at a DOE laboratory and the other end at a U.S. or European university or research institute. Today virtual circuits can be set up across the many different networks connecting collaborating science institutions. For example, OSCARS circuits move all data collected at the LHC (the world's largest scientific experiment) in Switzerland to Fermi National Accelerator Laboratory (FNL) and Brookhaven National Laboratory, two Office of Science labs where those data are cached for U.S. physicists to analyze.

Decomposing the machine room: Grid and cloud

The emergence of high-speed optical networks in the early 1990s led researchers to imagine new models of computing that leveraged those networks to connect scientists, computers, data stores, and instruments regardless of location. Computers, data, or software could be accessed on demand from remote systems, for example to analyze data generated by scientific instruments in real time or to perform computations too large for local computers. This vision of a new computing fabric, which became known as the “grid” by analogy with the on-demand access to electricity provided by the electrical grid, was realized through ASCR's National Collaboratories Research Program and SciDAC program.

DOE Science Grid, ANL, LBNL, ORNL, PNNL

An influential event in this development was the I-WAY experiment in 1995, in which 60 groups across the U.S. collaborated to connect and apply a grid network linking 17 sites. This grew into an ANL-led effort to develop the Globus Toolkit. ASCR's DOE2000 and SciDAC initiatives supported projects such as the DOE Science Grid at ANL, LBNL, ORNL, and Pacific Northwest National Laboratory (PNNL) then developed and deployed authentication and authorization, data transfer, remote computation, and resource discovery tools to permit coordinated, secure, and reliable use of distributed grid resources. In 1999, the ASCR-funded the multi-lab Particle Physics Data Grid project, which along with other efforts in Europe and the U.S., was built on this foundation to establish what became the Open Science Grid (OSG) in the U.S. and the Worldwide Large Hadron Collider Computing Grid (WLCG). OSG provided much of the data processing capabilities for major experiments at FNL. OSG and LCG (linear congruential generator) ultimately connected more than 300,000 computers worldwide to enable the data processing used to discover the Higgs Boson, leading CERN Director General Ralf Heuer to declare in 2012 that the Higgs discovery was only possible because of the "extraordinary achievements of grid computing."

ASCR support enabled other groundbreaking efforts, such as the NEOS (Network Enabled Optimization System) service created at ANL in 1996 pioneered access to optimization software. Now hosted at the University of Wisconsin, this pioneering scientific software-as-a-service has enabled tens of thousands of people to access advanced optimization capabilities.

Another pioneering ASCR-funded project was the ESG (Earth System Grid), established in 1999 by ANL and LLNL in collaboration with NCAR (National Center for Atmospheric Research) and the University of Southern California to enable sharing of large climate model datasets. ESG and the successor global ESG Federation, have enabled worldwide distribution of the terabytes (initially) and petabytes (today) of climate simulation data used in assessments by the Intergovernmental Panel on Climate Change.

ASCR-funded researchers were at the forefront of new models of distributed computing. Consequently, they were well positioned to exploit commercial cloud computing (grid computing with a business model) when this emerged in the 2000s. One influential development, established by the ASCR National Collaboratories Research Program, was the Globus Online service (now called Globus), launched in 2010, Globus leverages software-as-a-service concepts to provide identity management, data transfer and sharing, and research automation services for scientists at thousands of research institutions. Globus has become central to scientific practice in many science disciplines and, when used in conjunction with Science DMZ (Demilitarized Zone), makes data movement at multiple gigabyte/s routine.

The Science DMZ: Bringing big data to the campus

Another reason for poor end-to-end network performance is often the dreaded "last mile" (the path from wide area network to computers and storage systems within a laboratory or university campus).

This network often features devices poorly configured for, or even incapable of, receiving high-speed data streams from wide area network devices. In addition, site security at universities and laboratories is typically handled by a relatively low performance firewall through which all traffic must pass to get to computing systems on site. Such firewalls became a severe impediment for high-speed, long-duration data streams.

In response, ASCR-supported researchers at ESnet conceptualized and codified a new network architecture that came to be known as the Science DMZ. This architecture locates a wide area network-capable network device and a small number of high-performance data transfer nodes (DTN) within a special campus network domain outside the site firewall. The DTNs typically have a connection to the campus local area network that bypasses the site firewall, but with several strong security features, including the constraint that data transfers in either direction have to be initiated from within the site. Integration with the Globus service, with control channels passing through the site firewall, then allows high-speed transfers to be performed end-to-end without firewall slowdowns.

First demonstrated and evaluated at DOE laboratories, with support from the National Science Foundation, the Science DMZ architecture has since been deployed at more than one hundred university campuses. Science DMZ has been proven useful in a variety of research environments, including supercomputer centers, experiments, campus HPC centers, and laboratories.

Protecting networked computers from bad actors

Scientific research is founded on the free exchange of information and ideas among collaborators world-wide, and the Internet has made such collaborations far easier and more effective than ever before. Today's scientists depend critically on full and open access to the Internet. Such unfettered access also opens the doors to incessant network attacks, and scientific institutes such as DOE national laboratories find themselves under great pressure to secure their network access without choking off its richness.

Many such environments are both safer and more open today because they run the Bro network intrusion detection system (developed at LBNL and the International Computer Science Institute). Bro is an ASCR-funded open-source security monitor that can be used to inspect, in depth, all traffic on a network link for signs of suspicious activity. Unlike conventional network intrusion detection systems that support simple character matching on network traffic, Bro provides a general framework for analyzing network activity, supporting richly customizable higher-level detection based on semantics extracted from network traffic, then automatically blocking suspect traffic and alerting network security staff. Work on Bro dates back to 1995. Today, the Bro system enables (via parallelism, using many analysis servers) real-time network traffic analysis at more than 100 gigabit/s.

Laying the foundations for scientific AI

A report on ASCR computing would not be complete without a discussion of how ASCR research relates to recent advances in AI (artificial intelligence) that have already delivered dramatic

improvements in areas such as image recognition and machine translation, and now seem poised to have major impacts in scientific research.

ASCR work in AI dates back to the 1960s, when a team, led by pioneering logician Larry Wos, at ANL started development of automated reasoning systems, such as the AURA (Automated Reasoning Assistant) and Otter automated theorem prover. Otter was subsequently used to prove open problems in mathematics, including the Robbins conjecture in 1997 and the XCB problem in equivalential calculus in 2002. Methods pioneered in these systems also contributed to the subsequent development of satisfiability modulo theory (SMT) solvers such as Microsoft's Z3, and seem likely to influence the integration of symbolic and connectionist approaches to AI that will be required for AI-enabled science.

The more recent deep learning revolution also depends increasingly on technologies that have resulted from ASCR research. As neural networks and the associated computations become larger, we see increased use of technologies described earlier in this chapter, including MPI, parallel I/O, and high-performance linear algebra. Automatic differentiation is another ASCR-supported technology that is central to modern deep learning. The training of multi-layer neural networks involves propagating updates to network weights backwards through a network, which in turn requires efficient codes for computing derivatives of functions executed when running the network forward. Automatic differentiation methods are widely used for that purpose.

The past 40 years of ASCR-supported research in computer science at DOE laboratories has not only enabled a remarkable flowering of high-performance computational and distributed science, but also contributed to the establishment of the new AI methods that seem likely to be transformative in the next 40 years.

2.4 COMPUTER ARCHITECTURE

The Department of Energy's (DOE) Advanced Scientific Computing Research (ASCR) program is widely known for the leadership class computers housed at the Office of Science laboratories. These machines are the most visible aspect of ASCR's long history of impacts on computer architecture. These impacts range from fundamental research concepts to transformational paradigm shifts. ASCR's influence has come through a variety of modalities including basic research grants, hands-on exploration of emerging hardware, codesign collaborations with expertise from hardware architects to application scientists, and large procurements that have shaped the commercial marketplace.

Context

In the four decades since the DOE was created, the microelectronics revolution has reshaped the world. In 1965 Gordon Moore made the techno-economic observation that transistors would shrink at an exponential rate, and his projections have held true for five decades. In 1974 Robert Dennard added a critical insight that the power required to drive a transistor is proportional to its size. Moore's Law and Dennard Scaling combined to allow for exponential growth in transistor density with no overall increase in power consumption (at least until Dennard scaling reached its limit a bit over a decade ago). These technology advances have led to exponential growth in computing capability that, in turn, yielded a rapid, predictable, and affordable improvement of computing performance for mobile devices, desktop machines, data centers, and high-performance computers for scientific discovery and national security.

ASCR has been a major beneficiary of these advances in technology and the faster computers they enabled. But ASCR has also helped drive and shape the computing architectures built from these microelectronic advances. Importantly, the capability of the world's fastest computers (many procured by ASCR facilities) has grown even faster than Moore's Law and Dennard Scaling would have predicted. Innovations driven by ASCR-funded research, development, and procurement efforts have helped define the landscape of HPC (high performance computing). Moore's Law is coming to its end as transistors are approaching the atomic scale. Without an unforeseen new transistor technology the rate of improvement in microelectronics will slow drastically and computing advances will have to come predominantly from architectural innovation. ASCR-funded researchers will need to contribute to these efforts to ensure the capabilities demanded by scientific applications are addressed.

Paradigm shifts

One likely response to the end of Moore's Law will be architectural specialization, designing more specialized hardware to accelerate specific computations. We are already in the early stages of this transition to heterogeneous computing in which different processing elements work on different elements of an overall computation. ASCR facilities have embraced this trend with machines that combine traditional central processing units (CPU) with graphics processing units (GPU). GPUs can execute some important scientific computing kernels more quickly while using less power than CPUs.

In 2012, the Titan computer at the Oak Ridge Leadership Computing Facility (OLCF) became the first ASCR flagship platform with GPUs. Titan's successor, Summit, also employs GPUs and is currently the fastest computer in the world.

Effective use of Titan and Summit has required considerable changes to existing software. New algorithms have been developed to get optimal performance from GPUs, and simulation codes have been refactored into GPU-targeted and CPU-targeted kernels. This refactoring has often been done by hand, but this manual process is tedious and not scalable in a future where computers have several specialized accelerators and the stable of ASCR computers have different kinds of accelerators. New approaches to building software will be needed in which the functional specification is separated from the details of the computer architecture. ASCR has contributed to the development of programming models that incorporate this approach, including open standards like OpenMP and OpenACC, and metaprogramming solutions such as Kokkos from Sandia National Laboratories (SNL) and Raja/Umpire from Lawrence Livermore National Laboratory (LLNL). New approaches to system software will also be needed to dynamically manage the allocation of heterogeneous resources. ASCR is funding research at SNL, Oak Ridge (ORNL), and Argonne (ANL) national laboratories to address these challenges.

Although the end of Moore's Law will lead to disruptive changes in computer architecture, this will not be the first time that ASCR and its predecessor organizations have helped steer the community through a major paradigm change in computer architecture. In the 1960s, 1970s, and 1980s, the fastest computers were machines made by Control Data Corporation (CDC) and Cray. These machines used vector processing, an approach similar to GPUs today, to allow multiple operations to be performed in a single clock cycle. By the late 1980s this vector approach was under threat by the rapid advance in the performance of simpler computers due to the progress of Moore's Law and Dennard Scaling. Systems built with simpler commodity microprocessors were more cost effective than specialized machines since their development costs were underpinned by a larger commercial market. In principle, the disadvantages of slower cores could be ameliorated by having multiple cores work together to solve a problem. But this promising concept of parallel computing needed to be validated for ASCR-supported applications through redesign and testing.

In the 1970s and 1980s, ASCR's predecessor organization funded several computer architecture projects that would have major impacts on DOE and the broader research community. The Caltech Concurrent Computing Project (C3P) headed by Geoffrey Fox explored parallel computing and demonstrated its potential for many scientific computing applications. This effort built on top of a DARPA (Defense Advanced Research Projects Agency)-funded activity, but greatly deepened and broadened the understanding of how parallelism could be applied to ASCR problems. C3P led to several commercial products and paved the way for distributed memory computing architectures, algorithms, and software.

ASCR also funded the CEDAR project at the University of Illinois at Urbana-Champaign headed by Dave Kuck. Completed in 1988, CEDAR had a strong focus on system software for vector and parallel

computing and led to the Alliant commercial product. The compiler concepts pioneered in CEDAR continue to be important today.

Another important ASCR-supported architecture activity was the NYU (New York University) Ultracomputer project at the Courant Institute of Mathematical Sciences first headed by Jack Schwartz and later by Allan Gottlieb. Although there was no commercial computer spawned directly from the Ultracomputer, the project had substantial influence on the parallel processing community.

An ASCR-funded graduate student at Courant, Joseph (Josh) Fisher, was developing a micro-coded clone of the CDC 6600 computer (considered to be the first successful supercomputer and one of the fastest machines in the world at the time). A key performance technique in this system was increasing concurrency by moving micro-operations speculatively above branches. Fisher's research produced a practical method for automating that process and others, coined the term "instruction level parallelism", and laid the groundwork for what he later called VLIW (Very Long Instruction Word) architectures. His insights are ubiquitous in modern architectures and are used directly in massive-volume embedded CPUs. His software design is a key technology in GPUs. It is satisfying to note that GPUs are the performance engines of current leadership class DOE supercomputers, demonstrating the outsized impact of fundamental advances.

Building upon this early foundational work in parallelism, ASCR then provided funding for the CSCC (Concurrent Supercomputing Consortium). CSCC was a broad effort with multiple funders and a range of participants including SNL, ORNL, ANL, Los Alamos (LANL), and Pacific Northwest (PNNL) national laboratories, as well as Intel, NASA's Jet Propulsion Laboratory, and Rice and Purdue universities. The CSCC further developed ideas pioneered by the C3P and CEDAR projects and resulted in the Intel Touchstone Delta, which was the basis for the Intel's Paragon family of distributed memory supercomputers.

Computer architecture testbeds

Parallelism was emerging in the 1980s as the likely future of high performance computing (HPC). Many uncertainties existed about which architectures and programming models would be most viable and appropriate for ASCR applications. A number of small companies began offering a diverse array of products reflecting distinctive technological paths. To help sort through this complexity, ASCR supported the creation of Argonne's Advanced Computing Research Facility (ACRF) in 1984. The goals of ACRF were:

- To encourage experimentation on computers with innovative designs,
- To assess the suitability of diverse machines for specific applications,
- To assist research in parallel computation,
- To encourage the incorporation of state of the art computational techniques in research, development, and industry,
- To provide leadership in enhancing computing environments, and
- To operate as a national user facility in parallel computing.

ACRF procured a range of machines for experimentation and analysis and helped provide greater clarity about the strengths and weaknesses for ASCR applications of these architectures and approaches. In so doing, ACRF paved the way for what is now a recurring element of ASCR's investment strategy—exploratory architecture testbeds for benchmarking and evaluation.

Benchmarking and evaluation efforts at ACRF, SNL, and other laboratories helped shape the inevitable shakeout and selection of early parallel vendors and approaches. The commercial advantages of commodity microprocessors eventually drove early specialized hardware out of the market—the triumph of the “killer micros.” But ASCR's investments in architectural research and early prototyping had prepared the laboratories for this outcome. DOE has been a leader in message passing parallel computing ever since.

The importance of benchmarking did not end with the transition from vector to parallel computing. Early evaluation of HPC technologies and systems has continued to play a central role in the trajectory of systems procured by DOE facilities, as well as on the resulting software created to use these systems. Nearly all large DOE system procurements have been preceded by hands-on studies of smaller systems and emerging technologies. This is particularly true for first-of-a-kind systems pioneered by the DOE national laboratories like Cray's X1, IBM's BlueGene, Cray's XT, and the NVIDIA GPU-based systems. These evaluations have provided DOE management and the research community with specific, credible evidence on the performance of important DOE applications on these new architectures for procurements, application readiness, and codesign interactions with vendors. For example, pathfinding work at ORNL's Experimental Computing Laboratory on GPUs led directly to the 2010 Gordon Bell Prize and the selection of NVIDIA GPUs in the OLCF Titan system. Other examples include LBNL's (Lawrence Berkeley National Laboratory) investigations of large-scale vector and scalar systems, ORNL's study of the Cray X1, Cray XT, and BlueGene/P systems, the SNL design of the Cray Red Storm machine, and LANL's evaluation of IBM's Roadrunner built with Cell processors.

These detailed application studies have revealed algorithmic limitations to scalability and cases where an application is perfectly suited to architectural features such as vectors, local stores, and interconnection networks. The head-to-head comparison of real applications across architectures, combined with a careful analysis of the applications' characteristics, has provided invaluable insight into the suitability of a given type of machine for classes of scientific methods. The novel aspect of this research in the context of large-scale platforms is the emphasis on full applications, run with real input data and at the scale desired by domain scientists. These research activities are also much more difficult to perform, compared with traditional evaluations of benchmarking suites, because they require coordination, communication, and management of a team of application and computer scientists from highly disparate backgrounds. These interdisciplinary collaborations are the heart of the codesign process that allows architectures, system software, and applications to evolve jointly.

In addition to understanding the architectural impact on computational science applications, it is increasingly important to evaluate the impact of novel architectural designs on machine learning and

data-intensive scenarios, such as analyzing experimental and simulation data. Scientific computing workflows are getting more complex and this has ramifications for future procurements.

With the looming end of Moore's Law, the need for prototype evaluation and codesign is greater than ever before. The pace of architectural innovation is likely to quicken and the increasing need for power efficiency will accelerate recent trends towards architectural diversity through new interest in specialization and tighter system integration. It is essential to understand the tradeoffs of these computing approaches in the context of high-end numerical simulations, machine learning, and data intensive science.

Architectural testbeds have been and will continue to be essential to deepen understanding of the potential of emerging technologies, driving collaboration and codesign between laboratories and vendors, informing large procurement decisions, and adapting applications for new machines. ASCR's investments in these activities have been recognized by the HPC community via multiple best-paper awards at leading conferences.

Government and Industry Partnerships

Recent years have seen a staggering increase in the computational resources of the Office of Science for key computational science applications. At the OLCF, ALCF (Argonne Leadership Computing Facility), and NERSC (National Energy Research Scientific Computing Center), the computational science community has access to more than 200 petaflop/s of computational performance. These facilities are actively used by a wide community of researchers to simulate complex scientific problems at an unprecedented scale, resolution, and accuracy, which would be unattainable without the effective utilization of such facilities. Achieving this level of supercomputing resources was extraordinarily difficult, as it required consistent growth in the scalability and integration of all system components, including applications, system software, and hardware architecture. The ASCR community has been a key contributor to the success of these platforms, working closely with vendors, application scientists, and the HPC community to ensure that all requirements were met to allow the success of these new capabilities. In addition to procuring these flagship supercomputers, ASCR has been funding important activities that enhance the usability and effectiveness of these machines. Through the SciDAC (Scientific Discovery through Advanced Computing) program, the HPCS (High Performance Computing Systems) conference, the INCITE (Innovative and Novel Computational Impact on Theory and Experiment) program, and base funding, ASCR has enabled the development and deployment of scalable hardware, software, and applications that together have fostered high-end scientific computation at an unprecedented scale.

Examples of partnerships include the collaboration between ASCR, the NNSA (National Nuclear Security Administration), ANL, LLNL, and IBM in the revolutionary design of the Blue Gene architecture, which required 100x improvement in application scalability as well as reliability. This was achieved by the use of system-on-chip (SOC) coupled with keeping the design simple (low parts count); making the system software hierarchical with simple, highly replicated components; and integrated codesign of the system with the applications. The strong partnerships and funding for the

Blue Gene series of energy-efficient supercomputers enabled innovations that have been central to transformational scientific research around the world. In October 2009, the Blue Gene series was recognized by President Barack Obama as a recipient of the National Medal of Technology and Innovation, the nation's highest honor for technological achievement.

Another key ASCR-funded research program that included partnerships between ORNL, ANL, NERSC, SNL, PNNL, and Cray resulted in the Office of Science Cray XT systems' capabilities increasing from a few teraflop/s to systems capable of performing petaflop/s. Numerous collaborative avenues addressed the challenges of effectively scaling application performance three orders of magnitude in five years. A key component has been the DOE SciDAC program providing the resources to effectively develop an entire generation of new science codes, libraries, and tools for modern computer architectures. These investments were coordinated by ASCR and were critical to the OLCF and NERSC flagship supercomputers' success. Other examples include the NERSC, ORNL, and Cray partnership that developed a more functional Cray Linux Environment, an ultra-lightweight version of the standard Linux operating system on the Cray XT systems, as well as the Computational Liaison program which embeds computational scientists from OLCF in science teams to maximize application performance.

