LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-717318

Strengthening LLNL Missions through Laboratory Directed Research and Development in High Performance Computing

D. K. Willis

January 10, 2017

Disclaimer

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Strengthening LLNL Missions through Laboratory Directed Research and Development Investments in High Performance Computing

A White Paper

December 2016

Introduction

High performance computing (HPC) has been a defining strength of Lawrence Livermore National Laboratory (LLNL) since its founding. Livermore scientists have designed and used some of the world's most powerful computers to drive breakthroughs in nearly every mission area. Today, the Laboratory is recognized as a world leader in the application of HPC to complex science, technology, and engineering challenges. Most importantly, HPC has been integral to the National Nuclear Security Administration's (NNSA's) Stockpile Stewardship Program—designed to ensure the safety, security, and reliability of our nuclear deterrent without nuclear testing.

A critical factor behind Lawrence Livermore's preeminence in HPC is the ongoing investments made by the Laboratory Directed Research and Development (LDRD) Program in cutting-edge concepts to enable efficient utilization of these powerful machines. Congress established the LDRD Program in 1991 to maintain the technical vitality of the Department of Energy (DOE) national laboratories. Since then, LDRD has been, and continues to be, an essential tool for exploring anticipated needs that lie beyond the planning horizon of our programs and for attracting the next generation of talented visionaries.

Through LDRD, Livermore researchers can examine future challenges, propose and explore innovative solutions, and deliver creative approaches to support our missions. The present scientific and technical strengths of the Laboratory are, in large part, a product of past LDRD investments in HPC.

Here, we provide seven examples of LDRD projects from the past decade that have played a critical role in building LLNL's HPC, computer science, mathematics, and data science research capabilities, and describe how they have impacted LLNL's mission.

Highlighted Projects

Lawrence Livermore uses LDRD to sustain excellence and advance science and technology in critical areas, such as HPC, computational mathematics, and data science, that distinguish the Laboratory and enable it to meet enduring mission needs. In addition, work funded by LDRD positions the Laboratory to support emerging needs and to provide opportunities for program growth. LLNL depends on LDRD funding to investigate and develop expertise in a given area that will eventually aid the Laboratory in developing greater simulation capabilities, faster application development, innovative computer science research, and opportunities for national leadership and visibility.

The following seven project synopses capture the origins of the LDRD proposals, what was achieved with the investments, and how the outcomes have impacted the mission space at LLNL.

First Principles Molecular Dynamics Scaled Up
PI: Jean-Luc Fattebert

Biologists, chemists, and physicists use first-principles molecular dynamics (FPMD) codes to calculate properties of materials and molecules. These codes are based on quantum mechanics modeling, which is computationally demanding because it calculates many interactions between all the atoms and electronic orbitals in the problem and requires intensive communication between processors and memory. With traditional algorithms, simulation size is limited to a few hundred atoms, which is much too small for simulating complex systems or realistic materials.

Livermore researchers, led by Jean-Luc Fattebert and funded by LDRD for three years beginning in 2011, developed new algorithms with reduced computational complexity and better parallel scaling to efficiently use the largest supercomputers available now and in the future to run FPMD simulations with many thousands or even millions of atoms. With these algorithms, the number of atoms that can be simulated is directly proportional to the number of processors available. The method represents only the interactions between an atom's electrons and a certain (and adjustable) number of its neighbors. In a large system, interaction across wide distances is less significant than between nearby electrons and can be safely neglected. With this approach, the number of atoms and the complexity of the systems that can be simulated increase dramatically.

The team demonstrated excellent scaling using the novel algorithm and an accompanying new FPMD code called MGmol. For example, simulating the motion of 1,200,000 atoms for 98 time steps on roughly 1,500,000 processors of the Sequoia machine took less than 3 hours. If the team were to solve the same problem using traditional algorithms and the same number of processors, the problem would take about one hundred years to solve and require a computer with far more memory than any currently available. This advancement earned Fattebert and his team a 2016 Gordon Bell Prize finalist spot for outstanding achievement in HPC.

Of particular note is that the team achieved their scaling results without sacrificing accuracy, as has occurred in many previous attempts to speed up FPMD calculations. Fortunately, validations of their solutions for nonmetallic systems, used for studying

problems such as the dilute solutions in energy-storage devices, are progressing well. For instance, when calculating a problem involving 512 water molecules, both MGmol and Livermore's FPMD code QBox, which uses more traditional algorithms, produced virtually identical results. For this system size, time to solution using MGmol is already a factor 2 faster than using QBox, but more significant gains are obtained for larger problems.