In collaboration with NNSA's ASC (Advanced Simulation and Computing) program, ASCR has recently funded three successive programs to accelerate technology development in the computing industry. These competitive programs, called FastForward, DesignForward, and PathForward, have covered specific research, design, and engineering costs to help critical hardware and software technologies mature from research ideas toward commercial products. These investments have improved DOE's options for subsequent procurements while also enhancing U.S. industrial competitiveness. The FastForward program awards were announced in 2012 and supported five computing companies with a focus on improvements to memory systems and node architectures. DesignForward awards for five vendors were announced in 2013 and targeted interconnection networks. The six recipients of PathForward grants were announced in 2017 and cover a broad range of technologies including memory interfaces, advanced interconnects, resilience, and integration. These investments accelerated the development and commercial availability of exascale technologies and substantially enhanced the procurement landscape for early exascale systems.

The overall impact of ASCR's industrial partnerships has been enormous. The coordination of DOE program offices, basic and applied research programs, laboratories, vendors, and computational scientists has moved the Office of Science further ahead in the realm of scientific computing than many thought possible. ASCR facilities now have world-leading computational capability, the best numerical libraries and tools, and most importantly, the scalable applications to meet the demands of science communities who use these resources. The Office of Science is able to lead at a time where the world's most challenging problems in climate modeling, energy assurance, and basic science are demanding answers through modeling and simulation, and has the capability to analyze massive datasets resulting from Office of Science-supported user facilities.

2.5 ASCR COMPUTING AND NETWORKING FACILITIES IN SEVENTH DECADE OF DRIVING SCIENTIFIC DISCOVERY

When the 54th edition of the twice yearly TOP500 list of the world's top supercomputers was posted in November 2019, the top two slots were held by the U.S. Department of Energy (DOE) supercomputers. Summit, installed at Oak Ridge National Laboratory (ORNL), was still first in the ranking. Sierra, housed at Lawrence Livermore National Laboratory (LLNL), maintained its hold on the No. 2 slot. Throughout the lifetime of the list, DOE systems have consistently placed at or near the top, based on their performance in running the LINPACK benchmark application.

From the earliest days of computing machinery in the 1950s, DOE and its predecessor agencies have led in the development and deployment of HPC (high performance computing) systems as well as the software tools and science applications necessary to advance scientific discovery. These advancements, complemented by DOE's high performance scientific computing initiatives and user facilities, have been key to creating the field of computational science.

Today, ASCR supports thousands of researchers at national laboratories and universities, as well as industrial partners, with three supercomputing user facilities—the Argonne and Oak Ridge leadership computing facilities (ALCF; OLCF), and the National Energy Research Scientific Computing Center (NERSC). Each of these facilities is at the forefront of scientific computing and their organizations have been, throughout the history of DOE's Office of Science, leading the way in supporting the application of HPC for science and engineering.

As the computing facilities were playing an increasingly important research role, the Office of Science also pioneered communications links connecting scientists to the centers and to one another to support collaborative science. Today, the Office of Science operates the Energy Sciences Network (ESnet), the world's fastest network dedicated to science. Due to its unique role, ESnet is considered a national user facility.

ASCR's HPC facilities have pioneered computer systems from the earliest days of pre-transistor systems built with vacuum tubes and plug boards to systems built with individual transistors or integrated circuits to more recent advanced machines such as vector systems, microprocessor-based systems, multicore systems, and even the accelerated systems that today boast hundreds of petaflop/s of computing power. Each of the centers, with the associated legislative measures that defined them, have made DOE a leader in HPC, enabling breakthroughs in materials science, nuclear energy, chemistry, climate science, physics, artificial intelligence (AI), data science, and more over the past four decades. These efforts have also been emulated by other organizations, further testifying to DOE's leadership.

Early History

In the early 1950s, Hungarian mathematician John von Neumann proposed the U.S. create a Mathematics program to investigate the use of digital computers for scientific, mathematical, and

technological problems. With his support, the Atomic Energy Commission (AEC) launched the Mathematics program.

In the 1950s, “computer” was a job description, not a physical system for performing calculations electronically. But that began to change with work at the national laboratories. At that time staff at Argonne National Laboratory (ANL) were developing the AVIDAC (Argonne Version of the Institute’s Digital Arithmetic Computer), a system that cost \$250,000. AVIDAC worked tens of thousands of times faster than an employed human computer and was used to solve problems in physics. That same year ORNL built its first system ORACLE (Oak Ridge Automatic Computer and Logical Engine), which was capable of performing 100 person-years of computing in only eight hours.



Figure X: ANL AVIDAC. ANL computer scientist Jean Hall works on AVIDAC, Argonne's first digital computer installed in 1953

A decade later, the first megaflop machine, the Controlled Data Corporation (CDC) 6600, was installed at LLNL and is now known as the first successful supercomputer. The CDC 6600’s logic module was designed by Seymour Cray, who would later be known as the father of supercomputing and went on to create a series of successful supercomputers bearing his name. A new era of computing had arrived.

Bringing Computation to Open Science

The founding of the first unclassified supercomputing center began in 1973 when Dr. Alvin Trivelpiece, then deputy director of the Controlled Thermonuclear Research (CTR) program of the AEC, solicited proposals for a computing center that would aid in reaching fusion power, giving the magnetic fusion program under CTR access to computing power similar to that of the defense

programs. LLNL was chosen as the site for the new center which would be called the CTR Computer Center (CTRCC), later renamed the NMFEECC (National Magnetic Fusion Energy Computer Center), and eventually NERSC. Starting with a cast-off CDC 6600, within a year from its inception the center added a new CDC 7600 and provided, for the first time, a remote access system that allowed fusion energy scientists at ORNL and Los Alamos National Laboratory (LANL), as well as the General Atomics research center in southern California, to communicate with the centralized computers.

The center continued to deploy leading-edge systems and in 1978, NMFEECC developed the Cray Time Sharing System (CTSS), which allowed remote users to interface with its Cray-1 supercomputer. At the time, computers were essentially custom machines, delivered without software, leaving centers to develop their own. Due to its success, CTSS was eventually adopted by multiple computing centers, including the National Science Foundation (NSF) centers established in the mid-1980s in Chicago, Illinois and San Diego, California. In 1985, when ORNL deployed a Cray X-MP vector processing system, the system also ran CTSS. NERSC next deployed the first four-processor system, the 1.9-gigaflop Cray-2, which replaced the Cray X-MP as the fastest in the world. Having been prepared for multitasking, CTSS allowed users to run on the Cray-2 just one month after delivery.

In 1983, the NMFEECC opened its systems to users in other science disciplines, allocating five percent of system time to the other science offices in DOE's Office of Energy Research, paving the way for a broader role of computation across the research community. By 1990, the center was allocating computer time to such a wide range of projects from all of the Office of Energy Research offices that the name was changed to NERSC.

The growing number of users and increased demand for computing resources led Trivelpiece, then head of DOE's Office of Energy Research, to make another decision that mapped out a path for making those resources more widely accessible. He recommended that DOE's Magnetic Fusion Energy network (MFEnet) be combined with the High Energy Physics network (HEPnet), to become ESnet (Energy Sciences Network) in 1986. ESnet's roots stretch back to the mid-1970s, when staff at the CTRCC installed four acoustic modems on the center's CDC 6600 computer.

Working in Parallel

Along with the Cray-2 and Cray X-MP vector systems, the DOE laboratories began to deploy parallel machines and accompanying software in the 1980s as it became clear that parallelism would be required to dramatically increase the performance of these large systems to meet the needs of DOE-supported researchers.



Figure X: Staff at the National Magnetic Fusion Energy Computer Center install a Cray-2 machine.



Figure X: Ewing "Rusty" Lusk at ANL's Advanced Computing Research Facility which fielded an array of early parallel systems.

ANL established its parallel computing program in the 1980s, the Advanced Computing Research Facility (ACRF). By the late 1980s, this facility housed as many as 10 radically different parallel computer designs—nearly every emerging parallel architecture—on which applied mathematicians and computer scientists could explore algorithm interactions, program portability, and parallel programming tools and languages. By 1987, ANL was hosting a regular series of hands-on parallel computing training courses on ACRF systems for attendees from universities, industry, and research laboratories.

ORNL invested heavily in a series of parallel computer systems, including the Intel Personal Supercomputer (iPSC) models (iPSC1, iPSC2, and iPSC860). ORNL also took delivery of its first shared memory parallel computer, a Sequent Balance 8000.

Several years later, as part of the High Performance Parallel Processing project with LANL, NERSC

Scaling Science to Teraflop/s

To increase the understanding of material properties, an ORNL team, led by Malcolm Stocks, developed the Locally Self-consistent Multiple Scattering (LSMS) electronic structure code in 1995 to solve the runaway problem of calculating the electronic structure for large numbers of atoms by assigning each individual atom to a single computer node. After porting their code to the Cray T3E massively parallel architecture, the team's modeling of metallic magnet atoms was run on progressively more powerful Cray T3E supercomputers, starting with NERSC's 512-processor machine, and won the 1998 Gordon Bell Prize with a top performance of 657 gigaflop/s.

However, the group later gained access to a machine on Cray's manufacturing center floor and achieved 1.02 teraflop/s (trillions of calculations per second), making the code the first application to reach a sustained performance of 1 teraflop. In 2000, this achievement was recognized with a ComputerWorld/Smithsonian Institution Award. In addition to receiving a medal in recognition of their achievement, the group's work on the project was accepted as part of the Smithsonian's permanent collection.

deployed a 128-processor Cray T3D machine, the first large-scale, parallel system from Cray Research, in 1994. The machine was used in a national laboratory-industry partnership to advance the development of parallel codes and upgraded to 256 processors within a year.

In 1987 the White House Office of Science and Technology Policy (OSTP) published “A Research and Development Strategy for High Performance Computing,” describing the strategies for supporting HPC, the software necessary for solving Grand Challenges, and the network and resources that would be crucial going forward. Congress passed the High-Performance Computing and Communications Act of 1991 (Public Law 102-194) to enact these recommendations in the form of a program called the HPCC (High Performance Computing and Communications) in 1994.

Big Changes in the 1990s

The HPCC program was originally developed by four agencies, ARPA (Advanced Research Projects Agency, NASA (National Aeronautics and Space Administration), DOE, and NSF. HPCC included efforts in four areas, High Performance Computing Systems (HPCS), Advanced Software Technology and Algorithms (ASTA), NASA Research and Engineering Network (NREN), and Basic Research and Human Resources (BRHR). The organization’s goals were to develop HPC and communication technologies, implement NREN to enable collaboration and networking amongst top researchers, provide training and education on HPC topics, and draw up solutions to the Grand Challenge problems.

The Grand Challenges for the ASTA component, requiring development of algorithms and software, included computational fluid dynamics and combustion dynamics, numerical tokamak computations, materials science, computational chemistry, quantum chromodynamics, global climate modeling, computational biology, and groundwater remediation.

Following the HPCC Act that solidified these goals, MICS (Mathematical, Information, and Computational Sciences), a forerunner of ASCR, issued a solicitation for High Performance Computing Research Centers (HPCRC), and in 1992 proposals from ORNL and LANL were selected.

ORNL joined with Ames, Brookhaven, and Sandia national laboratories as well as seven universities—Rice, SUNY Stony Brook, Texas A&M, University of South Carolina, University of Tennessee—Knoxville, University of Wyoming, and Vanderbilt—to submit the Partnership in Computational Science (PICS) proposal. ORNL proposed an HPCRC to deploy large-scale parallel systems, as development platforms for a new generation of highly parallel applications, to solve three Grand Challenge problems—groundwater remediation, materials modeling, and quantum structure. The Center for Computational Sciences (CCS) at ORNL was created as a user facility the same year to allow DOE and other federal agencies, academia, and industry access to a large parallel system and to explore additional parallel architectures for future systems.

In its first year, the CCS deployed a 32-node Kendall Square Research (KSR-1) system featuring a shared memory architecture in which all the memory in the system appeared as cache that could be

directly addressed by any of the system's processors. It also deployed two Intel Paragon systems of five and 35 gigaflops.

The Advanced Computing Laboratory was created at LANL as the second HPCRC and deployed a large-scale parallel system from Thinking Machines. Together the Paragons at ORNL and the Thinking Machines system at LANL provided the generally accessible, large-scale, parallel systems to develop the system software, algorithms, and applications to address the grand challenge applications listed in the HPCC program. In 1992, ANL acquired an IBM SP, the first scalable, parallel system to offer multiple levels of I/O (input/output) capability essential for increasingly complex scientific applications. The HPCRCs at LANL, ORNL, and ANL started focusing on experimental, production-oriented parallel computing for Grand Challenges, in addition to computer science, and emphasized collaborative research with computational scientists.




Figure X. ORNL Paragon. ORNL's Intel Paragon system was installed in 1995, comprising 3,072 processors to support research into Grand Challenge problems.

With modeling and simulation gaining traction in the HPC world, ORNL's CCS established the Computational Center for Industrial Innovation (CCII) to help industry increase its competitiveness. CCII focused specifically on advanced materials and environmental waste remediation.

In 1995, a 150-gigaflop/s Intel Paragon system was installed at ORNL. When it was delivered, with 1,024 nodes and 3,072 processors, it was among the fastest supercomputers in the world. The Intel Paragon machines at ORNL were largely used for the problems highlighted in the PICS proposal. These same problems outlasted the 1990s generation of supercomputers and helped lay the foundation for many of the codes used in modeling and simulation today, codes that would eventually need to be adapted to run on petascale, and ultimately exascale, systems.

In 1996, after NERSC and ESnet had been relocated to Lawrence Berkeley National Laboratory (LBNL), the CCII acquired the Cray T3E-600 system, its first massively parallel processor architecture machine which was upgraded to a T3E-900 the following year. The system brought with it a fundamental change in the computing environment, making it possible for scientists to perform larger and more accurate simulations. It also had the largest I/O system built to date with 1.5 terabytes of disk



Fig. X: T3E 900. NERSC's Cray T3E-900 provided the development platform for the first scientific code to reach a sustained performance of 1 teraflop.

storage and a read/write capability of 800 megabytes. Ranked No. 5 on the TOP500 list, this system, named MCurie, was the most powerful computer for open science in the U.S. NERSC's upgraded T3E-900 provided the training platform for a materials science project led by ORNL's Malcolm Stocks, whose code was the first application to reach a sustained performance of 1 teraflop.



Fig. 6: HPSS. HPSS like this grew out of a consortium launched in 1992 and today provide critical data storage management for DOE supercomputing centers and other institutions around the world.

As the laboratories developed their respective computing centers and associated research, they also shared a common challenge. The growth in computing power led to increasingly large data archives. Since no available commercial product could provide a highly scalable high-performance storage system, LBNL, LLNL, LANL, ORNL, Sandia National Laboratories (SNL), and NERSC partnered with IBM Houston Global Services to launch the High Performance Storage System (HPSS) Consortium in 1992. The goal was to provide scalable hierarchical storage management, archive, and file system services. Today the collaboration includes 39 research laboratories, universities, and companies in Asia, North America, and Europe.

ASCR Cements DOE's Role in HPC

As the Office of Science supercomputers continued to evolve into increasingly more capable systems serving an increasingly large community of scientists, DOE saw the need for an organization that could manage and coordinate the activities at its multiple institutions. In 1999, the MICS office was reorganized as ASCR to support research in eight Grand Challenge areas, computational accelerator physics; computational chemistry; high energy and nuclear physics; genome analysis; materials, methods, microstructure, and magnetism; numerical tokamak turbulence; particle physics; and protein dynamics and biocatalysis. ASCR would act as a conduit for research and development in these and other relevant technological areas. In the early 2000s, the CCS

at ORNL upgraded its IBM RS/6000 SP Eagle machine, featuring a Power3 architecture and became the first Office of Science system with a peak performance exceeding a teraflop. In 2001, NERSC deployed the Seaborg supercomputer, which was ranked the second most powerful system in the world. Seaborg was an IBM SP Power3 system with 3,792 nodes that achieved 3 teraflop/s on the TOP500 list in June 2001.

By 2003, NERSC was supporting more than 4,000 users from all the Office of Science program offices, and requests for time on its systems were three times what was available. At the direction of Office of Science Director Raymond Orbach, NERSC launched the INCITE (Innovative & Novel Computational Impact on Theory & Experiment) program, which created a system for scientists to apply for and receive large allocations of time on NERSC computing resources. INCITE was expanded to include the leadership computing facilities (LCF) in 2006. Now supported by the ANL and ORNL facilities, INCITE continues to provide critical HPC systems and support to a wide range of industry researchers, helping the nation maintain its science and engineering leadership across a range of industries.

In 2003, ORNL completed construction of the most advanced facility in the country for unclassified scientific computing. The National Center for Computational Sciences was designed specifically for leadership-class computer systems, providing ample physical space, cooling capacity, and electrical power to upgrade and add computers.

In 2004, the Office of Science issued a call for proposals in response to the DOE High-End Computing Revitalization Act to create a leadership computing facility. ORNL and ANL partnered on this proposal and won the competition. The CCS at ORNL was renamed the OLCF, which pursued two architectural paths, the expansion of the Cray X1 vector processing system to a Cray X1e system, and to deploy a Cray XT3 MPP system of 50 teraflop/s with planned expansion to more than 100 teraflop/s. Ultimately users were able to get much higher performance from the XT3 parallel system, called Jaguar, and the X1e vector system was retired. The XT3 system was Cray's commercial version of the Red Storm machine developed in a collaboration between Cray and SNL. In 2005, the Energy Policy Act of 2005 amended the High Performance Computing Act of 1991 (Public Law 102-194), solidifying the place of HPC systems at DOE by permanently authorizing the provisioning of advanced HPC systems, software, and applications.

High performance data storage for HPC

By the early 1990s, HPC systems were playing an increasingly important role in scientific discovery. More powerful computers meant researchers could study problems in fusion energy, combustion, astrophysics, subatomic particles, accelerators and other areas in ever-greater detail. However, these simulations generated ever-larger data sets and computer performance was slowed by the systems' inability to quickly write the data into storage.

In 1992, IBM joined with Lawrence Livermore, Los Alamos, Oak Ridge, and Sandia national laboratories and NERSC to develop the HPSS. Released to the HPC community in 1996, HPSS received an R&D100 Award in 1997.

"We were just starting to use parallel computers at that time and needed a plan for storing and retrieving big files," said Buddy Bland, now program director for the OLCF. "If this behind-the-scenes infrastructure isn't there, then scientists can't get their work done."

The resulting HPSS architecture was able to write and retrieve data in parallel to the storage system, initially at one gigabyte/s. The HPSS architecture is also scalable—in capacity, bandwidth, and metadata—as the computers scale up in size making it well positioned to support the next-generation exascale supercomputers. Bland said the ongoing collaboration, which expanded to include Lawrence Berkeley National Laboratory in 1997, is a result of the partnership between IBM and the national labs. Industry wouldn't have done the project alone and the laboratories can't devote the necessary level of staffing. "It's a model for collaboration between the national labs and industry," Bland said. "The sustained investment by ASCR and NNSA means that as the labs deploy very large cutting-edge systems, the necessary infrastructure for storing and accessing data is on the floor and ready to go."

With these leadership computing resources came the need for increasingly complex and scalable codes. ANL formed the Blue Gene Consortium with IBM and other national laboratories to design, evaluate, and develop code for a series of massively parallel computers. ANL's Laboratory Computing Resource Center then installed a five-teraflop IBM Blue Gene/L in 2005, a prototype and proving ground for what became the ALCF in 2006.

As the demand for access to HPC systems continued to grow, requests for time on the LCF systems rapidly exceeded the capability of the INCITE program to supply those resources. The LCFs worked with ASCR leadership to create two new allocation programs for access to the LCF supercomputers. The ALCC (ASCR Leadership Computing Challenge), administered by ASCR, and the Director's Discretionary (DD) program was established at ALCF and OLCF in 2006. ALCC supported Office of Science programmatic needs while DD awarded computing time to R&D, industry projects, development work for INCITE projects, or pilot projects.

In 2008, ALCF installed a 100-teraflop Blue Gene/P system, Intrepid, and was named the fastest supercomputer in the world for open science—and third fastest machine overall—on the TOP500 list. Intrepid entered production in 2009 and topped the first Graph 500 list in 2010 and 2011. In 2008, OLCF upgraded Jaguar to a delivered performance of over 1 petaflop/s, ranking as the world's second




Figure X: Roadrunner. The Roadrunner machine at LANL was the first system to achieve petaflop/s performance. Roadrunner was a hybrid system using AMD Opteron and IBM Cell processors.

fastest system. LANL ranked fastest in the world with its Roadrunner system, the first machine to deliver over 1 petaflop/s on the TOP500 benchmark. In 2012, Jaguar was upgraded one final time and renamed Titan. During this upgrade, an NVIDIA K20X Graphics Processing Unit (GPU) accelerator



Figure X: Titan. By upgrading its Jaguar machine and installing NVIDIA GPUs, ORNL created Titan, which reclaimed the No. 1 position on the TOP500 List.

was added to each node. Titan had a peak performance of 27 petaflop/s and was again the world's fastest supercomputer. While Titan was not the first GPU-based system, it was the first GPU-accelerated computing system that was generally available to the scientific user community through a DOE user facility. The same year, the ALCF's 10-petaflops IBM Blue Gene/Q, Mira, ranked third on the June TOP500 list and entered production in 2013.