This capability not only allows scientists to more efficiently use the supercomputers presently available and in the future, it enables research with more realistic models of matter than we use today, involving, for instance, more realistic defects and more complicated molecular structures in materials, or calculation of the equation of state of polymers with realistic molecular structure.

The approach also opens up new classes of FPMD simulations, such as dilute aqueous solutions. Dilute solutions is an area where a large number of atoms is needed to model realistic physical problems. Sea water, for instance, contains 1 salt molecule (1 sodium atom and 1 chlorine atom) per 100 water molecules, so that 300 atoms (3 atoms per water molecule) are needed for every salt molecule in the system. Saline solutions used in medicine (e.g., intravenous infusions) are even more dilute, with 360 water molecules for every salt molecule. Thus, at least tens of thousands of atoms are needed to draw statistically meaningful conclusions for such a problem. The ability to model dilute solutions is relevant for many scientific fields (such as dilute alloys in metallurgy), and accurate, large-scale MD simulations of millions of atoms may shed light on problems not fully understood today.

Many other fields will benefit from accurate, large-scale MD simulations. In molecular biology, several ions with orders-of-magnitude differences in their concentrations are regulated with remarkable accuracy via extremely complex interrelated processes (e.g., K+/Na+ ion pump). Even simulating part of such an arrangement will require a million-atom system size. In addition, atomistic models of electrochemical interfaces (e.g., energy storage or conversion devices) have not developed to the same level of clarity as, say, semiconductor physics. Ion concentration can vary significantly near the electrode–electrolyte interfaces, and MD simulations of millions of atoms will be required for a faithful description.

LDRD funding of this project supported the hiring of postdoctoral researcher Daniel Osei-Kuffuor, who has since accepted a full-time staff appointment at LLNL.

The team is now using LDRD investments to develop similar methods for simulating metallic systems—a more challenging problem because fewer electron interactions can be ignored, making computational cost reductions more difficult to achieve. If successful in this metallic realm, very large and accurate simulations will help researchers better understand alloy properties and perhaps help create better alloys.

Metagenomics and New Computational Methods for Scalable Genome Discovery
PI: Jonathan Allen

Starting in 2012, LDRD funding has been used to address a key need in the Laboratory's biosecurity mission space by improving the capability to detect and diagnose engineered

and emerging diseases, thus improving the ability to respond to a biological attack. In addition, it has bolstered the Laboratory's core competency in information systems and data science by demonstrating the use of low-cost, multicore compute nodes augmented by persistent flash computer memory.

The team, led by bioinformatics scientist Jonathan Allen, developed a novel approach to metagenomic sequence analysis using flash drives as a supplemental memory source to more efficiently search very large datasets. Their efforts led to the creation of the Livermore Metagenomic Analysis Toolkit (LMAT), a genome analysis technology that addresses the fundamental scaling limitations of pre-existing metagenomic classification methods by using a custom reference database with a fast searchable index.

Metagenomic classification involves comparing millions to billions of genetic fragments with unknown origin to a large database of reference genomes. To avoid long runtimes, the reference database is limited in size and scope and usually stores only a representative set of bacterial strains. Organisms with larger genomes, such as fungi, are analyzed in rare cases only when using sufficiently large computer clusters. Previous metagenomic approaches could not exploit important genetic relationships between the different reference organisms, a technique that enables efficient fragment matching of multiple reference organisms. LMAT, in contrast, compares each reference genome sequence across the entire database and tags sub-sequences that uniquely identify a biological strain or species. This process helps accelerate the classification process and improves the accuracy of the results.

A major challenge is to keep pace with the data deluge that the sequencing process produces. To overcome the related memory and data storage problems, the LMAT team turned to innovations in supercomputing hardware and software and began applying cutting-edge storage and access methods, using Livermore's Catalyst machine as a testbed. An essential innovation of the project was the development of a data structure that was optimized to store the searchable index on flash drives as if it were in memory. When performing the lookups, the database has been memory mapped from the flash drive into a single address space, treating it as if it were in memory. Unlike conventional storage structures that use disk memory, LMAT's use of both flash (NVRAM) and main (DRAM) memories increases processing. Additional algorithmic developments improved classification accuracy, while a method for splitting index data improved data retrieval time. The new software enables searching more complete reference genome databases than would otherwise be possible on a single compute node.

The team's research results proved that fast, accurate, and scalable microbial detection can be achieved with substantially reduced compute resources. The effort has since led to the development and deployment of a microbial threat detection pipeline, a technology that is used by government agencies for sequence-based pathogen surveillance, diagnostics, and characterization. The tool is designed to operate with 2 TB of flash memory and 384 gigabytes of DRAM, which is roughly equivalent to one Catalyst node.

This work helped support a NASA Jet Propulsion Laboratory effort to characterize microbes in the air and dust filters on the International Space Station.

In 2015, LDRD invested in an effort to develop a genomic sequence DeBruijn graph for antibiotic microbial resistance detection. The use of novel graph structures and

algorithms for genome analysis will make use of the unique HPC available at LLNL and advance the state of the art in Livermore's HPC, simulation, and data science core competency. The team expects to deliver a new capability to the scientific community that enables searching large genomic data sets against genome databases that were previously considered too large to search.

It is relatively uncommon to develop high performance algorithmic tools in a biological space. With the support of LDRD, Lawrence Livermore is one of a very few institutions developing these capabilities and using NVRAM as memory mapping for pathogen-detection efforts.

This LDRD effort partially funded the hiring of postdoctoral researcher Sasha Ames, who has since converted to a full-time LLNL staff member.

New Discretization Framework Leads to Better Simulations and More Efficient Use of HPC

PI: Tzanio Kolev

The mission of the Laboratory is to solve the nation's most difficult national security problems. By nature, these problems lack a well-defined path to solution. In fact, the path is often completely unknown. LDRD funding helps support these riskier, but potentially high payoff projects.

Such was the case in 2009, when LDRD investments were made in a team of computational mathematicians, physicists, and computer scientists led by applied mathematician Tzanio Kolev, who sought to bring more modern discretizations to simulations of compressible flow and other problems that arise in conditions of extreme pressure, such as those in the National Ignition Facility (NIF). The team's plan—to develop new finite element discretization algorithms for Lagrangian shock hydrodynamics and address several longstanding issues with the hydrodynamics simulation codes at LLNL—was risky because finite elements had not been applied to these types of problems before, but the potential payoff was recognized and the project funded.

Hydrodynamics simulations are of critical importance in numerous LLNL applications, including stockpile stewardship and inertial-confinement fusion. For example, at NIF, researchers who study the complex shock wave interactions occurring within high-pressure experiments have found that standard low-order finite difference simulation techniques sometimes fail to satisfactorily represent material behavior under experimental conditions, particularly at material interfaces. In addition, significant human intervention was required when a mesh tangled during the simulation.

The team's solution was to create a new, high-order Arbitrary Lagrangian–Eulerian (ALE) framework. Standard low-order ALE methods are based on approximating the solution to a complex equation by breaking down the problem into many smaller parts (elements) using an unstructured grid, or mesh. The team's high-order ALE framework uses elements with more control points within each mesh element, allowing them to curve element boundaries and the geometry inside them.

Their results have far exceeded expectations. The team has successfully demonstrated that the high-order ALE framework can produce accurate and robust simulations of various shock hydrodynamics problems through the method's implementation in the Livermore-developed BLAST code. Their work has also benefitted research on parallel performance and next-generation computer architectures. Higher-order methods have greater FLOP/byte ratios, meaning that more time is spent on floating-point operations relative to memory transfer, an important characteristic of numerical algorithms for extreme-scale computing.

With additional LDRD funding, the team focused on developing and integrating a new remapping algorithm into the framework, allowing BLAST to simulate larger time steps and to more accurately model multiple materials within one element. Overall, the remapping algorithm has demonstrated excellent parallel scalability, geometric flexibility, and accuracy on several model problems and single- and multiple-material ALE simulations. In 2015, the ASC Program began funding this work under the Advanced Technology Development and Mitigation (ATDM) program element, recognizing it as foundational to the strategy being explored for next-generation simulation capabilities.

With the investments of an LDRD Strategic Initiative (SI), the team is applying high-order solution methods to other types of physics components required for realistic multiphysics simulations, beginning with radiation diffusion.

These LDRD projects has spurred several collaborations with national laboratories and university partners, including Los Alamos; Sandia; Argonne; UC San Diego; the University of Illinois Urbana-Champaign; Virginia Tech; University of Tennessee, Knoxville; Colorado University, Boulder; and the Rensselaer Polytechnic Institute. It has also allowed LLNL to hire two excellent postdoctoral researchers who have since transitioned to full-time staff—Veselin Dobrev and Vladimir Tomov.

An additional outgrowth has been the recent DOE funding for Lawrence Livermore to lead a high-order co-design center to foster greater collaboration between computer hardware, code, and algorithm developers as part of DOE's Exascale Computing Project (ECP).