Figure X: ANL Mira. Mira, a 10-petaflop/s IBM Blue Gene/Q, went into production at the ALCF in 2013.

In 2014 ORNL, ANL, and LLNL began working together on the acquisition of a pre-exascale system with performance targets in the 150–300 petaflop/s range for delivery in 2017–2018. CORAL (Collaboration of Oak Ridge, Argonne, and Livermore) issued a joint request for proposals to industry. The ALCF selected an Intel system, Aurora, based on Intel’s Xeon Phi (Knights Hill) processor. Both LLNL and the OLCF selected IBM systems based on the IBM Power9 CPU and NVIDIA Volta GPUs. Summit (OLCF) debuted at the top of the June 2018 TOP500 list, while Sierra (ALCF) entered the list at No. 3.

In November 2015, LBNL opened Shyh Wang Hall, a 149,000-square-foot facility housing NERSC, ESnet, and researchers in the laboratory’s Computational Research Division. The facility is one of the most energy-efficient computing centers anywhere, tapping into the San Francisco Bay’s mild climate to cool NERSC’s supercomputers and eliminate the need for mechanical cooling and earned a Gold LEED (Leadership in Energy and Environmental Design) certification from the U.S. Green Building Council.

The facility soon became home to NERSC’s next system, Cori, a 30-petaflop/s Cray system with Intel Xeon Phi (Knights Landing) processors. With 68 low-power cores and 272 hardware threads per node, Cori was the first system to deliver an energy-efficient, pre-exascale architecture for the entire Office of Science HPC workload. In 2019, NERSC began preparing for the installation of its next-generation, pre-exascale Perlmutter system, a Cray Shasta machine which will be a heterogeneous system comprising both CPU-only and GPU-accelerated cabinets.

Performance Modeling

Another way ASCR has influenced computer architecture is by investing in predictive models of system performance. These models can be used to predict performance before a system is built, guiding and supporting design and procurement decisions. They can also be used to assess limits to scalability of algorithms, to support application tuning, and to understand tradeoffs in system balance that are otherwise impractical to gather empirically. A number of important ASCR-funded performance modeling activities have had significant impacts.

The performance profile of a given system/application combination depends on numerous factors, including system size, system architecture, processor speed, multi-level cache latency and bandwidth, interprocessor network latency and bandwidth, system software efficiency, algorithms used, programming language, problem size, amount of input and output, and more. Indeed, a comprehensive model must incorporate most if not all of the above factors. Given this complexity, models have historically been limited in scope and required large amounts of computer time and extensive expertise from researchers. ASCR investments have led to unprecedented capability which has significantly impacted procurements and applications. These models have come in several forms: runtime estimation, system design, system tuning, application tuning, and system procurement.

Runtime estimation. The most common application for a performance model is to enable a scientist to estimate the runtime of a job when the input parameters for the job are changed, or when a different number of processors is used. One can also estimate the largest size of system that can be used to run

a given problem before the parallel efficiency drops to an unacceptable area. Around the turn of the century, the state of the art at LANL, and at the Performance Modeling and Characterization (PMAc) laboratory at the University of California San Diego, had improved to the point where many models have proven 95 percent accurate and the production of new models is automated or semi-automated such that new models can be produced quickly without specialized domain-science expertise.

System design. Performance models are frequently employed by computer vendors in their design of future systems. Typically, engineers constructed a performance model for one or two key applications, then compared future technology options based on performance model projections. Now that performance modeling techniques are better developed, it is possible to target many more applications and technology options in the design process. For example, the PMAc and LANL independently forecast the performance of ORNL's Jaguar post-quadcore upgrade, helping to provide confidence in the technical design decisions. When the system was built the predictions were shown to be within five percent accuracy on strategic applications, including S3D, a parallel direct numerical simulation solver developed by SNL.

System tuning. A performance model can also be used to diagnose and rectify system problems. LANL discovered that when a recent system was installed its performance fell below model predictions by almost a factor of two. Further analysis uncovered some system difficulties, which, when rectified, improved performance to nearly the level the model predicted. When observed performance of a system falls short of that predicted by a performance model, it may be the system that is wrong, not the model.

Application tuning. Models can be used to anticipate how application performance will be impacted by algorithm choices like the cache-blocking factor. This approach allows for much more efficient exploration of algorithmic tradeoffs and performance optimization. Simple performance models can even be incorporated into an application code, permitting on-the-fly selection of different program options. The Performance Engineering Research Institute (PERI) a research collaboration funded through DOE's SciDAC and SciDAC2 programs, used models to guide tuning that resulted in greater than 10 percent performance improvements on S3D. The PMAc Laboratory had the Gordon Bell Prize finalist entries in 2007 and 2008 for application performance based on model-guided tuning resulting in world records for size and detail of a non-hydrostatic weather model (WRF), and the frequency of a seismic wave propagated through the Earth (SpecFem3D).

More recently, the Roofline performance model, from LBNL, has provided guidance to applications users on performance expectations for their applications. Roofline is a throughput-oriented performance model used to compare application performance to machine capabilities along the axes of performance, bandwidth, and data reuse. Roofline has proven exceptionally nimble in its ability to be applied to ever more powerful CPU and accelerated architectures running throughput-oriented scientific and deep learning applications. It can also allow users to select appropriate algorithms before going to the trouble of implementing all the possible choices.

System procurement. Arguably, the most compelling application of performance modeling is to inform the selection process of a new computing system. Procurement decisions rely heavily on benchmark performance. But it can be time consuming for vendors to optimize good sets of benchmark codes and often the final system hardware does not exist at the time bids need to be submitted. Recent advances have greatly improved this process and several federal agencies, including the DOE, Department of Defense, and National Science Foundation, now use ASCR-funded modeling methods to streamline and improve procurements.



Figure X: Cori. Installed at NERSC, the 30-petaflop/s Cori system provided the first pre-exascale architecture for the broader Office of Science research community.

Getting Big Science out of Big Data

In recent years, ASCR's user facilities have continued their leadership in HPC, deploying resources that will be capable of solving new problems in data science with the use of AI (artificial intelligence) and machine learning.

To help scientists address their growing data needs, NERSC created a Data Department in 2015 and has developed a vigorous program on deep learning for science. At the same time, OLCF began supporting users with end-to-end workflows that enable data-intensive computing by creating the Advanced Data and Workflows Group. Driven by the growth of experimental and

observational data across DOE facilities and scalable analytics requirements, OLCF began actively offering data science toolkits and resources for its users. In 2016, the ALCF Data Science Program, an initiative to explore new ways to foster data-driven discoveries, was launched with an eye toward

growing a new community of HPC users. The program targets users with “big data” science problems and provides time on ALCF resources, staff support, and training to improve computational methods across all scientific disciplines.

ESnet continues to provide DOE’s state-of-the art external networking capabilities to help scientists move data among the laboratories and users for analysis and archiving. ESnet connects many of the DOE laboratories using 100 Gbps links. In December 2015, ESnet launched four transatlantic links with a total capacity of 340 Gbps to provide researchers in the U.S. with secure, high-speed access to data from experiments at the LHC (Large Hadron Collider) in Switzerland. In 2017, the total transatlantic capacity was upgraded to 400 Gbps.

Typically, most HPC centers export more data than they import, as researchers create detailed simulations and move the resulting data to their home institutions. There is now a net flow of data into NERSC due to HPC’s increasing role in the analysis of data from experimental facilities. NERSC has a long history of supporting data-driven science and many high energy and nuclear physics teams have used NERSC data systems and the Parallel Distributed Systems Facility cluster for their analyses.

Each of the ASCR facilities work with DOE’s experimental facilities to provide data storage and analysis. NERSC and OLCF maintain close collaborations with DOE’s Joint Genome Institute, which conducts high-throughput DNA sequencing, synthesis, and analysis in support of DOE’s Biological and Environmental Research (BER) program. NERSC provides the primary storage and analysis site while OLCF provides geographically diverse backup copies of the data. Another BER program, the Atmospheric Radiation Measurement (ARM) program, gathers climate data for research and makes it available to researchers around the world.

All three of ASCR’s computing facilities have begun executing several efforts to support programmatic job submission and monitoring interfaces, gateways, flexible scheduling, and analysis and visualization. One ALCC project, led by ANL computational scientists in collaboration with researchers from Fermi National Accelerator Laboratory and the LHC ATLAS (A Toroidal LHC ApparatuS) experiment, aimed to develop an end-to-end simulation workflow on DOE supercomputers that delivers improved HPC support to particle physicists. This work led to the Singularity service developed at LBNL several years later to provide a framework for creating and running containers—platforms designed to package code and its dependencies facilitating fast and reliable switching between computing environments.

In 2015, NERSC developed and released Shifter, a Docker-based container technology that allows users to bring their custom compute environment to NERSC’s supercomputers. Shifter was originally inspired by the need to improve the flexibility and usability of HPC systems for data-intensive workloads, but use cases are expanding to include general HPC workloads. In 2016 NERSC demonstrated that Shifter could be used to run complex, scientific, Python-based codes in parallel on more than 9,000 nodes on Cori. The OLCF worked closely with IBM to scale their container technology to the Summit system.

In addition to the container technology, each of the ASCR facilities is providing a variety of data analysis tools and frameworks that work at large scale to support the needs of their users, for example R in a scalable form (parallel big-data R, or pbdR). Other analytics resources include the support for native and containerized operations of leading open source frameworks for machine learning and deep learning including TensorFlow, PyTorch, Keras, and others. To enable data coordination and operations across facilities, user-operated workflow and workload programs such as Singularity, Pegasus, Fireworks, BEAM, OpenShift, Kubernetes, Slate, and Panda are supported by facility staff.

In the emerging field of deep learning for science, training performance on the Cori supercomputer at NERSC was increased to 15 petaflop/s in 2017, giving climate scientists the ability to use machine learning to identify extreme weather events in the output of huge climate simulations. In 2018, ORNL and LBNL users each demonstrated machine learning performance in excess of one exaflop/s of performance using the Summit system. Both of these projects won ACM Gordon Bell Prizes.

In 2017, ALCF launched Theta, an Intel/Cray machine, doubling the capacity to do impactful science and operate at the frontier of data-centric and high performance supercomputing. The ALCF continues to add new services, helping researchers manage workflow execution of large experiments, and co-schedule jobs between ALCF systems. Additionally, through substantial awards of supercomputing time, the ALCF makes the facility's computing and data analysis resources available to scientists in distant places who are tackling large and complex problems in science and engineering.

As part of the effort to prepare the leadership scientific codes for exascale systems, NERSC is currently evaluating applications for the NERSC Exascale Science Application Program (NESAP) that will prepare application workloads for Perlmutter, its future HPC system that will feature AMD CPU's and NVIDIA GPU's. NESAP will include three designated project categories, simulation, machine learning, and data processing and analytics. Both OLCF and ALCF are also evaluating submissions for early science programs for the Frontier and Aurora exascale systems.

To exascale...and beyond

The ASCR facilities are part of a broader community working to achieve a capable exascale computing ecosystem for scientific discoveries. The benefits of exascale computing—computing capability that can achieve at least a billion billion operations per second—is primarily in the applications it will enable. The facilities seek to contribute to the emerging convergence of supercomputing, big data analytics, and machine learning across a wide variety of science and engineering domains and disciplines.

Both the ALCF and OLCF have contracts for their future exascale systems, Aurora and Frontier, respectively. Aurora and Frontier will use innovative technologies to provide more than 1,000 petaflop/s for R&D in three areas, simulation-based computational science, data-centric and data-intensive computing, and learning (machine learning, deep learning, and other AI techniques). Both systems are expected to be delivered in 2021, with Aurora available to users in 2021, and Frontier available in 2022.



Figure X: ANL Aurora. Scheduled for delivery to the ALCF in 2021, Aurora will be the Office of Science's first exascale computing system.

The ALCF and OLCF have each inaugurated an Early Science Program to prepare key applications and libraries for their exascale systems. ALCF computational scientists and performance engineers are working closely with Argonne's Mathematics and Computer Science Division, the seedbed for such groundbreaking software as BLAS3, p4, Automatic Differentiation of Fortran Codes (ADIFOR), Portable, Extensible Toolkit for Scientific Computation (PETSc), Globus, and MPICH implementation of the Message Passing Interface. With Argonne's Computational Science and Data Science and Learning divisions, the Early Science Program aims to advance the boundaries of high performance computing technologies ahead of Aurora. The OLCF Center for Accelerated Application Readiness team is working with application scientists from the OLCF's Early Science Program to port dozens of GPU-enabled applications from Titan and Summit to the Frontier system.

With the end of Moore's Law in sight, LBNL and other DOE Labs are conducting research in quantum computing, developing quantum chemistry and optimization algorithms, as well as prototype superconducting quantum processors. Lab researchers proved the viability of their work by using these algorithms on a quantum processor comprising two superconducting transmon quantum bits to successfully solve the chemical problem of calculating the complete energy spectrum of a hydrogen molecule. In 2017, DOE funded two projects to build on this momentum. One team received \$1.5 million over three years to develop novel algorithms, compiling techniques, and scheduling tools that will enable near-term quantum computing platforms to be used for scientific discovery in the chemical sciences. The other team will work closely with these researchers to design prototype four- and eight-qubit processors to compute these new algorithms. This project will last five years and the researchers

will receive \$1.5 million for their first year of work. By year five, the hardware team hopes to demonstrate a 64-qubit processor with full control.

Networking at the Speed of Science: The Energy Sciences Network

ESnet and LBNL have pioneered many advances in networking to support scientific research. ESnet is an Office of Science national user facility and provides high-bandwidth connections between more than 50 sites, including the entire national laboratory system, its supercomputing facilities, and its major scientific instruments, as well as more than a dozen U.S. university high energy physics groups, and connections to collaborators around the globe.

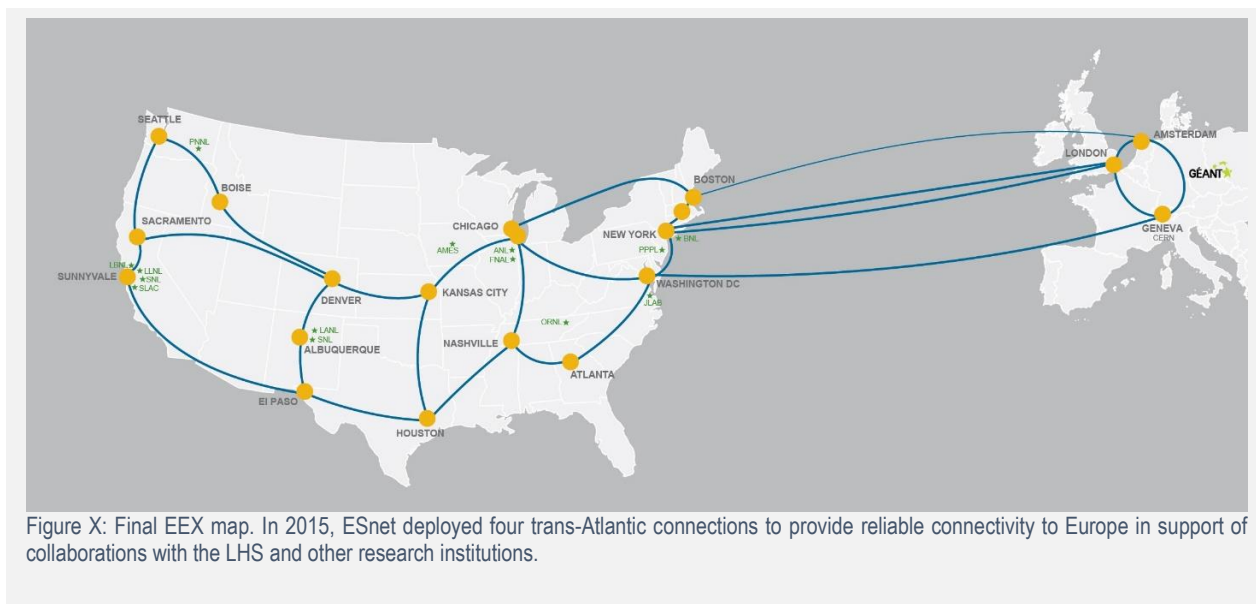
ESnet is the largest mission-science network in the world, consisting of multiple hundred-gigabit links around the U.S. and to Europe. ESnet currently transports about 3 petabyte/s per day between ESnet sites and collaborators, and the level of ESnet's traffic has grown exponentially over the past three decades. This growth is driven by the rising tide of data produced by advanced scientific instruments, data-intensive science operated by global collaborations that can involve thousands of researchers, specialized facilities like the LHC, light sources like the LCLS (Linac Coherent Light Source), digital sky surveys, more powerful supercomputers, etc.



Figure 2. ESnet ca 2017.

As important as bandwidth is, global connectivity or reach is just as important for collaborative science. As a result of 30 years of working with the U.S. and global research and education networking community, ESnet is a highly connected network with linkages to more than 150 research and commercial networks as well as cloud providers. ESnet provides 100Gbps connectivity to most of the major North American and European research institutions, as well as networks in Japan and Korea,

ensuring that DOE scientists can collaborate and exchange data with scientists and instruments almost everywhere.

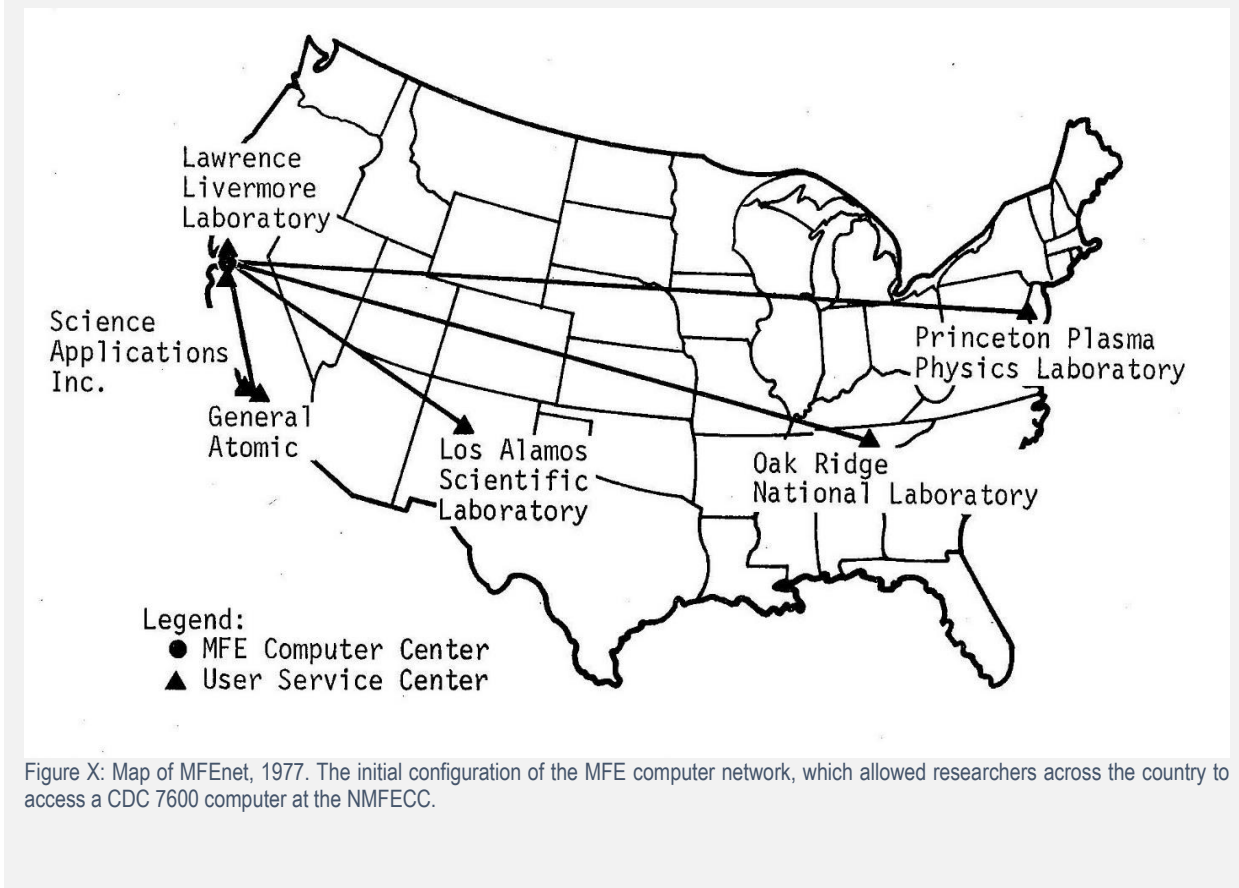


The First DOE Science-Oriented Network

In May 1974, LBNL connected its CDC 6600 computer to the ARPANET (Advanced Research Projects Agency Network)—the first generation of the Internet—thereby launching the concept of today’s ubiquitous networked computing. Two months later, NERSC went online at LLNL, allowing fusion researchers at the Princeton Plasma Physics Laboratory (PPPL) to access the machine via four dial-up modems.

Planning began almost immediately to create a new communications network, the Magnetic Fusion Energy (MFEnet), to serve the new computing center. Program Data Processor-10s (PDP-10s) were installed at six sites beginning in July 1975, interconnected by 50 kilobit/s links. By October of that year, scientists at PPPL, LANL, ORNL, LLNL, General Atomics, and Science Applications Inc. could access a handful of systems using MFEnet.

By 1985, MFEnet connected dozens of sites, including 11 national laboratories, DOE headquarters, 12 universities, and six companies engaged in fusion research, as well as collaborators in Japan. In 1986, due to increasing computing power at NMFEECC, and the allocation of computing time to researchers in other scientific disciplines, DOE Office of Energy Research Director Dr. Alvin Trivelpiece directed the Office of Scientific Computing Research to move forward to implement the Energy Sciences



Network, which combined MFEnet with the High-Energy Physics network (HEPnet). Operated by NMEFCC, ESnet began providing networking services in January 1988.

Just as ESnet was launching as a dedicated science network, the burgeoning Internet, heralded as the “Information Superhighway” as it carried streams of data at an unprecedented 56 kilobit/s, itself was facing imminent collapse due to network congestion.

Saving the Internet

In October 1986, the Internet stream slowed to a trickle. For example, the transmission rate between LBNL and the University of California at Berkeley, only a quarter-mile away, slowed to 320 bit/s. Electronic mail that had gone through in minutes took an entire day leading many to say the Internet could not scale. The whole concept of network communication was threatened. Van Jacobson, a network researcher at LBNL, took on the problem and designed the solution that is still in use today.

The cause of “congestion collapse” was not a network problem, but rather an end-computer problem due to a design glitch in TCP (transmission control protocol). Faced with a congested network, the old TCP hammered at the network over and over, trying to get its data through. With every computer attached to the Internet doing the same thing, the network got more and more congested—to the point of not transporting anything.

Jacobson, who led LBNL's Network Research Group, worked with Mike Karels of the University of California at Berkeley UNIX development group, to come up with the solution. They redesigned the TCP software to introduce a slight and variable delay before a packet was retransmitted into a congested network, which solved the problem. Jacobson and Karels outlined their "slow start" approach in their landmark 1988 paper "Congestion Avoidance and Control." They then worked with computer vendors to get the slow start algorithm into the vendor TCP implementations.

TCP remains the dominant protocol for moving data on the Internet, and Jacobson is widely credited with enabling the Internet to expand in size and support increasing data speed demands.

DOE's decision to transfer management of ESnet to LBNL in 1995 allowed for the laboratory's expertise in networking research to mesh with ESnet's operational expertise, resulting in technology innovations, support for increasingly collaborative research, and tools that have benefitted the global Research & Education networking community.

Pushing Networking Technology

Over the years, ESnet has consistently introduced new networking technologies, including Asynchronous Transfer Mode (1995), differentiated services (1998), IPv6 (2002), 10G optical networking (2003), network monitoring software (2006), bandwidth reservation software (OSCARs) and the start of Software-Defined Networks (2009), the Science DMZ (Demilitarized Zone) architecture for high-speed data access (2010), the first 100G network for science (2012), the first 400G production link for science (2015), demonstration of end-to-end science workflows targeted towards exascale (2018), and using telecom dark-fiber to be a seismic sensor (2019).

ESnet in the 2000s: Advanced networking for data-intensive science

The LHC at CERN in Switzerland is the world's largest scientific experiment in terms of the size and cost of the accelerator and the detectors, which are critical for two large-scale experiments, ATLAS and CMS (Compact Muon Solenoid), in terms of the volume of data generated and analyzed, and size of the collaboration. The data analysis of the steady state 50 Gbps of data from the LHC detectors is performed via the highly coordinated use of 20,000 CPU's at 200 institutes in 35 countries, some 500 petabytes of disk storage, and 800 petabytes of tape storage. Realizing that it was no longer feasible to ship data tapes out for analysis, the LHC analysis systems were designed to be network-based from the beginning. Much of the LHC data is stored at DOE's Brookhaven and Fermi national laboratories.

Starting in 2005, ESnet saw an exponential increase in LHC-related network traffic from operational testing at scale using simulated data. This continued to, and through, the LHC turn-on in 2010 and proved to be a good prototype for network demands by other DOE experiments on the horizon.

To prepare for this era of big data science, ESnet deployed a series of metropolitan area ring networks to connect national laboratories directly to the wide-area core network using dual independent connections. ESnet then designed and started implementing a national optical fiber network core that

would support multiple 100 Gbps network paths. ESnet also embarked on a program to provide tools to support very high-speed end-to-end data transfers.

Getting ready for 100G networking

With funding from the American Recovery and Reinvestment Act of 2009, ESnet launched its Advanced Networking Initiative in 2012 with the goal of building the world's fastest science network. To ensure the network would work as planned and to avoid disrupting the existing network, ESnet staff first built two test systems:

- a 100 Gbps long-haul prototype network linking ASCR supercomputing centers in California, Illinois, and Tennessee with a peering network in New York, and
- a rapidly reconfigurable, high-performance, national-scale test bed available for use by the research community and industry.

By the end of 2012, ESnet was again operating the world's fastest science network. ESnet5, served the entire national laboratory system, its supercomputing centers, and its major scientific instruments at 100 Gbps—10 times faster than its previous generation network—and connected all of these facilities to major scientific institutions around the world.

Evolving to 400 Gbps networks

In 2015, ESnet and NERSC at LBNL built a 400 Gbps super-channel, the first-ever 400G production link to be deployed by a national research and education network. The connection, nicknamed the BayExpress, provided critical support for NERSC's 6,000 users as the supercomputing center moved from its location in Oakland, California to the main campus of LBNL, leaving the tape storage system in Oakland.

The project was also important to the mission of the Office of Science. As scientists generate massive datasets at experimental facilities and through modeling and simulation at supercomputing centers like NERSC and the LCFs, the growth in data is threatening to oversaturate that capacity—especially as supercomputers approach exascale capabilities—and ESnet needs to prepare for the next level of capacity.

To develop the 400 Gbps production connection as a step toward the next level of networking, ESnet and NERSC joined with vendors Ciena and Level 3 Communication to deploy a pair of 200-Gbps-per-wavelength-per-optical channel links, which allowed for a simple-to-deploy dual subcarrier 400 Gbps solution. As part of the project, the team also set up a 400 Gbps research testbed for assessing new tools and technologies without interfering with production data traffic and allowing staff to gain experience operating a 400 Gbps optical infrastructure

Services for Advanced Networking Supporting Data-Intensive Science

ESnet has pioneered and provides a number of services to facilitate advanced, very-high-bandwidth use of the network in support of big data science.

OSCARS, NSI and software defined networks

In 2006, ESnet and Internet2 demonstrated an automated system for providing on-demand, end-to-end, multi-network bandwidth reservation service to support large-scale and time-critical data transfers. Three years later ESnet rolled out a production version called OSCARS (On-Demand Secure Circuit and Advance Reservation System), enabling network engineers and users to provision end-to-end dynamic circuits across multiple networks when and where they need them.

OSCARS, an early example and forerunner of what is now called Software Defined Networking (SDN), received a 2013 R&D 100 Award. Two years later, former U.S. Secretary of Energy Ernest Moniz recognized ESnet staff members with a DOE Secretarial Honor Award for their development of OSCARS.

OSCARS has since evolved into a standard called Network Services Interface (NSI) and has been adopted by multiple Research & Education networks around the world. Currently, the team is working on taking this model end-to-end through an ESnet research project called SDN for End-to-end Network Services at Exascale (SENSE) that is being integrated into DOE's Exascale Computing Project's Data Analytics at Exascale for Free Electron Lasers project in conjunction with the LCLS-II (Linac Coherent Light Source).

Designing the Science DMZ

A nearly insurmountable task for sites involved in data-intensive science was to get the data into and out of the site across the security perimeter, frequently a firewall that could not handle high-bandwidth data streams and a router not optimized for high-speed, wide-area network transport.

ESnet developed a network architecture design pattern, the Science DMZ, which enables a site's high performance data servers to securely connect directly to the wide area network. Science DMZ provides secure access into the site for access to storage and compute resources, and provides a highly secure, outward-facing interface for remote access to data located on data transfer nodes within the Science DMZ.

ESnet's Science DMZ Breaking Down Barriers, Speeding up Science

As science has become more collaborative and more data intensive, traditional network architectures designed for normal network traffic can create nearly insurmountable hurdles to discovery. Typical firewalls can't handle high-bandwidth data transfers, nor are routers optimized for high-speed, wide-area network transport.

To address this situation without sacrificing security, ESnet developed the Science DMZ.

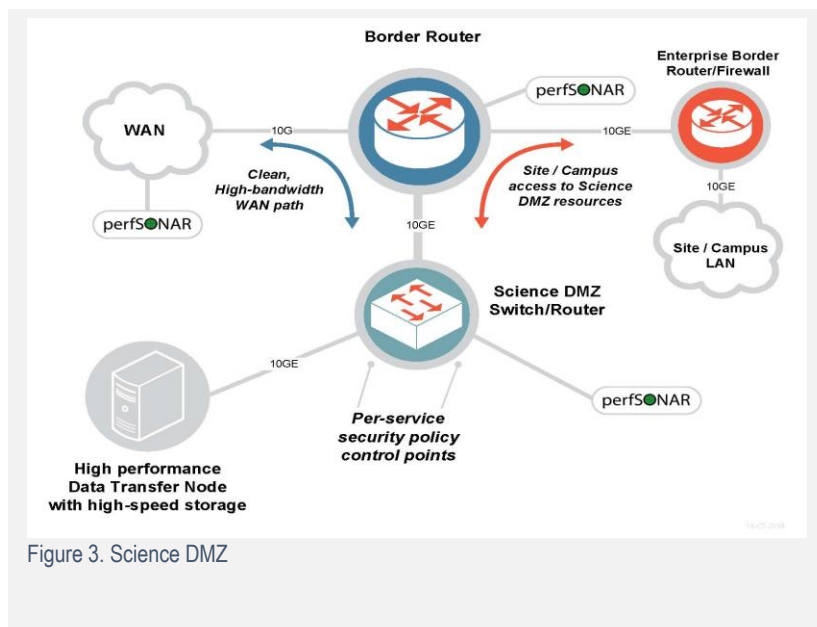
Science DMZ was first presented to the networking community in 2011. Since then it has been adopted by most DOE sites and endorsed by NSF, which has funded implementations at more than 100 universities. The architecture is also a key part of the Pacific Research Platform, an NSF-funded, science-driven high-capacity data-centric "freeway system" linking universities, DOE national labs and supercomputing centers on the West Coast.

"[ESnet] defined the Science DMZ and took it to the DOE science community. NSF has now cloned this approach through the Campus Cyberinfrastructure—Network Infrastructure and Engineering (CC-NIE) program over the past three years. It's been built out on over 100 campuses and these Science DMZs are all based on ESnet's approach," said University of California at San Diego Professor Larry Smarr, founder and director of the California Institute for Telecommunications and Information Technology (Calit2) and founding director of the National Center for Supercomputing Applications.

To further spread the word about the Science DMZ, ESnet partnered with Internet2 and Indiana University to co-lead 23 workshops as part of the Operating Innovative Networks series, which ended in 2018. Through both in-person and online workshops, which devoted the first day to the design, the organizers reached an estimated 750 network employees at 360 institutions in 39 states and 38 other nations.

perfSONAR Performance Monitoring Software

Monitoring and testing the network end to end is the only way to keep multi-domain, international-scale networks error-free, which, because of how TCP works, is essential for high-speed data transfers.



ESnet was instrumental in the design and implementation of perfSONAR, a software package which provides a standardized way to test, measure, export, catalogue, and access performance data from many different network domains (service providers, campuses, etc.). A multi-domain network monitoring system, perfSONAR is deployed extensively throughout LHC-related networks, in international networks, and at end sites.

ESnet's Sixth-generation network built on SDN expertise

ESnet's next-generation network, ESnet6, is designed to help the DOE research community handle the ever-increasing data volumes of advanced science. ESnet6 will provide more bandwidth, greater flexibility, and faster data transfer capabilities.

With a projected availability in 2023, ESnet6 will feature an entirely new software-driven network design that enhances the ability to rapidly invent, test, and deploy new innovations. Generally speaking, SDN is an emerging technology paradigm aimed at making it easier for software applications to automatically configure and control the various layers of the network to improve flexibility, predictability, and reliability. ESnet6 represents a transformational change in network capacity, resiliency, and flexibility that will bring tangible benefits to the DOE mission, including exascale computing and data.

ASCR's Supercomputing Landscape

Today, ASCR's three supercomputing centers are among the world's leading computational facilities in terms of systems, software, and user support. Many of the supercomputers available at ASCR's facilities feature groundbreaking architectures.

The 7,000 users at NERSC, oldest of the three, make it the most scientifically productive HPC center in the world. In 2018 NERSC users produced more than 2,500 refereed publications; over the years the center has been associated with six Nobel Prize-winning scientists or their teams.

NERSC's current supercomputers were the first of their kind. The recently retired Edison, a Cray XC30, included the first two cabinets of its line developed by Cray and DARPA (Defense Advanced Research Projects Agency). Cori, a Cray XC40, is the first and largest system with Intel Xeon Phi (Knights Landing) processors, debuting as the fifth most powerful supercomputer in the world in 2016.



Figure X: ORNL Summit. Installed in 2018, Summit at the OLCF has held the No. 1 position on the TOP500 List four consecutive times. Using IBM Power9 CPUs and NVIDIA Tesla V100 GPUs, Summit achieved performance of 148.6 petaflop/s.

In addition to its computing capabilities, Cori was the first large system to deploy an all-FLASH burst buffer that provided a worlds-best 1.7 terabyte/s system bandwidth.

OLCF's current systems were also novel architectures: the 27-petaflop Cray XK7 Titan was DOE's first ever hybrid CPU-GPU architecture supercomputer; and Summit, its IBM AC922 machine, is the first system to use IBM's Power9 architecture and NVIDIA's Volta GPUs, debuting as the fastest supercomputer in the world on the June 2018 TOP500 list.

Additionally, each facility maintains a large-scale data analysis and storage center based on the HPSS consortiums' work. NERSC's tape archive holds 150 petabyte/s of data, some of which date back to the early days of the center in the 1980s. Recent upgrades to the OLCF's storage system resulted in a data center that can also hold up to 150 petabyte/s of data, and the ALCF's HPSS can hold 64 petabyte/s of data.

Today, the resources at the three centers are allocated by various programs. NERSC's systems are primarily allocated by Office of Science program managers through an annual process known as the Energy Research Computing Allocations Process (ERCAP) to support program priority research. Eighty percent of time is allocated through ERCAP, an additional 10 percent of time goes to the competitively awarded ALCC program, and the final 10 percent is distributed at the discretion of the NERSC director. OLCF and ALCF allocate 60 percent of their computing time to the INCITE program, 20 percent to the ALCC program, and 20 percent to the DD program. Each center also has a number of industrial collaborations and small business users.

Conclusion

From its earliest days, DOE has pioneered the use of computing systems to tackle science and engineering problems critical to the national interest. As supercomputers became more common and offered increased capabilities, ASCR supported research into both furthering those capabilities with an eye toward scientific discovery and developing the necessary software and hardware. As a result, DOE's national labs have deployed systems that have consistently been among the most powerful in the world, as recognized by the TOP500 list, but have also pushed the boundaries of performance in running scientific applications, as evidenced by many ACM Gordon Bell Prizes awarded to users of DOE systems.

As demonstrated over the years, ASCR has supported research into new architectures, often when such ideas are in the very early stage. This has helped the Office of Science consistently stay on the leading edge of computational science, rather than working to keep up. This was the case when DOE computing moved from hand-built machines to vector systems, then to massively parallel computers and on to hybrid systems. Now as the era of exascale looms on the near horizon, ASCR is looking farther forward toward quantum computing and beyond Moore's law. And to ensure that such systems can be effectively used, ASCR is funding the adaptation of existing codes and development of new applications to maximize the returns on these critical investments.

3.0 ASCR INVESTMENTS HELP DRIVE INDUSTRIAL DISCOVERY AND ECONOMIC COMPETITIVENESS

As the single largest supporter of basic research in the physical sciences in the U.S. Department of Energy's (DOE) Office of Science supports tens of thousands of researchers across the country. Their findings and science results are published in publicly available journals for use by other scientists in both the public and private sector. High performance computing (HPC) facilities, managed by the Office of Science's Advanced Scientific Computing Research (ASCR) office provided the computational resources critical to many of these results. For example, each year the 7,000 users of NERSC (National Energy Research Scientific Computing Center) consistently publish more than 2,000 scientific papers citing the center's role in their work.

While industry has access to these papers, ASCR also supports a range of research projects and programs that engage industry, with ASCR-funded programs and facilities working directly with industry researchers. This support ranges from computing time for large-scale manufacturers to collaborations at the individual researcher level, from making powerful software libraries publicly available to seed funding for small startup companies.

Software Libraries Boost Circulation of ASCR Expertise

A great deal of activity in algorithms and software for solving linear algebra problems by ASCR-funded researchers over the past 30 years has led to the development of some of the most accurate and efficient solvers for studying complex scientific problems. ASCR has supported this work through its longstanding applied mathematics programs and SciDAC (Scientific Discovery through Advanced Computing) collaborations. Although originally developed for scientific uses, the solver libraries are also being used by industry to improve vendor software and tackle science and engineering problems.

LAPACK and ScaLAPACK software – Setting the standard for dense matrix solvers

Dense linear algebra forms the core of many scientific computing applications and there is continuous demand for increasingly better algorithms in terms of improved reliability, accuracy, robustness, ease of use, and most importantly, more efficient use of computing resources to speed up computation.

The linear algebra community recognized the need for help in developing algorithms into software libraries. A de facto standard for basic operations required in linear algebra algorithms and software was developed, resulting in the LAPACK and ScaLAPACK libraries. Known as dense matrix solvers, LAPACK and ScaLAPACK are suitable for scientific problems in which every part of the problem interacts with nearly all of the others. These libraries are widely used for numerical linear algebra within DOE and in many scientific applications on a wide range of advanced-architecture computers.

Developers of LAPACK and ScaLAPACK perceived the key to attain high performance was to minimize the frequency with which data is moved between different levels of the memory hierarchy. By using block-partitioned algorithms, the developers reduced the total amount of data moved, as well

as the total number of messages needed to move it, both of which can represent significant fractions of the total runtime. Secondly, the performance of an algorithm can be tuned by varying the data layout parameters.

More than 273 million web hits at Netlib.org for the associated libraries LAPACK, ScaLAPACK, CLAPACK and LAPACK95 demonstrate the demand for such solvers. These libraries have been adopted by many vendors and software providers as the basis for their own libraries. The companies include AMD, Apple (under Mac OS X), ARM, Cray, Dell, Fujitsu, HPE, IBM, IMSL, Inspur, Intel, Lenovo, Mathematica, Microsoft, MSC Software (in Nastran and Marc), NEC, NVIDIA, several LINUX distributions, NAG, the MathWorks (producers of MATLAB), InteractiveSupercomputing.com, SIMULIA (in Abaqus), Sunway, and PGI.

SuperLU and PETSc – Successful tools for solving sparse matrix problems

Sparse linear systems are at the core of many computational science and engineering problems. High fidelity simulations often encompass multi-physics components in multiple spatial and time scales. Unlike dense matrix solvers, sparse matrix solvers are used to tackle problems in which many of the values on the grid are zero. Although it sounds trivial, solving for zero times zero or zero divided by zero over thousands of grid points can dramatically slow down a computer's performance or even cause the algorithm to "blow up" by solving too many equations in which the answer is "not a number."

The ASCR program has invested heavily in the research and development of sparse solver libraries. In particular, the SuperLU and PETSc (Portable, Extensible Toolkit for Scientific computation) projects are critical software libraries that enable robust, scalable solutions for thousands of applications users. These open source libraries are a resource for users throughout the world and represent the state of the art in parallel algorithms research and library software implementation.

The SuperLU libraries are sparse direct solvers and preconditioners that can solve the most challenging linear systems that can overwhelm iterative methods (the name refers to the supernodal technique used in the algorithm and the LU matrix decomposition method used to solve linear systems). First released in 1997, SuperLU is downloaded more than 30,000 times a year and has been included in many vendors' mathematical libraries, including Cray LibSci, FEMLAB, HP's MathLib, IMSL, NAG, OptimaNumerics, Python (SciPy), and Interactive Supercomputing. SuperLU has been adopted in industrial mathematical libraries and simulation software, including AMD (circuit simulation), Boeing (aircraft design), Chevron and ExxonMobil (geology), and Walt Disney Animation.

PETSc is a software package created for the scalable solution of large-scale numerical systems arising in complex simulations in climate, combustion, fusion, and geoscience research. PETSc enables scientists and engineers to focus on their primary scientific interests while having access to state-of-the-art solvers, which reduces implementation costs and achieves faster and better results. Solving massive algebraic systems for more than three billion unknowns, PETSc allows far more accurate simulations than ever before and is now used throughout the DOE science and engineering research complex, as well as by industry.

PETSc solvers were cited as one of the Top Ten Computational Science Accomplishments of DOE in 2008. PETSc won an R&D 100 Award in 2009 and the SIAM/ACM Prize in Computational Science and Engineering in 2015.

PETSc has been widely implemented and supported by vendors. Commercial simulation packages, including FIDAP 8.5, TigerCAD, and RF3P, use PETSc for their algebraic solvers. Also, products from the companies Tech-X and Actel (now Microsemi) have used PETSc. Cray has provided and supported PETSc with their products and Microsoft has distributed previous versions of PETSc with its HPC software suite, the Microsoft Cluster CD. PETSc has been used by Boeing for computational fluid dynamics simulations and Shell for solving inverse problems for oil reservoir management. Among the research areas where PETSc has been used are acoustics, aerodynamics, air pollution, arterial flow, bone fractures, brain surgery, cancer surgery, cancer treatment, cardiology, cell simulation, combustion, concrete modeling, corrosion, dentistry, seismology, economics, linguistics, magnetic films, medical imaging, ocean dynamics, oil recovery, polymer injection molding, polymeric membranes, semiconductors, rocketry, and surface water flow.

ASCR is supporting the continued development of both SuperLU and PETSc by through support in the software technology focus area under DOE's Exascale Computing Project.

ASCR Computing Facilities Help Drive Industrial Innovation

The most prominent features of ASCR's HPC (high performance computing) program are its three leadership-class computing centers—the Argonne and Oak Ridge Leadership Computing Facilities (ALCF; OLCF) and NERSC (National Energy Research Supercomputing Center) at LBNL (Lawrence Berkeley National Laboratory). The centers provide users with both leading-edge computing and data storage resources as well as specialized support services, often provided by scientists with knowledge of both science domains and computer science. Through programs like ASCR's INCITE (Innovative and Novel Computational Impact on Theory and Experiment) and HPC4Mfg (HPC for Manufacturing), these resources are made available to companies whose ability to innovate plays a critical role in the nation's economic and technological leadership. The following examples illustrate some of the impacts of ASCR's leadership computing facilities.

Caterpillar and Cummins gain edge through computer modeling and analysis

By collaborating with Argonne National Laboratory (ANL) and software developer Convergent Science, Inc., heavy equipment manufacturing giants Caterpillar and Cummins gained access to cutting-edge computer modeling and analysis tools and expertise that allowed them to achieve major advances in fuel economy and reduce development costs and time-to-market for engines.

ANL, through its Virtual Engine Research Institute and Fuels Initiative (VERIFI), has developed engine models and software for large-scale computer simulations that provide, in virtual space before costly physical production ever begins, a better understanding of how internal combustion engine parameters interact.

As part of a Cooperative Research and Development Agreement (CRADA), Caterpillar, ANL, and Convergent Science, Inc., conducted simulations of the 580-horsepower Caterpillar C15 diesel engine to identify the bottlenecks and gaps for engine simulations on HPC platforms. Such simulations reduce




Figure X. The random nature of fuel spray leads to significant cycle-to-cycle and fuel consumption variations in engines. High-resolution computer simulations conducted on Argonne's supercomputers depict the flow structures that occur during fuel injection. Note the differences in individual fuel spray plumes.

the time and cost of the design cycle for new engines, allow faster adaptation of fuels from new sources, and lead to substantial increases in fuel economy while meeting future emissions standards.

As part of another CRADA, Cummins, ANL, and Convergent Science, Inc., conducted simulations on Cummins engine products, focusing specifically on the fluid dynamics of fuel injectors. ANL's fuel spray modeling capabilities and advanced load-balancing algorithm technologies allow for more reliable analytical calculations, lower

dependence on experimental engines and prototypes, and accelerate new engine concepts and design to market.

Industry and LBNL partner to model paint behavior

The process of painting is one of the biggest energy expenditures in the U.S., accounting for .01 percent of all energy nationwide and about 70 percent of every automotive assembly plant's energy costs. One of the world's largest coatings manufacturers, PPG, is working with ASCR-funded applied mathematicians at Lawrence Berkeley National Laboratory (LBNL) on computational models that will help optimize paint design and application to reduce energy consumption.

As they are being made, cars are painted using a device called an electrostatic rotary bell atomizer. In the atomizer, paint flows to a cup rotating at speeds of 10,000 to 70,000 rpm and then, driven by centrifugal forces, the paint sheets across the inside surface of the cup, atomizing into droplets at the edge. Cars are painted in a specialized spray booth and then go through an oven the length of a football

field where the paint sets—both of these areas must be tightly humidity and temperature controlled and thus use huge amounts of energy.

Finding a way to reduce the time case spend in the painting process could reduce energy use. PPG sought to understand the key mechanisms driving the atomization process so paint flow could be sped up without reducing quality.




Figure X: High-resolution petascale modeling of liquid atomization, showing a thin film of liquid (blue) traveling up the inside surface of the cup (the modeling domain is a small wedge), over the beveled edge of the bell, and then detaching into a sheet that subsequently breaks up into thousands of droplets. The side-on view shows how the droplets are pushed up by shaping air currents.

PPG and LBNL are developing an advanced, tailor-made HPC numerical PETSc framework to study rotary bell atomization at industrially relevant conditions, a process that will require millions of CPU hours on NERSC systems, as well as other ASCR supercomputing facilities. The goal is to be able to analyze and predict atomization behavior as a function of fluid parameters, such as density, viscosity, and surface tension. The model will then be used to discover the key mechanisms driving the atomization process, with the goal of optimizing for higher paint flow rates in order to speed up the assembly line, reduce energy use, increase process quality, and reduce defects. Following the success of the first year of the project, DOE’s HPC4Mfg extended funding for the collaboration for an additional year.

Computing cleaner, quieter jet engines

When General Electric (GE) and Pratt and Whitney researchers were separately looking for ways to create more energy efficient, lower-noise jet engines, ASCR's INCITE (Innovative and Novel Computational Impact on Theory and Experiment) program provided them access to the ALCF's high performance computers to perform virtual prototyping. Physical tests of complex systems such as jet engine combustion are difficult, expensive, and time consuming. High-resolution, 3D virtual tests enable the companies to more rapidly predict design performance and tweak designs on the fly before building working prototypes.


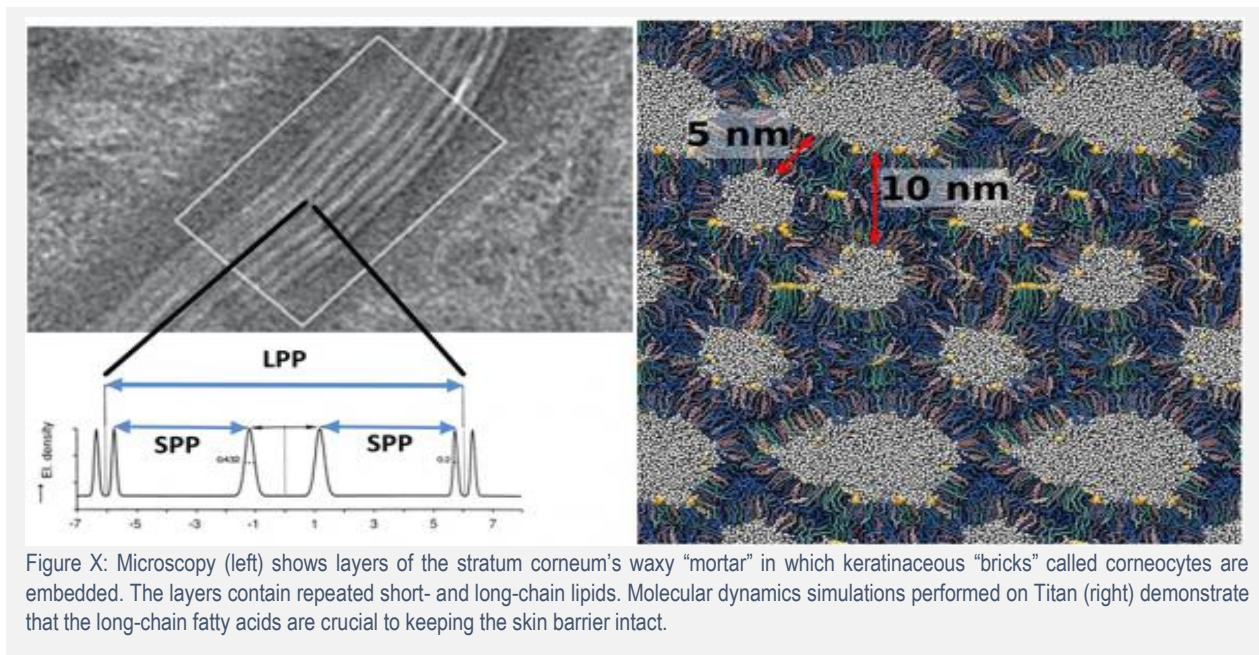


Figure X: Jet engine simulations conducted at the ALCF have been integrated into engine design and the resulting improvements are now in use in aircraft.

In the case of jet engines, a combustor combines air flowing faster than a hurricane with swirling fuel to generate the extraordinary release of heat that ultimately powers the aircraft. Using ALCF's Mira and later Blue Gene/P supercomputers, GE researchers predicted the jet engine noise of different designs of exhaust nozzle shape by modeling the turbulent flow features, and measuring quantities such as velocity, temperature, and pressure that create jet engine noise. Similarly, Pratt and Whitney researchers reduced the solution times for their 3D combustor simulations and provided key insights in the design of their next-generation engines. More fuel-efficient, lower noise next-generation jet engines are aiding these companies, and their customers, to reach greater heights in the extremely competitive aircraft and jet engine global market.

Going with a better flow in consumer products

Many of Procter & Gamble's (P&G's) products for the body and the laundry comprise systems of lipids that dictate the flow, thickness, performance, and stability of the products. In 2011, P&G researchers, with colleagues at Temple University, used Oak Ridge National Laboratory's Jaguar supercomputer to model and simulate microstructural arrangements of lipid vesicles called liposomes. Liposomes encapsulate and suspend ingredients in manufactured products but also can fuse with each other and threaten product performance.




To speed discoveries of why vesicle systems have stability problems the researchers needed to simulate the reorganization of vesicles over time. Before vesicles fuse their surfaces strain and break, exposing water-avoiding parts of molecules to an aqueous environment. P&G researchers hoped such simulations would help to predict performance attributes before a product is manufactured, to enable delivery of high-quality products faster than competitors and earn greater revenues in the marketplace. But the magnitude of the challenge exceeded the capabilities of even P&G's powerful supercomputers.

Using 69 million core hours on the ASCR-funded Jaguar supercomputer over two years, P&G researchers were able to set the stage for simulations of large, complex systems of lipid assemblies. Until then, no one had ever been able to simulate vesicle fusion in realistic systems, such as biological vesicles whose fusion is triggered by proteins.

Squeezing out energy savings in the paper industry

The papermaking industry ranks third among the country's largest energy users, behind only petroleum refining and chemical production, according to the U.S. Energy Information Administration. To

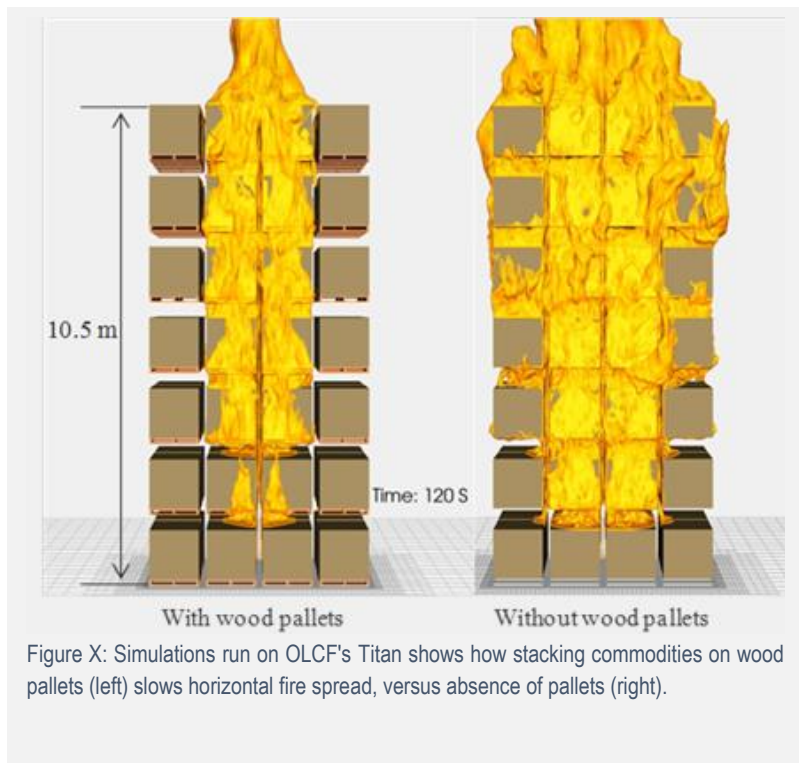
address this issue, researchers from Lawrence Livermore National Laboratory (LLNL) and LBNL used advanced modeling techniques to identify ways that paper manufacturers could reduce energy and water consumption. The first phase of the project targeted “wet pressing”—an energy-intensive process



in which water is removed by mechanical pressure from the wood pulp into press felts that help absorb water from the system like a sponge before it is sent through a drying process. If manufacturers could increase the paper’s dryness by 10-15 percent, it would save paper manufacturers up to 20 percent of the energy used in the drying stage—up to 80 trillion BTUs (thermal energy units) per year—and as much as \$400 million for the industry annually.

The researchers used a computer simulation framework developed at LLNL that integrates mechanical deformation and two-phase flow models. To model the complex pore structures in the pore felts, the team used a flow and transport solver in complex geometries developed by LBNL and based on the Chombo software libraries developed with support from ASCR’s SciDAC programs. The simulations were then run on the Cori supercomputer at NERSC.

The project was one of the seedlings for DOE’s HPC4Mfg initiative, a multi-lab effort to use HPC to address complex challenges in U.S. manufacturing. Through HPC4Mfg, LBNL, and LLNL partnered with the Agenda 2020 Technology Alliance, a group of paper manufacturing companies with a roadmap to reduce their energy use by 20 percent by 2020.



Industrial fire, flood simulations help save lives, money

Insurer FM Global scientists wanted to simulate fires in their clients new, largest mega warehouses. INCITE time allocations allowed them to scale-up their in-house FireFOAM code on OLCF’s Titan supercomputer. FireFOAM is FM Global’s flagship code for simulating the complex physics that occur during an industrial fire in fine temporal and spatial resolution.

An OLCF combustion computational scientist provided expertise on how to efficiently

scale-up FireFOAM on Titan to simulate a seven-tier, 35-foot high, storage area—something impossible on FM Global’s in-house HPC system. FM Global scientists discovered that stacking storage boxes on wooden pallets slows horizontal flame spread, enabling them to offer better fire protection for clients, a third of all Fortune 1000 companies, saving both insured and insurer tens-of-millions of dollars. The results from this first-of-its-kind virtual fire modeling have enabled FM Global, one of the largest property and casualty insurers in the world, to better protect clients.

Similarly, environmental risk start-up KatRisk used Titan to scale up the application of its flood risk software and create the first high-resolution 10-by-10 meter flood risk maps for the U.S. and 90-by-90 meter maps worldwide. The supercomputer modeling helped establish KatRisk as one of the nation’s leading risk catastrophe modeling companies providing insurance companies and U.S. public agencies with critical information for analyzing flood risk at the scale of city blocks.

Geo mapping software opening up new energy fields

In supporting companies’ global competitiveness with leadership class HPC, sometimes it is not just the raw computing power that counts, but also the advanced software. ASCR is the world leader in

developing such software and sharing it with industry to provide HPC-leveraged solutions to a wide range of commercial challenges.

Traditional seismic imaging used in sub-sea oil and gas exploration could not distinguish between a reservoir of oil, gas, water, or brine. This inability was a costly problem as a single failed drill could cost up to \$100 million and six months of lost time. Now, the Electromagnetic Geological

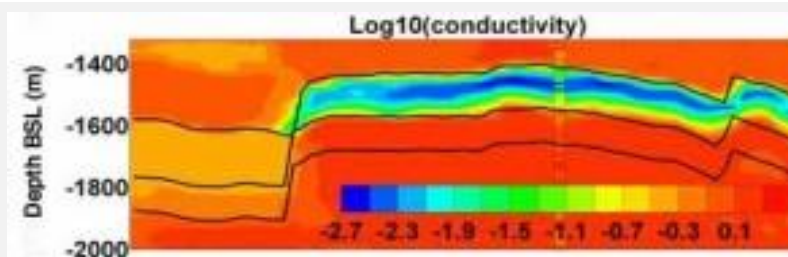


Figure X: Plots of electrical conductivity over the Troll Field in the North Sea produced by analyzing 3D electromagnetic field data using EMGeo software.

Mapper (EMGeo) software developed by researchers at LBNL enable the world's largest oil companies, including ExxonMobil, ConocoPhillips, and Chevron, to overcome this issue by pinpointing hydrocarbons with unparalleled accuracy. The NERSC supercomputers were used to develop EMGeo by interpreting seismic data in combination with extremely large data sets of electromagnetic measurements collected at the same time. This produces a clearer subsurface view at a scale and resolution previously unavailable. The resulting EMGeo software, which can be run on smaller in-house computers, has been licensed to more than 10 companies and is reducing the environmental impact of unnecessary drilling while boosting companies' bottom lines.

Helping Small Businesses Innovate

The Small Business Innovation Research (SBIR) program is a highly competitive program that encourages domestic small businesses to conduct R&D with potential for commercialization. Through an awards-based program, SBIR helps small businesses explore their technological potential and provides the incentive to profit from its commercialization. Bringing qualified small businesses into the nation's R&D arena helps stimulate innovation and encourage an entrepreneurial spirit to meet the nation's specific R&D needs.

DOE is one of nine cabinet-level departments supporting the SBIR program, along with the Environmental Protection Agency, NASA (National Aeronautic & Space Administration), and the National Science Foundation. Within the Office of Science, 13 program offices, including ASCR, work collaboratively with SBIR.

Tech-X helping drive scientific modeling and simulation

Ten years after John Cary and Svetlana Shasharina founded Tech-X Corp., with help from a DOE SBIR grant, the company was front-page news. Nature magazine featured an image from an experiment in creating laser wakefield acceleration, which offers the possibility of a compact, high-energy accelerator for probing the subatomic world, for studying new materials and new technologies, and for medical applications.

Scientists at LBNL guided and controlled extremely intense laser beams over greater distances than ever before to produce high-quality, energetic electron beams. The experimental results were then analyzed by running the VORPAL plasma simulation code developed by Tech-X for use on supercomputers at DOE's NERSC. The results were published in the September 30, 2004 issue of Nature, and the cover image was of simulations created using the VORPAL code. This modeling allowed the scientists to see the details of the evolution of the experiment.



Figure X: Geddes Guided Accel Nature 2004

“With VORPAL, one can see the laser pulse breakup and the injection of particles into the laser-plasma accelerator,” said Cary, CEO of Tech-X Corp. and a professor of physics at the University of Colorado. “This allows one to understand how the injection and acceleration occur in detail to help the experiment’s designers figure out how to optimize the process.”

As 2019 came to a close, Cary reflected on his company’s 25 years of accomplishments. Success, he said, is measured on three counts. First, the Boulder, Colorado company has grown to support about 40 full- and part-time employees and consultants. Second, Tech-X has contributed to the knowledge base and shared that through collaborations and presentations and through the natural process of employee relocation. Last, the company has developed a high level

of technical excellence and developed software that both DOE researchers and companies want to use. Tech-X is an example of successfully transferring its technology to the private sector.

“The SBIR programs has been extremely important to our success, helping us both contribute to DOE’s science mission and compete in the commercial world,” Cary said, adding that Tech-X has continued to receive SBIR grants over the years. “The program is a way for us to contribute to physics research and high-end computing.”

Tech-X develops a variety of scientific software, including physics simulation software and HPC in the cloud, and provides expert physics simulation consulting services in plasma and beam physics and electromagnetics, including photonic integrated circuits.

Cary, who earned his Ph.D. at University of California Berkeley and did his thesis work at LBNL, said his positions at the University of Colorado and Tech-X are another way to advance research. Being a professor allows him to mentor students and post-docs at the beginning of their careers, while at Tech-X, he co-mentors others in the next stages of their careers. “We help them become leaders themselves by developing projects and having a broader impact, whether they remain at Tech-X or move on to another organization.”

Voltaiq: Making the development of next-generation batteries more efficient

Over the past 20 years, batteries have taken on an increasingly important role in our daily lives. They drive our cars, power our phones and tablets, and provide storage for electricity generated from

renewable resources. With this growing importance has come an accelerating effort to improve performance while driving down the costs of next-generation batteries.

Developing new batteries is engineering-intensive, requiring rapid design iterations followed by months or years of testing to identify the optimal material configurations, size, and operating parameters that will meet the specified application requirements. Computation,

both for developing new battery chemistries and for analyzing the performance of new designs, has helped power this research. The increasing volumes of data generated during battery development have made data collection and management disproportionately time-consuming, limiting advanced analysis and the key insight needed to get new products to market.



With the goal of advancing battery development through the power of big data analytics, two ARPA-E (Advanced Research Projects Agency–Energy) funded engineers founded Voltaiq, a LBNL start-up. In 2013, an SBIR grant from ASCR helped Voltaiq founders develop a software environment to simplify and streamline the process of battery test data collection, visualization, and analysis. Collaborating with researchers at four national laboratories, Voltaiq successfully deployed its software platform, with initial customers realizing up to a 70 percent reduction in development time.

The success of this initial product, Voltaiq Core, spurred the development and distribution of three additional products, Voltaiq Analytics, Voltaiq Notebook, and Voltaiq Reports. The list of Voltaiq customers now includes three of the largest consumer electronics companies in the world, global automotive original equipment manufacturers in Detroit and Germany, and dozens of startups and leading university laboratories that are developing the next generation of batteries.

“The SBIR grant was instrumental in helping to get us going,” said Voltaiq CEO Tal Sholklapper. “We were working nights and weekends and the ASCR SBIR funds really helped provide a runway to commercialize our software.”

Power Info, LLC: -- Data-driven visualization of large power grids

In order to keep the nation’s power grid operating, system operators and analysts must make quick and accurate decisions as they perform mission-critical tasks. Electric utilities and grid coordination organizations are looking to advanced tools to help operators better perform these tasks. Visual analytics holds tremendous promise in this area.

Traditionally, visualization of power grids has relied heavily on human designers. Building and maintaining the visualization displays is both labor-intensive, error-prone, and restricted to a limited number of pre-defined patterns.



Figure X: Power transmission lines.

To overcome these shortcomings, researchers at Power Info, LLC developed a data-driven approach for visualizing large power grids. Power Info's data-driven visualization algorithm uses empirically or mathematically derived data to formulate visualizations on the fly. The resulting visual presentations emphasize what the data is rather than how the data should be presented, thus fostering comprehension and discovery.

The software tool resulting from this research is now being leveraged by more than 70 utility organizations in North America and Europe. License sales tripled in the first half of 2018 compared to the same period the previous year. Dominion Virginia Power created a custom solution from Power Info's tool which auto-generates a large number of high-quality visualization displays at a fraction of the traditional cost. Work that used to take more than six months for two full-time employees was completed by a college intern within a month and completely eliminated errors.

With over three years of ASCR SBIR support Power Info, LLC has seen its revenue increase by a factor of almost seven, growing from \$60,000 per year to \$400,000.

4.0 ASCR@40: BROADER ACHIEVEMENTS AND CONTRIBUTIONS

The bulk of this report describes the impact of the Office of Advanced Scientific Computing Research's (ASCR) 40 years of “products”—i.e., the research breakthroughs and facility accomplishments highlighted in section 2 and elsewhere in this report that ASCR has funded and facilitated. In addition, ASCR's “process” achievements have not only established the office as a global leader in the field, but has also significantly benefitted the broader research enterprise. In particular, ASCR's efforts in strategic planning, community standards, and research modalities have advanced the frontiers of research within ASCR's portfolio and beyond.

ASCR and its advisory committee, Advanced Scientific Computing Advisory Committee (ASCAC), have a long tradition of community-driven strategic planning. This planning is the result of an open, transparent process with an emphasis on community engagement and independent leadership. ASCR not only benefits from this thinking in setting its own agenda but also provides guidance for the broader community through reports that are broadly available and used by others, domestically and internationally. These reports articulate emerging challenges and set future research directions to address those challenges.

ASCR has similarly asserted its leadership by articulating and championing community standards and approaches. Using its convening authority to drive consensus standards and community directions (e.g., in partnerships with vendors and professional societies), ASCR facilitates convergence and efficiency within its own portfolio while also showing the way for other players.

Finally, ASCR has championed and used a diverse set of research modalities. These modalities emphasize a balanced and evolving portfolio of fundamental, discovery-driven, principal-investigator research, collaborative teaming, and centers. These elements bring together diverse experts, large-scale facilities, and the knowledge to use them well. Again, these modalities have not only served ASCR well but also set an example that other agencies and institutions have emulated.


“Imitation is the sincerest form of flattery...,” (Oscar Wilde) and the extent to which others have followed ASCR's lead reinforces the defining role ASCR has played in shaping community directions. Given the ever-growing scale and importance of simulation and computational research, and the diverse and evolving hardware technology landscape, ASCR needs to redouble its strategic planning and community standards-settings efforts. By defining research directions, setting future standards, and innovating research modalities that accelerate progress ASCR will continue to shape directions for the broader research community for decades to come.

Strategic Planning

ASCR and ASCAC's long tradition of community-based strategic planning relies on workshops and town hall meetings for broad input and independent leadership for objectivity. Summarizing the content and impact of each ASCR workshop report is far beyond the scope and scale of the current document; more than 100 have been published in the last 15 years alone. Here, we highlight several

exemplars that have stood the test of time and span the breadth of ASCR’s portfolio to illustrate ASCR’s strategic planning process and its impacts on the broader research community.

The Lax report, more formally the “Report of the Panel on Large Scale Computing in Science and Engineering,” a study led by Peter Lax and completed in 1982, emphasized the importance of large-scale computing. The report highlighted two key problems: access to supercomputing was inadequate in many disciplines, and supercomputing capacity, including planned future capacity at the time of the report, was inadequate to meet the needs of the research community. The report particularly noted the lack of supercomputing facilities at universities and argued that scientists were not choosing to work on problems that required the largest computing resources. As a result, U.S. leadership in associated fields was being put at risk. Lax’s study played a significant role in motivating and shaping the planning for National Science Foundation (NSF)- funded computing centers that emerged several years after the report was issued.



Similarly, David Brown led a study in 2008 that articulated “Future Directions for DOE Applied Math.” The report highlighted the key role that applied mathematics plays in the ASCR research portfolio and discussed how these approaches could benefit other DOE programs in the Office of Science and the applied programs. Brown’s report challenged the applied mathematics community to leverage its strengths in predictive simulation and modeling to more fully address systems, developing mathematical approaches that embrace the full complexity of these intricate systems. The report also noted that the mathematics for analyzing very large amounts of data, generated by either simulations or experimental observations, required significant development. A decade later, these themes are still motivating the ASCR applied mathematics portfolio and the

broader applied mathematics research community.

Most recently, ASCR championed a series of workshops spanning the exascale needs of each of the Office of Science’s directorates and the associated development of seven exascale requirement reports.

The capstone of this effort, published in 2017, was the “Exascale Requirements Reviews Crosscut Report.” This report, and the efforts that underpinned it, identified mission-critical scientific problems within the U.S. Department of Energy’s (DOE) Office of Science, including its experimental facilities, and determined the requirements of the exascale ecosystem needed to meet these challenges. Importantly, those requirements were not limited solely to facilities but spanned computing, software, data, and workforce development needs. Given how new this report is, it is premature to assess its long-term impact. However, the fact that ASCR partnered with domain-science communities to articulate challenges and set requirements for the exascale era should benefit ASCR and the broader research community for years to come.

Community Standards

Throughout its entire 40-year history, ASCR has played a key role in setting community directions by asserting its convening authority. Since the 1970s, ASCR’s predecessor organizations articulated the value of, and incentivized, the open source software. In the 1980s, with the emergence of parallel computing, ASCR recognized the need for a distributed-memory programming model, then drove and helped support convergence towards the Message Passing Interface (MPI) as that standard. Since the 1990s, ASCR has championed the existence of a community metric used to track facility performance, the TOP500 list, which is based on the LINPACK Benchmark (developed with DOE support). ASCR-supported research underpins each of these efforts. ASCR was able to influence the community beyond its own cohort of researchers and facilities only by combining research outcomes with convening authority.

Open Source

The U.S. federal government has a history of advancing innovations and discoveries that were simply too big or too risky for any single private enterprise to undertake. Understanding the importance of infusing critical research and development operations with resources, federal departments like DOE have jump-started everything from clean energy to scientific computing by providing funding in the form of competitive grants. This first critical boost often makes the difference between success or failure for ideas and technologies that private industry and entrepreneurs can then pick up and develop into hardened standards and marketable products. This is especially true in the rapidly advancing fields of high performance computing (HPC) and computational science. Many of the scientific applications, software libraries, and software standards, created under these federal grants, are open source and can be rapidly adopted and deployed by the community at large. The collection of software is fundamentally important as it addresses just about every key computing need—compilers, package managers, numerical methods, data analysis tools, visualization tools, input/output benchmarking, data storage, parallel performance, workload management, math and physics codes, and more. Given the breadth of this kind of software and software development, it makes sense to employ a community-based, open source approach when developing and maintaining such a large and strategically important software stack.

Open source as a software strategy

The open source approach to software development engages a community of interested users and developers, including government agencies and private industry, in a collaborative effort in which anyone can improve upon the code and share the changes with the community. Whenever someone in the community makes an improvement, the whole community benefits, and high usage combined with external contributions make the results incredibly stable and widely relevant to many computer platforms, programming languages, and simulation frameworks. This enables the development and maintenance of a high-quality product at a low cost to individual users. With open source software, agencies and contractors can align their goals and make the most of external resources by sharing projects and avoiding redundancies. This not only reduces barriers to collaboration and reduces duplication of effort, but it helps establish the reputation of individuals, organizations, and entire nations as software leaders in a field. A huge part of the reason ASCR continues to drive the global vision of HPC is its dedication to pursuing open source solutions.

Open source roots at DOE

The open sourcing of DOE software began in the late 1970s. To provide even broader access to this software, DOE began creating software repositories.

One such repository, the Argonne Code Center, was established by Margaret Butler at Argonne National Laboratory (ANL). This effort later became the National Energy Software Center. In 1991, the center moved to the Office of Scientific and Technical Information (OSTI) headquarters in Oak Ridge, Tennessee, and was renamed as the Energy Science and Technology Software Center. The latest iteration of this effort, still housed at OSTI, is now known as DOE CODE. This repository allows authorized users to search through DOE's software portfolio and to submit their own projects.

Netlib is another early repository of open source software. Netlib repositories contain collections of mathematical software, papers, and databases and were initially established in 1984 by Eric Grosse at Bell Labs and Jack Dongarra at ANL and partially funded by ASCR, based on a suggestion from Gene Golub, a professor at Stanford University. Each element of the repository combines to make a collection of high-quality mathematical software available via electronic mail. The Netlib collection quickly spread to mirror servers around the world and as the Internet became ubiquitous other access methods were added.

Projects like the Argonne Code Center and Netlib have succeeded through volunteerism, with encouragement by colleagues and goodwill from employers. Advancing to the next stage of growth for computational simulation and modeling will require the solution of hard, basic research problems in computer science and applied mathematics as well as creating and promulgating a new paradigm for the development of scientific software. The obvious questions are how should this kind of activity be supported, and where should it be done (e.g., universities, national laboratories, software companies)? Funding agencies often take the view that such efforts are not research, yet the results significantly

benefit the research community. With declining industrial research support for broad ecosystem efforts future progress will likely require a greater level of sustained funding from government sources.

Message Passing Libraries and MPI

The 1980s saw distributed-memory parallel computers and the emergence of networks of workstations as a resource for parallel computation. With these developments the need for parallel programs that used a distributed-memory programming model emerged. In this approach, processes are distributed across the nodes of the distributed-memory parallel computer. Each process has direct read/write access only to its own local memory. This means that whenever data residing in the memory of another process is needed, the data must first be communicated and brought into the local memory. Such a programming model is known as a message-passing model. The exchange of messages requires a communication, or message-passing, library (i.e., a set of functions that perform the needed communication operations). These libraries have well-defined syntax and semantics and an efficient implementation of those functions that exploits the capabilities of the networking hardware and minimizes the overhead associated with communication.

Early efforts

When distributed-memory parallel computers initially appeared, there was no standard or well-defined way of doing message-passing in a parallel program. As a result, each vendor of parallel systems created its own message-passing library with its own API (application programming interface) which was available only on that vendor's systems. Examples include NX from Intel, MPL from IBM, CMMD from Thinking Machines, and VERTEX from nCUBE. As a result, programs written for one vendor's systems could not be run on another vendor's system. If a particular vendor went out of business, application developers were left with code that could not run elsewhere unless additional effort was made to port it for use in another vendor's library.

The newly emerging field of distributed-memory parallel programming also attracted the interest of many researchers in parallel programming models. Their efforts focused on two main areas: the development of portable message-passing libraries that provided a single API and could run on multiple platforms, thereby overcoming the problem of lack of portability in the vendor libraries; and the research question of what functionality and features a message-passing library must provide to meet application needs, provide ease of programming, and (importantly) achieve the highest performance. Prominent research libraries of the time included Parallel Virtual Machine (PVM), p4, Chameleon, TCGMSG, Zipcode, and others. PVM, p4, and Chameleon were all funded, or partially funded, by ASCR.

PVM, developed at Oak Ridge National Laboratory (ORNL) and funded by ASCR, the University of Tennessee, and Emory University, was among the early parallel programming environments that grew to be enormously popular on a global scale. PVM worked on all platforms, all parallel computers in existence at the time, on clusters of workstations, and also on a heterogeneous collection of systems of

different types. PVM made distributed-memory parallel programming widely available and accessible to people all over the world and led to the development of many parallel applications.

Three distinct message passing efforts were developed at ANL under ASCR funding. P4 was a parallel programming library that included both message passing and shared-memory components and was portable to many parallel computing environments, including heterogeneous systems. Chameleon was a high performance portability package for message passing on parallel supercomputers implemented as a thin layer over vendor message-passing systems (e.g., Intel's NX, TMC's CMMD, IBM's MPL) for performance and over freely available systems (e.g., p4 and PVM) for portability. TCGMSG was another message-passing library, developed by the Theoretical Chemistry Group at ANL, intended for use in computational chemistry applications.

MPI Forum and the MPI standard

By 1992, there were many message-passing libraries in existence and also significant user experience in developing applications with these libraries. However, the lack of a single, standard message-passing API that was portable across all platforms and enabled high-performance implementations, was hindering the development of portable application codes. The HPC community, including application developers, vendors, and researchers, was determined to address this problem.

In April 1992, the National Science Foundation's Center for Research in Parallel Computation, led by Rice University, sponsored a one-day workshop on "Standards for Message Passing in a Distributed-Memory Environment." The result of that workshop, which featured presentations of many message-passing systems, was a realization that a great diversity of good ideas existed among message-passing systems, and that people were eager to cooperate on the definition of a message-passing standard. At the Supercomputing '92 conference in November 1992, a committee was formed to define a message-passing standard with the following goals:

- To define a portable standard for message passing; not an official, ISO-like standard but attractive to both implementers and users;
- To operate in a completely open way, allowing anyone to join the discussions, either by attending meetings in person or by email; and
- To finish in one year.

This international committee was led by Jack Dongarra, Bill Gropp, Rusty Lusk, and David Walker who were funded by the ASCR program. The MPI (message passing interface) effort was lively as a result of the tensions among these three goals. The MPI standardization attracted a wide class of vendors and users because the MPI Forum itself was so broadly based. At the original (MPI-1) forum, the parallel computer vendors were represented by Convex, Cray, IBM, Intel, Meiko, nCUBE, NEC, and Thinking Machines, as members of the groups associated with the portable message-passing libraries (PVM, p4, Chameleon, Zipcode, PARMACS, TCGMSG, and Express) and a number of parallel application developers. Equally important was an early commitment to producing a reference implementation, MPICH, by a group from ANL and funded by ASCR. MPICH helped demonstrate

that the implementation of MPI on multiple platforms was feasible. MPI absorbed the best ideas from existing systems and broke new ground with novel important concepts to aid scalability and portability.

The first version of the MPI standard (MPI-1) was completed in May 1994. The release of the MPI standard combined with the availability of the MPICH implementation that worked on all platforms led to the quick adoption of MPI by both users and vendors. All vendor systems released after 1994 supported MPI, in many cases using MPICH as a basis for the vendor-derived implementation, and application developers began using MPI in their codes. As a result, MPI gained in popularity quickly and soon became ubiquitous.

The Forum reconvened during 1995–1997 to extend MPI to include additional functionality, such as remote memory operations, parallel Input/Output, dynamic process management, and a number of features designed to increase the convenience and robustness of MPI. The resulting standard (MPI-2) was released in 1997 and remained the definition of MPI for the next eleven years. A third Forum was constituted in 2008, again consisting of vendors, computer scientists, and computational scientists. During 2008–2009, the Forum updated the MPI-2 functions to reflect recent developments and to correct errors, culminating in the release of MPI 2.1 (2008) and MPI 2.2 (2009). The Forum continued to meet, substantially extending MPI with new operations (nonblocking collectives, neighborhood collectives, improved one-sided communication, etc.), releasing the MPI-3 Standard in 2012. MPI 3.1, with minor updates to MPI-3, was released in 2015 and is the current official release of the MPI Standard. The Forum continues to meet to define new features that will be part of the next major version, MPI-4.

MPI Implementations: MPICH

The MPICH project, supported by the ASCR program began in 1992 with the goal of providing a portable implementation of MPI. The idea was to provide early feedback on decisions being made by the MPI Forum and provide an early implementation allowing users to experiment with the definitions even as they were being developed. Targets for the implementation were to include all systems capable of supporting the message-passing model. MPICH is a freely available, complete implementation of the MPI specification, designed to be both portable and efficient, and is both a research project and a software development project. As a research project, its goal is to explore methods for maximizing the communication performance achievable by a parallel program on a given hardware architecture. MPICH has resulted in a number of research publications, including 150 from ANL staff alone. As a software project, MPICH's goal is to promote the adoption of the MPI Standard by providing users with a free, high performance implementation on a multitude of platforms while aiding vendors in providing their own customized implementations. MPICH won an R&D100 Award in 2005.

MPICH is widely used and many vendors use MPICH as the basis for their own vendor-tuned MPI implementations (e.g., Intel, Cray, IBM [Blue Gene series], Microsoft). MPICH-based implementations are routinely used in the majority of the top 10 systems in the TOP500 list. For example, in the June 2018 edition of the TOP500 list, eight of the top 10 systems used MPICH-based implementations. MPICH continues to evolve to exascale with support from DOE's Exascale Computing Project (ECP). The first exascale system (Aurora), to be deployed at ANL in 2021, will use

an MPICH-based MPI implementation with close cooperation and co-development between the MPICH group at ANL and the system vendor, Intel.

MPI implementations: Open MPI

In addition to MPICH, there were other MPI implementations in the late 1990s and early 2000s, such as LAM/MPI from Indiana University, FT-MPI from University of Tennessee, and LA-MPI from Los Alamos National Laboratory. Around 2004, these groups decided to abandon their separate efforts and instead work together on a single, new, open-source MPI implementation called Open MPI, with a number of participants funded by ASCR. Since then, Open MPI has grown into a major effort with contributions from many others around the world, including Oak Ridge and Sandia national laboratories. Open MPI has similar goals as MPICH in terms of portability and high performance. It has a strong presence in the cluster computing world, particularly on InfiniBand clusters. A recent—and perhaps the most high-profile—example of an Open MPI implementation is IBM’s Spectrum MPI being used on the Summit system at ORNL and on the Sierra system at Lawrence Livermore National Laboratory. The K Computer in Japan also uses an Open MPI-based implementation. Like MPICH, Open MPI is supported by ECP.

TOP500

The TOP500 list of supercomputers has served as a defining measure for supercomputing for over 25 years. The collected data have enabled clear, and often early, detection and analysis of important technological and architectural trends. Its simplicity has, over time, invited many critics but has also allowed it to stay useful and useable during the advent and age of giga-, tera-, and peta-scale computing.



The TOP500 project has been tracking supercomputer installations since 1993, publishing a list of the fastest systems and some of their key characteristics twice a year. Systems are ranked according to their performance using the High Performance LINPACK (HPL) benchmark, funded in part by ASCR, which solves a linear system of equations.

Early on, tracking supercomputer installations was somewhat arduous, with ill-defined bounds. In the mid-1980s, Hans Meuer, a professor at the University of Mannheim, and Erich Strohmaier, then a graduate student and now at LBNL, started collecting statistics about the numbers, locations, and manufacturers of supercomputers worldwide. They collected these data from vendors and colleagues in academia and research centers. Initially, it was relatively obvious which systems should be considered “supercomputers.” This label was reserved for vector processing systems from companies like Cray, CDC, Fujitsu, NEC, and Hitachi, all of which competed in the same market—with each claiming they had the fastest system for scientific computation by some selective measure. By the end

of the 1980s, the situation became increasingly complicated as smaller vector systems became available from some of these vendors and new competitors (e.g., Convex, IBM) and as massively parallel systems (MPP) with single instruction, multiple data (SIMD) architectures (Thinking Machines, MasPar) or multiple instruction, multiple data (MIMD) systems based on scalar processors (Intel, nCube, and others) entered the market. Simply counting the installation base for these systems of vastly different scales did not produce any meaningful data about the market. A new criterion was needed. Over two years of experimentation and looking at various metrics and other long-term approaches, ranking according to the actual performance achieved emerged as the new criterion.

The simplest and most universal ranking metric for scientific computing is floating-point operations per second. More specialized metrics (e.g., time to solution or time per iteration and per grid-point) can be more meaningful for their specific application domain and allow for more detailed comparisons (e.g., between alternative algorithms with different algorithmic complexities). They are, however, harder to define properly, more restricted in their use, and, due to their specialization, not applicable for the overall market. Similarly, it is imperative to use actual measured performance values to avoid contaminating the results by then-common outlandish, and often unsubstantiated, claims about estimated performance for systems which often did not reliably function or even exist.

A decision was made by the TOP500 organizers—Meuer, Strohmaier, and Dongarra, and subsequently Horst Simon (Dongarra, Simon, and Strohmaier are funded by ASCR)—to select and require the use of a single benchmark for all listings. This benchmark would not be chosen to represent performance of an actual scientific computing application, but should very coarsely embody the main architectural requirements of scientific computing. To encourage participation, a well-performing code that would showcase the capability of systems while not being overly harsh or restrictive was favored. The purpose of using a single benchmark in the TOP500 was never to claim such representativeness but to instead collect reproducible and comparable performance numbers.

An evaluation of the benchmarks suitable for supercomputing in the early 1990s found that the LINPACK benchmark had the most documented results by a large margin and therefore allowed immediate ranking of most of the systems of interest. The LINPACK benchmark, developed by Dongarra and three co-authors, solves a dense system of linear equations, which is sometimes criticized as an overly simplistic problem. However, it is by no means embarrassingly parallel and worked well with respect to reducing the rankings of loosely coupled architectures, which were of limited use to scientific computing in general. The HPL implementation of LINPACK came with a self-adjustable problem size, which allowed it to be used seamlessly on systems of vastly different sizes. As opposed to many other benchmarks with variable problem sizes, HPL achieves its best performance for large problems that use all of the available memory and not for small problems that fit into the cache. HPL also encodes an important set of optimization parameters and allows substantial performance optimization through their adjustments. All of these features made LINPACK the obvious choice.

Leveraging HPL, the TOP500 collection has become the leading metric for the HPC community. The trends it exposes, the focused optimization efforts it inspires, and the publicity it brings to the community are very important. As we are entering a market with growing diversity and heterogeneity

in hardware architectures, a careful selection of appropriate metrics and benchmarks that match the needs of applications is more necessary than ever. HPL encapsulates some aspects of real applications (such as strong demands for system reliability and stability, for floating point performance, and to some extent for network performance) but no longer adequately tests memory performance. Alternative benchmarks, as a complement to HPL, could provide corrections to individual rankings and improve understanding of systems but are much less likely to change the magnitude of observed technological trends.

The gap between HPL predictions and real application performance is expected to increase in the future, which creates a risk that future architectures that target good HPL performance will not be a good match for scientific applications. With ASCR support, the community has been working towards a new metric with a stronger correlation to the HPC application base that can, therefore, drive system designers in more broadly constructive directions. The High Performance Conjugate Gradient (HPCG) benchmark, a candidate for this new metric, is a simple program that generates a synthetic sparse linear system that is mathematically similar to a finite element, finite volume, or finite difference discretization of a 3D heat diffusion problem on a semi-regular grid. The problem is solved using domain decomposition with an additive Schwarz preconditioned conjugate gradient method, where each subdomain is preconditioned using a symmetric Gauss-Seidel sweep.

The HPCG Benchmark can help alleviate many of the problems described above by providing comprehensive yet minimal coverage of the code types that test major communication and computational patterns, rewarding investment in high performance of collective communication primitives and rewarding investment in local memory system performance.

Research Modalities

Throughout its history, ASCR has championed and utilized a diverse set of funding modalities to accomplish its research objectives, emphasizing a balanced and dynamic portfolio of fundamental, discovery-driven PI research, collaborative teaming and centers that bring together diverse experts for enhanced impact, and large-scale facilities and the expertise to use them well. ASCR has also championed workforce development through efforts like the CSGF (Computational Science Graduate Fellowship) program. These diverse approaches have been a key element of ASCR's success.

As highlighted throughout this report, many of ASCR's, and the computing research community's, largest and most visible successes began with research breakthroughs by individual investigators. These discoveries, spanning things like algorithmic advances, software strategies, and hardware technologies, ultimately led to huge impacts but began as high-risk endeavors by visionary scientists. ASCR's commitment to high-risk, high-reward foundational research, stewarded by technically cognizant program managers, is one of the key ingredients to ASCR's success.

ASCR has also recognized the essential importance of collaborative research and codesign in its many forms. By creating a variety of center-like research modalities, from the SciDAC (Scientific Discovery through Advanced Computation) Centers and Institutes, to the initial exascale codesign centers, to the

present ECP effort, ASCR has both accelerated progress on key questions and enhanced the scale of activity consistent with research priorities. In addition to bringing multiple investigators together, ASCR has fostered cross-discipline collaboration between domain science experts and ASCR-funded investigators. This emphasis on codesign is a key feature of ECP and will play a central role in defining the exascale ecosystem in the years to come.

ASCR's commitment to large-scale computing facilities, NERSC (National Energy Research Scientific Computing Center) at Lawrence Berkeley National Laboratory and the leadership computing facilities at ANL and ORNL, has differentiated DOE's HPC efforts from those of other government agencies. Not only does the scale of these facilities dwarf those of NSF's university-based centers, but the facilities have also made a robust commitment to ensuring user success by providing resident applications and computational science expertise in addition to operational leadership. Further, by operating in an open, peer-reviewed, competitive environment, ASCR has ensured broadly available resources for tackling the most compelling problems.

In addition, ASCR has made a robust and long-standing commitment to workforce development. The CSGF program has benefitted many members of the research community, including some of today's most senior leaders. In addition to the advantages of having CSGF alumni working at DOE labs, others have assumed leadership roles across the full breadth of the computer science ecosystem, from labs to academia to industry, further enhancing ASCR's impact. Another impact has been the adoption by the National Nuclear Security Administration of a program modeled after CSGF. The recent addition of Office of Science Early Career awards has increased ASCR's commitment to workforce development as well as single-investigator innovation.

ASCR's commitment to collaborative centers, large-scale, user-friendly computing facilities, workforce development, and single-investigator research has played a defining role in its success. Other agencies and other countries are emulating these approaches. As computational research continues to expand and diversify, including into areas such as data analytics and artificial intelligence, ASCR will need to continue to innovate collaborative research modalities to ensure its competitiveness and community leadership.

One such model being discussed is the "Superfacility." Increasingly, users of DOE's national user facilities access ASCR's 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. The Superfacility framework would further integrate 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.

Fully realizing the Superfacility capabilities will require significant research, strong partnerships with experimental facilities and innovations in networking, software stacks, computing architectures, and high performance computing facilities policies.

Conclusion

ASCR's "process" achievements have significantly and positively impacted ASCR's success as well as that of the broader research enterprise. In particular, ASCR's efforts in strategic planning, community standards, and research modalities have advanced the frontiers of research both within ASCR's portfolio and beyond. Given the ever-growing scale and importance of simulation and computation research and the diverse and evolving hardware technology landscape, ASCR needs to redouble its strategic planning and community standards-settings efforts—both in defining research directions and future standards and in innovating research modalities that accelerate progress—so that ASCR can continue to shape directions for the broader research community for decades to come.

Mirroring the growth in capabilities for modeling and simulation, ASCR's HPC resources are increasingly playing a critical role in supporting scientific discovery from experiments are the Office of Science's light sources, accelerators, and nanoscience centers, as well as from international facilities. In order to help researchers get the science out of their data, ASCR is supporting projects to advance artificial intelligence and machine learning. Training supercomputers to more quickly and accurately characterize and analyze data is critical to maintaining the nation's scientific leadership. In 2019, ASCR supported the convening of AI Town Hall Meetings at ANL, LBNL, and ORNL, with attendees meeting in a fourth session in Washington, D.C. to present their findings and recommendations as the path(s) forward are mapped out. By combining past process-approaches with the expertise being developed across the Office of Science labs, ASCR can play a leading role in shaping the nature of computational science in the years to come.

5.0 LESSONS LEARNED AND CHALLENGES OF THE FUTURE

It is impossible to collect all of this historical information about the Office of Advanced Scientific Computing Research's (ASCR) impacts without also recognizing some broader lessons that enabled these successes. These lessons and challenges have relevance for ensuring ASCR's continued success.

Lesson 1: A Compelling and Consistent Vision can Drive Scientific Revolutions

ASCR and its predecessor organizations have had a consistent belief that computing was a key driver of science. This belief long predated the broad recognition of the field of computational science. It was instantiated in various ways throughout the decades as technologies and disciplines evolved, but it has always been ASCR's guiding principle.

This sustained commitment drove nearly all of ASCR's investments—the development of advanced mathematical techniques, the evolution (and in some cases revolution) of computer architecture, the creation of state-of-the-art networking capabilities, an array of innovative computer science concepts, the development and support of powerful software libraries, and an interdisciplinary workforce. The integration of all of these capabilities with the U.S. Department of Energy (DOE) applications has driven a scientific revolution. It is nearly inconceivable today that any area of science can push forward without advanced computing.

ASCR, in partnership with the scientists at the DOE labs, displayed admirable vision of a more focused nature at key points in the evolution of technology. Well before it was broadly embraced, ASCR recognized the importance of parallel computing and funded critical work that helped the community navigate the transition to parallelism. ASCR appreciated the emerging importance of data science and the need to make large investments in networking technologies and infrastructure. And ASCR identified the critical role that uncertainty quantification would have to play as computer models were used for critical decision making. In all these areas and others, a clear vision and sustained commitment were essential to community progress.

This clarity was made possible through the hard work of technically knowledgeable program managers in close collaboration with the research community.

Lesson 2: Diverse Funding Models are Required for Diverse and Impactful Outcomes

In its long history, ASCR has employed a wide variety of different funding models—short- versus longer-term, open-ended versus narrowly-targeted, large collaborations versus single investigators, etc. The broad scientific changes ASCR has driven required all of these different funding approaches. Support for graduate and postdoctoral fellowships have helped build the required workforce and attract them to the DOE laboratories. Fundamental mathematical and computer science advances have been made possible by patient, sustained investments in single principal investigators or small teams. Interdisciplinary collaborations have been created by larger, cross-institutional investments.

Software is an essential element of ASCR’s capabilities, and software development and support are best enabled by targeted funding. Competitive processes drive excellence and innovation, but it is also essential to have mechanisms that sustain long-term assets in researchers and software. Networking, HPC (high performance computing) facilities, and HPC platforms require yet another funding model. ASCR’s greatest successes have been enabled by this broad ecosystem of funding modalities.

Lesson 3: Workforce Investments have been Critical



Fellows in the CSGF Program meet annually to present their research and hear from invited speakers. The program requires them to also spend at least one summer in a practicum at a DOE national laboratory.

Nothing in science is possible without the right set of highly skilled people. ASCR’s vision has been ahead of the ability of academia to adapt. When trying to do things that universities were not yet embracing, ASCR had little choice but to invest in workforce development initiatives to meet its needs. The Computational Science Graduate Fellowship (CSGF) program has been the most visible and transparently successful effort in this regard, but the named postdoctoral positions at many of the laboratories have also been hugely impactful (e.g. the Wilkinson Fellowship at ANL, the Householder Fellowship at ORNL, and the von Neumann Fellowship at SNL). In addition, many research grants support students or postdocs, and DOE’s early career programs, have further developed promising young scientists. These investments have been central to staffing the DOE laboratories, but they have also had broader impacts. Scientists with computational skills are now in high demand in industry. And ASCR’s investments have played an important role in the creation of interdisciplinary programs and activities at a number of universities.



Lesson 4: Partnerships are Essential

The complex challenges faced by DOE can only be addressed by interdisciplinary teams that encompass diverse areas of expertise, a model started by Ernest Orlando Lawrence when he established his Radiation Laboratory in Berkeley, California in the 1930s. These scientific partnerships are best enabled by programmatic partnerships. SciDAC (Scientific Discovery through Advanced Computation) is a trail-blazing program that helped overcome organizational barriers in the Office of Science, and at some of the labs, and led to new kinds of science. As the potential impacts of computing continue to broaden, ASCR has embraced additional partnerships. The Exascale Computing Project is a close partnership with the National Nuclear Security Administration. And in appropriate contexts,

interagency partnerships have also been highly impactful. ASCR's willingness and ability to engage in partnerships has been a key element of many of its greatest impacts.

Lesson 5: Testbeds and Platform Access Funding Models are Important

At points of architectural uncertainty or inflections, it is critical to invest in small testbed systems to understand the strengths and weaknesses of different designs to make informed decisions about future directions. Based on the experiences with those systems, it is helpful to get somewhat larger "early access systems" where friendly users can try out full-scale applications and begin adapting them for the next generation of machines. These steps build confidence in the architectures ahead of an acquisition as well as buy-in from the users, and they allow vendors to learn from the experiences of early adopters. This approach has been embodied in ASCR's funding of codesign centers on the path to exascale systems.

As the growth in the scale of the largest systems continues to outpace that of commercial systems, ASCR must continue to invest in the R&D necessary to build and deploy systems with increased scientific capabilities but are also more energy efficient. This will take the form of both pathfinding R&D such as the PathForward program that focuses on developing new technology that can be used in many systems, and the non-recurring engineering associated with the acquisition of a specific system. Both of these types of investments will be needed for all aspects of the systems including hardware, tools, system software, programming environments, and applications.

From time to time, scientific computing centers at universities and national laboratories have tried to finance medium to large-scale systems by charging users for access. This business model has never worked for large-scale, leadership systems. To advance the state of the art in HPC, the funds to purchase the computer must be appropriated and access to the system must be free to users. Of course, any proprietary use of the computers should be reimbursed at a "full cost recovery" rate.

Looking forward, ASCR faces a number of issues due to changes in the scientific and technology landscapes. Lessons learned from past successes can shed useful light on these upcoming challenges.

Challenge 1: Technology Disruptions

The vast majority of performance improvements in supercomputers over the past several decades have come from ever-shrinking microelectronics and faster clock speeds. With the end of Dennard Scaling and the imminent demise of Moore's Law, these performance drivers are coming to an end. New ways must be found to squeeze further performance gains out of the machines. The scientific community is already embracing simple heterogeneity, but future machines will necessarily be much more complex. How will this complexity be managed? How should codes be written today to prepare for more frequent disruptions in the future? How should the role of leadership-class facilities evolve in an era of architectural uncertainty and disruption? How can we prepare for even more disruptive future technologies like quantum computing?

As we enter an era of great change, strategic clarity and vision from ASCR will be essential. Technology disruptions will also require innovative new ideas in mathematics and computer science. ASCR will need to sustain investments in creative individuals and high-risk concepts.

Challenge 2: Funding Balance

Another dramatic shift in the ASCR landscape is the rapid emergence of data science and machine learning in scientific workflows. With finite dollars, how should ASCR balance the need to support new areas with the need to continue supporting areas of historical strength? Again, strategic clarity will be critical to making these hard choices.

Challenge 3: Software Stewardship

The community has long struggled to settle on a good model for sustained support for key elements of the software ecosystem. This issue will get more acute as the Exascale Computing Project winds down and its large, focused software efforts are at risk of being left high and dry with no support for continued development or maintenance. ASCR needs to recognize that software is really a scientific facility that requires sustained investments in maintenance and support. Improved and sustainable funding mechanisms are required.

Closely related is the need for investments in improved software engineering practices reflecting profound changes in the way scientific software is developed and maintained. Modern scientific software is increasingly the product of large, dispersed teams and leverages a diverse suite of libraries and tools. Codes may include machine-specific optimizations and incorporate tools from the machine learning community. All of these trends add complexity to the software and to the development process. This complexity can be managed only through disciplined software engineering processes including thorough documentation, comprehensive regression suites, issue tracking and more. Investments in developing, applying, and advancing best practices in software engineering for scientific applications will be essential for continued progress.

Challenge 4: Broader Partnerships

The SciDAC program is rightly held up as a visionary success for bridging science domains across the Office of Science to build interdisciplinary partnerships and thereby transform science. There is a growing opportunity for simulation and machine learning in other areas of DOE, including the Applied Energy Offices and even some of the more operations-focused organizations. These organizations often have different value systems and funding models from the Office of Science. How can ASCR work to partner with this broader group of entities to maximize impact?

Challenge 5: A Sought-after Workforce

ASCR's workforce investments have played a major role in educating computational scientists and attracting them to DOE laboratories, but generational and technological changes will require fresh thinking about workforce issues. Computationally trained scientists have vastly more opportunities in industry and academia than they did a generation ago, which increases competition for the "best and the brightest." And changes to salaries and benefits have also eroded labs' ability to recruit and retain critical talent. What are the implications for ASCR of a more mobile and transient lab workforce?

Challenge 6: New Roles for Computing to Advance Science

From the 1950s through the 1990s, simulation was the dominant way in which science was enabled by advanced computing. The specific areas of research emphasis have evolved as computing has become more capable, e.g. as simulations became faster and more accurate, scientists have been able to focus on design optimization and uncertainty quantification. But for many decades the primary goal for ASCR and its predecessor organizations was to enable more rapid, detailed and accurate simulations.

In the past two to three decades, in response to scientific needs, ASCR has broadened its activities to support collaborative technologies and data-centric research and development. As described throughout this report, these investments have enabled scientific advances quite different from modeling and simulation. As science continues to evolve, ASCR will need to adapt to, embrace, and support new roles for computing in science.

Two trends are clear today, and others will undoubtedly emerge in the future. First, the explosive growth in the capabilities of machine learning and artificial intelligence (AI) is creating new routes to scientific discovery. ASCR researchers are exploring new computing workflows in which simulation and AI work together to generate insights that neither could produce alone. ASCR will need to nurture the young science of scientific machine learning with both fundamental and applied investments. These new workflows will likely drive new thinking about the design and usage models for advanced computers.

A second trend is the rapidly growing data streams coming from DOE's user facilities like light sources, accelerators, and telescopes. The volume and velocity of these data streams can only be addressed with some form of advanced computing. Increasingly, scientists are seeing value in using HPC to analyze experimental data in real time, and to combine experiments with simulation. As with AI, this trend will likely drive fresh thinking about the nature of advanced computing platforms and how they are integrated with data sources and human decision makers. The Superfacility concept discussed elsewhere in this report is an important exemplar of the path forward.

APPENDIX 1: CHARGE TO ASCAC

Department of Energy
Office of Science
Washington, DC 20585

December 7, 2017

Professor Daniel Reed
Chair, Advanced Scientific Computing Research
Advisory Committee
University of Iowa
Iowa City, Iowa

Dear Professor Reed:

I am writing to ask the Advanced Scientific Computing Advisory Committee to produce a report that assess and document the historical accomplishments of the Advanced Scientific Computing Research (ASCR) program and its predecessors over the past four decades. The report, to be produced during the coming year, should highlight outstanding examples of major scientific accomplishments emerging from ASCR's and its predecessor organizations that have shaped the fields of ASCR research. In addition to selected accomplishments, the report should identify the lessons learned from these examples to motivate ASCR investment strategies for the future.

As history has shown, basic research advances have been the bedrock of American innovation and prosperity. These advances often gave rise to new lines of scientific inquiry and led to inventions of new technologies and industries that transformed our society. Breakthrough discoveries emerging from Federal investment can have broader impacts beyond the original field of scope and have made Federal programs, such as ASCR, an essential part of the Nation's science and technology strategy.

The ASCAC 2007 "Opportunities and challenges of Exascale computing," 2012 "DOE Data-intensive Science and Exascale," and the 2013 "Ten Technical Approaches to Address the Challenges of Exascale Computing" reports have identified critical research opportunities for applied mathematics, computer science, computational partnerships and advanced networking during the exascale era. The upcoming ASCAC report on the opportunities and challenges for future high performance computing capabilities should further identify areas for ASCR investments in the "Beyond Moore's Law" era.

By examining past successes, I expect this report to illuminate the guiding strategies and approaches that will be key to ensuring future U.S. leadership, and more generally, U.S. leadership in the full range of disciplines stewarded by ASCR. Even more broadly, such a report will be timely to inform the future investment strategy for the Office of Science as it contributes to fulfillment of the Department of Energy's missions, especially in view of the Federal budget outlook. With these high-level objectives in mind, the report should provide technical details as needed for context but should be primarily concerned with the essence of each story as it relates to the larger progress of science.

♻️ Printed with soy ink on recycled paper

I ask ASCAC to consider the following questions in formulating the study plan:

- What are the major scientific accomplishments that have shaped the ASCR-supported disciplines in the past 40 years? How has ASCR contributed to these advances?
- What impacts have these advances had on the Department's missions in energy, environment, or security?
- What are the key aspects of the ASCR investment strategy that have had the greatest impacts?
- Looking to the future, and building on the ASCAC reports, what research areas and funding strategies to pursue those areas could further strengthen ASCR in serving the DOE's missions?

I would appreciate receiving a written report by December 31, 2018.

Sincerely,



J. Stephen Binkley
Deputy Director for Science Programs
Office of Science

cc: Barbara Helland, SC-21
Christine Chalk, SC-21.2

APPENDIX 2: CONTRIBUTORS TO THIS REPORT

The following individuals provided input that helped inform this document, but they are not responsible for any inaccuracies we may have introduced. We apologize for any names we have unintentionally overlooked.

Frank Alexander, BNL	Jeff Hittinger, LLNL
Jim Ahrens, LANL	Lee Holcomb, NASA (retired)
Jim Ang, PNNL	Hans Kaper, Georgetown University
John Bell, LBNL	Sid Karin, UC San Diego (ret.)
Pete Beckman, ANL	Scott Klasky, ORNL
Martin Berzins, University of Utah	David Kuck, Intel
Wes Bethel, LBNL	Harriet Kung, BES
Pavel Bochev, SNL	Alex Larzelere, Larzelere & Associates, LLC
David Brown, LBNL	Alan Laub, UCLA
Rich Carlson, ASCR	Randall Laviolette, ASCR
John Cary, Tech-X Corp.	Richard Lethin, Reservoir Labs
John Cavalliini, ASCR (retired)	Xiaoye "Sherry" Li, LBNL
Christine Chalk, ASCR	Ewing "Rusty" Lusk, ANL (retired)
Barbara Chapman, BNL	John Mandrekas, FES
Lali Chatterjee, HEP	David Martin, ANL
Jeff Christensen, EntryPoint Networks, LLC	Barney McCabe, ORNL
Eric Colby, HEP	Lois Curfman McInnes, ANL
Jim Crowley, SIAM	Cleve Moler, MathWorks
Lois Curfman McInnes, ANL	Jorge More, ANL
Rob Falgout, LLNL	Todd Munson, ANL
Josh Fisher, HP Labs (retired)	Jeffrey Nichols, ORNL
Geoffrey Fox, Indiana U.	Shelly Olsan, Krell Institute
Alan Gara, Intel	Mike Papka, ANL
Gary Geernaert, BER	David Park, Voltaiq, Inc.
Tim Germann, LANL	Vern Paxson, LBNL, UC Berkeley and International Computer Science Institute
Allan Gottlieb, NYU	Linda Petzold, UC Santa Barbara
Leslie Greengard, NYU	Walt Polansky, ASCR (retired)
William Gropp, University of Illinois Urbana- Champaign	James Pool, CalTech (retired)
Jim Hack, ORNL	Terri Quinn, LLNL
Aric Hagberg, LANL	Justin Rattner, Intel Corp. (retired)
Timothy J. Hallman, NP	Rob Ross, ANL
Jim Hack, ORNL	James Sethian, UC Berkeley and LBNL
Vijay Hanagandi, Optimal Solutions, Inc.	Tal Sholklapper, Voltaiq
Barb Helland, ASCR	Andrew Siegel, ANL
Mike Heroux, SNL	James Siegrist, HEP
Dan Hitchcock, ASCR (retired)	Barry Smith, ANL

Rick Stevens, ANL
Jim Stewart, SNL
Valerie Taylor, ANL
Rajeev Thakur, ANL
Suzy Tichenor, ORNL
David Trebotich, LBNL
Raymond Tuminaro, SNL
Kerstin Kleese Van Dam, BNL

Robert Ward, ORNL and University of
Tennessee
Sharlene Weatherwax, BER
Andy White, LANL (retired)
Sam Williams, LBNL
Margaret Wright, NYU
Kesheng "John" Wu, LBNL
Kathy Yelick, LBNL
Jun Zhu, Power Info, LLC

APPENDIX 3: FOR FURTHER READING

The work surveyed in this document spans an enormous range of time periods and scientific disciplines. Details about the technical advances and impacts can be found in the archival, peer-reviewed literature. Here, we provide some pointers to additional historical sources that shaped and reflected the growth of ASCR, but which may not be in the archival record or might be difficult to find. It is our intention that these documents and other primary materials we uncovered will be archived by DOE's Office of Scientific and Technical Information in close proximity to the official web home for this report.

National CTR Computer Center; LLNL Energy and Technology Review, Dec. 1975. This article documents the launch of the CTRCC, which evolved into NERSC, describing the first computer, a CDC 7600, and the "high-speed (50 000-bits/s) transmission lines." (Doc title: LLNL Report on Origin of CTRCC_1975.pdf)

The MFECC Electronics Group has Played a Vital Role in the Creation of the National MFE Computer Network; Electronics Engineering Department Quarterly Report No. 1 -- 1977, LLNL. Report describes the origins of today's ESnet, launched with dial-up modems to a full network of PDP-10-based user service centers connecting to the Magnetic Fusion Energy Computer Center over dedicated telephone circuits. (Doc title: Origin of MFEnet1977)

Report of the Panel on Large Scale Computing in Science and Engineering; Peter D. Lax, chair, December 1982. This report, sponsored by DOE and NSF in cooperation with DOE and NASA, recommended increased access to supercomputers for researchers, increased research in computational math, software and algorithms, more training, and R&D basic to the design and implementation of new systems with increased capability and capacity . (Doc title: Lax_Report_1982)

The Future of Intersite Networking; report from the First Annual Workshop on Energy Research Computing, Oct. 27-28, 1986, LBNL. Held shortly before ESnet was created, this workshop report summarizes discussions on future networking demands, consolidating ER networks, the possible role of commercial networks, network security and other issues. (Doc title: The Future of Intersite Networking.pdf)

ESNET (1987 Program Plan), U.S. Department of Energy Office of Science Scientific Computing Staff, William Boatwick, LANL, et al, June 1987. The report, the initial program plan for ESnet, was written by the cross-program Energy Sciences Network Steering Committee to codify overall ER computer network requirements, to document and set priorities for computer networking requirements including performance objectives. (Doc title: 1987 ESnet Program Plan)

A Research and Development Strategy for High Performance Computing; Office of Science and Technology Policy, Nov. 1987. This report was the result of a systematic review of the status and directions of high performance computing and its relationship to federal research and development. Conducted by the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), the review involved a series of workshops attended by numerous computer scientists and technologists from academia, industry, and government. (Doc title: R and D Strategy for HPC-Nov-1987-ED3360886)

Grand Challenges to Computational Science; Kenneth G. Wilson, Ohio State University, 1989, published in *Future Generation Computer Systems* 5 (1989) published by Elsevier. This paper, which focuses on the NSF centers launched in 1985, aims to define computational science and to categorize the principal problems it faces. (Doc title: Ken Wilson 2.pdf)

Federal High Performance Computing and Communications Program. The Department of Energy Component; DOE Office of Energy Research; June 1991. This report elaborates on the DOE research program in High Performance Computing and Communications (HPCC). The DOE HPCC program emphasizes research in four areas: (1) HPCC Systems; (2) Advanced Software Technology and Algorithms; (3) National Research and Education Network; and (4) Basic Research and Human Resources. (Doc title: DOE Component of Federal HPCC-June 1991-ED344569)

The DOE Program in HPCC: High-Performance and Communications; DOE Office of Energy Research, Office of Scientific Computing, March 1993. This document reports to Congress on the progress that the Department of Energy has made in 1992 toward achieving the goals of the High Performance Computing and Communications (HPCC) program. Its second purpose is to provide a picture of the many programs administered by the Office of Scientific Computing under the auspices of the HPCC program. (Doc title: DOE Program in HPCC-March 1993-ED371721)

Computing at the Leading Edge: Research in the Energy Sciences; Arthur Mirin, NERSC/LLNL, 1993. This report describes how computing is fundamentally changing research through large-scale "computational experiments" and how advances in microprocessor technology have further accelerated computation. (Doc title: Computing at the Leading Edge_Research in the Energy Sciences.pdf)

Argonne National Laboratory and the Emergence of Computer and Computational Science, 1946-1992; History Thesis by Charles Nelson Yood, Pennsylvania State University, 2005. Yood's This dissertation uses the Applied Mathematics Division of ANL as a window to explore the emergence of computer and computational science as independent scientific disciplines. (Doc title: Yood_Dissertation.pdf)

Scientific Discovery: A Progress Report on the U.S. Department of Energy's Scientific Discovery through Advanced Computing (SciDAC) Program; Jon Bashor et al., LBNL, 2007. This report documents the accomplishments of the projects supported during the first five years of the SciDAC program. (Doc title: Final SciDAC Report (SciDAC 1))

Applied Mathematics at the U.S. Department of Energy: Past, Present and a View to the Future; David Brown, et al. LLNL, May 2008. This report by an independent panel from the applied mathematics research community describes how DOE needs to leverage and expand its applied math strengths to more fully address research in complex systems. (Doc title: Applied Math Report_D. Brown.pdf)

Breakthroughs 2008; Report of the Panel on Recent Significant Advancements in Computational Science; This report identifies recent breakthroughs in computational science and enabling technologies, supported in a broad sense by the Office of Advanced Scientific Computing Research through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, the Scientific Discovery through Advanced Computing (SciDAC) program, and/or its base program. (Doc title: Breakthroughs_0902)

INCITE: In Review; 2011. This brochure surveys research breakthroughs supported by the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program has promoted transformational advances in science and technology through large allocations of computer time, supporting resources, and data storage. (Doc title: INCITE_In_Review_FINAL_7-19-11)

The TOP500 List: Twenty Years of Insight into HPC Performance; Jon Bashor et al, 2012. This brochure, marking the twentieth year of the TOP500 list, describes each of the systems to have achieved the number one ranking on the list between June 1993 and Nov. 2012. Of the 15 number one systems highlighted nine were at DOE national labs. (Doc title: TOP500_20th_Anniv_Brochure)

Adventures in parallelism: Celebrating 30 years of parallel computing at Argonne; Gail Pieper and Laura Wolf, May 2013. The 1980s saw a shift in computing architectures from serial machines, which relied on a single fast processor, to a model that employed many not-so-fast processors ganged together. Believing parallel computing to be the only realistic way to continue scaling up in power, Argonne established the Advanced Computing Research Facility (ACRF) in 1983. Housing as many as 10 radically different parallel computer designs, the ACRF enabled applied mathematicians and computer scientists to learn parallel computing and experiment with emerging parallel architectures. Available online at <https://www.anl.gov/article/adventures-in-parallelism-celebrating-30-years-of-parallel-computing-at-argonne>

NERSC: 40 Years at the Forefront (1974-2014); Jon Bashor et al., LBNL, 2014. This report documents the evolution of NERSC from its start with a cast-off CDC 6600 computer to its role as the primary computing facility for the Office of Science's research community. (Doc title: NERSC_40th_brochure.pdf)

Twenty-five Years of Leadership Science at the Oak Ridge Leadership Computing Facility; Katie Bethea, et al., OLCF, 2017. This report documents 25 years of research conducted at the OLCF since its founding as the Center for Computational Sciences in 1992, as well as the center's ability to deliver supercomputers of unprecedented capability. (Doc title: OLCF_25th_Anniversary.pdf)

For news from the ASCR National User Facilities:

Argonne Leadership Computing Facility (ALCF): <https://www.alcf.anl.gov/news>

Energy Sciences Network (ESnet): <https://www.es.net/news-and-publications/esnet-news/>

National Energy Research Scientific Computing Center (NERSC): <https://www.nersc.gov/news-publications/nersc-news/>

Oak Ridge Leadership Computing Facility (OLCF): <https://www.olcf.ornl.gov/olcf-media/olcf-news/>

APPENDIX 4: ACRONYMS

ACM	Association for Computing Machinery
ACRF	Advanced Computing Research Facility
ADIFOR	Automatic Differentiation of Fortran Codes
AEC	Atomic Energy Commission
AI	Artificial Intelligence
ALCC	ASCR Leadership Computing Challenge
ALCF	Argonne Leadership Computing Facility
AMR	Adaptive Mesh Refinement
AMS	Applied Mathematical Sciences
ANL	Argonne National Laboratory
ARM	Atmospheric Radiation Measurement
ARPA	Advanced Research Projects Agency
ARPA-E	Advanced Research Projects Agency—Energy
ARPANET	Advanced Research Projects Agency Network
ASC	Advanced Simulation and Computing
ASCAC	Advanced Scientific Computing Advisory Committee
ASCI	Accelerated Strategic Computing Initiative
ASCR	Office of Advanced Scientific Computing Research
AVIDAC	Argonne Version of the Institute’s Digital Arithmetic Computer
BDF	Backward Differentiation Formulas
BER	Office of Biological and Environmental Research
BES	Office of Basic Energy Sciences
BLAS	Basic Linear Algebra Subprograms
BRHR	Basic Research and Human Resources
BTU	British Thermal Unit
C3P	Caltech Concurrent Computing Project
CANDLE	CANcer Distributed Learning Environment
CCII	Computational Center for Industrial Innovation
CCS	Center for Computational Sciences
CDC	Control Data Corporation
CMS	Compact Muon Solenoid
CORAL	Collaboration of Oak Ridge, Argonne, and Livermore
CPU	Central Processing Unit
CRADA	Cooperative Research and Development Agreement
CSCC	Concurrent Supercomputing Consortium
CSGF	Computational Science Graduate Fellowship
CTR	Controlled Thermonuclear Research
CTRCC	Controlled Thermonuclear Research Computer Center
CTSS	Cray Time Sharing System
DAE	Differential-Algebraic Equations
DAKOTA	Design and Analysis toolKit for Optimization and Terascale Applications
DARPA	Defense Advanced Research Projects Agency
DASSL	Differential-Algebraic System SoLver
DD	Director’s Discretionary program
DOE	U.S. Department of Energy
DRS	Dynamic Right-Sizing

DTN	Data Transfer Nodes
E3SM	Energy Exascale Earth System Model
ECI	Exascale Computing Initiative
ECP	Exascale Computing Project
ERCAP	Energy Research Computing Allocations Process
ERDA	Energy Research and Development Administration
EPSI	Center for Edge Plasma Simulation
ESG	Earth System Grid
ESnet	Energy Sciences Network
FES	Office of Fusion Energy Sciences
FLOPS	Floating Point Operations per Second
FMM	Fast Multipole Method
Gbps	Gigabits per second
GPU	Graphics Processing Unit
GSI	Grid Security Infrastructure
HEP	Office of High Energy Physics
HEPnet	High Energy Physics Network
HPC	High Performance Computing
HPC4Mfg	HPC for Manufacturing
HPCC	High Performance Computing and Communications
HPCG	High Performance Conjugate Gradient
HPCRC	High Performance Computing Research Centers
HPCS	High Performance Computing Systems
HPL	High Performance LINPACK
HPSS	High Performance Storage System
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
I/O	Input/Output
IP	Internet Protocol
iPSC	Intel Personal Supercomputer
KSR-1	Kendall Square Research computer
LANL	Los Alamos National Laboratory
LBNL	Lawrence Berkeley National Laboratory
LCF	Leadership Computing Facility
LCG	Linear Congruential Generator
LCLS	Linac Coherent Light Source
LCLS-II	Linac Coherent Light Source – II
LEED	Leadership in Energy and Environmental Design
LIGO	Laser Interferometer Gravitational-wave Observatory
LHC	Large Hadron Collider
LHC ATLAS	LHC A Toroidal LHC ApparatuS experiment
LLNL	Lawrence Livermore National Laboratory
Mbps	Megabits per second
MFEt	Magnetic Fusion Energy Network
MICS	Mathematical, Information, and Computational Sciences
MIMD	Multiple Instruction, Multiple Data
MPAS	Model for Prediction Across Scales
MPI	Message Passing Interface
MPP	Massively Parallel Processing
NASA	National Aeronautics and Space Administration

NCAR	National Center for Atmospheric Research
NEOS	Network-Enabled Optimization System
NERSC	National Energy Research Supercomputing Center
NESAP	NERSC Exascale Science Application Program
NMFECC	National Magnetic Fusion Energy Computing Center
NNSA	National Nuclear Sciences Administration
NP	Office of Nuclear Physics
NRC	Nuclear Regulatory Commission
NREN	NASA Research and Engineering
NSF	National Science Foundation
NUCLEI	Nuclear Computational Low Energy Initiative
NVMe	Non-Volatile Memory Express
NYU	New York University
ODE	Ordinary Differential Equations
OLCF	Oak Ridge Leadership Computing Facility
ORACLE	Oak Ridge Automatic Computer and Logical Engine
ORNL	Oak Ridge National Laboratory
OSCARS	On-demand Secure Circuits and Advance Reservation System
OSG	Open Science Grid
OSTI	Office of Scientific and Technical Information
OSTP	White House Office of Science and Technology Policy
P&G	Proctor & Gamble
PDP-10s	Program Data Processor-10s
PERI	Performance Engineering Research Institute
PETSc	Portable, Extensible Toolkit for Scientific Computation
PGAS	Partitioned Global Address Space
PICS	Partnership in Computational Science
PMaC	Performance Modeling and Characterization
PNNL	Pacific Northwest National Laboratory
PPPL	Princeton Plasma Physics Laboratory
PVM	Parallel Virtual Machine
QCD	Quantum Chromodynamics
R&D	Research and Development
SBIR	Small Business Innovation Research
SciDAC	Scientific Discovery through Advanced Computing
Science DMZ	Science Demilitarized Zone
SCS	Scientific Computing Staff
SDC	Spectral Deferred Corrections
SDN	Software Defined Networking
SENSE	SDN for End-to-end Network Services at Exascale
SENSEI	Sub-Electron-Noise Skipper Experimental Instrument
SIAM	Society for Industrial and Applied Mathematics
SIMD	Single Instruction, Multiple Data
SISAL	Streams and Iteration in a Single Assignment Language
SMT	Satisfiability Modulo Theory
SNL	Sandia National Laboratories
SOC	System-on-Chip
SPD	Symmetric Positive-Definite
SSO	Single sign on

STEM	Science, Technology, Engineering, and Mathematics
TAMPR	Transformation Assisted Multiple Program Realization
TAO	Toolkit for Advanced Optimization
TCP	Transmission Control Protocol
UNEDF	Universal Nuclear Energy Density Functional
VERIFI	Virtual Engine Research Institute and Fuels Initiative
WLCG	Worldwide Large Hadron Collider Computing Grid