A New Method for Wave Propagation
PI: Anders Petersson

Seismic wave simulations are essential to several LLNL missions, including national security (e.g., ground motion monitoring, source physics experiments, explosion forensics, seismo-acoustic monitoring) and energy security (e.g., geothermal monitoring, hydraulic fracturing, seismic risk reduction).

In 2005, LDRD funding was granted to explore curvi-linear and mesh refinement methods to allow for efficient modeling of seismic wave propagation with realistic topography at the local to regional scale. The result was WPP, the massively parallel, three-dimensional, finite-difference, elastodynamic wave propagation code written in C++. WPP discretizes the seismic wave equations using a second-order accurate

numerical method that satisfies the principle of summation by parts, thereby guaranteeing energy stability and robustness of the numerical solution for general heterogeneous material properties.

WPP uses three-dimensional models of the local geology, topography, and the structure of the fault. Since its development, it has been validated against benchmark problems relevant to Laboratory program applications. For instance, WPP was the basis for accurate simulations of ground motions during the 1906 San Francisco earthquake, which featured prominently in the 100-year commemoration of the event. WPP was also used to study the influence of shear wave generation due to non-planar topography and aided in the analysis of seismic signals from the North Korean nuclear tests. The simulation capabilities of WPP support national security applications, emergency preparedness, earthquake engineering, and even building siting.

In 2012, in a unique collaboration with scientific and science education advisory panels and nearly 100 contributors from research institutions, including the U.S. Geological Survey (USGS), the University of California at Berkeley, and Lawrence Livermore, the California Academy of Sciences in San Francisco featured a major exhibit, called "Earthquake: Evidence of a Restless Plant". The exhibit included visualizations of WPP ground-motion simulations demonstrating seismic waves traveling through Earth during two major earthquakes (one historical and the other hypothetical). The visualizations were presented in the 75-foot-diameter hemi-spherical Morrison planetarium and were seen by more than two million people. The exhibit ran for 18 months, boosting the scientific literacy and earthquake preparedness of Bay Area residents and visitors, bringing wide visibility to the Laboratory's capabilities, and ultimately showcasing—in a very relatable way—what can be achieved by combining the world's most powerful computers with advanced scientific simulation codes.

In other applications, WPP is used by Livermore's Geophysical Monitoring Program for Source Physics Experiments (SPE) detonated at the Nevada Nuclear Security Site. The SPE shots consist of a series of six underground high-explosive detonations in hard rock that are designed to improve the United States' ability to detect and identify low-yield nuclear explosions amid the clutter of conventional explosions and small earthquakes. SPE team scientists are measuring the ground movement resulting from the explosions and comparing them to simulations backed by WPP. This work will help them develop predictable explosion S-wave and P-wave models to better understand how and where to monitor for nuclear tests.

Although still available and in use today, WPP has been upgraded into a more efficient, fourth-order accurate code called SW4. SW4 is used by the Computational Geophysics group and the Geophysical Monitoring Program within LLNL and by many academic and commercial users on six continents (no users in Antarctica).

With additional LDRD funding, the team generalized the modeling capabilities to study coupled seismo-acoustic wave propagation. Infrasound waves can be generated in the atmosphere by natural phenomena such as exploding bolides, and by man-made nuclear explosions in the atmosphere. During such events, part of the energy is transferred into the solid Earth where it propagates as seismic waves. The seismo-acoustic modeling

capability of the new code ELAC provides a capability unique to Lawrence Livermore that builds upon the technology developed in the WPP and SW4 codes.

DOE's ECP is currently funding further advancement of SW4's modeling capabilities through the Berkeley-led "High Performance, Multidisciplinary Simulations for Regional Scale Earthquake Hazard and Risk Assessments Seismic Simulation" project.

Over time, LDRD investments in wave propagation have helped Lawrence Livermore develop fundamental research in coupling and multi-physics capabilities. Today, Livermore is one of only a few institutions with the expertise and computational resources to advance the scientific understanding of various causes of ground motion through supercomputer simulations.

Exascale Computing Technologies (ExaCT)
PI: Bronis de Supinski

LDRD funding has often been used to help the Laboratory prepare for the changes in computer architectures that can significantly impact application simulation performance. In 2009, the Laboratory recognized the impending challenges that would be faced by the extreme scales inherent in future exascale computing platforms and funded a strategic initiative to address these issues.

This large, multi-directorate project encompassed several interrelated thrust areas in computer science and mathematics: code correctness and performance methods, application-level fault tolerance techniques, scalable mathematics algorithms for multicore systems, and their use in Laboratory applications. The goal was to create fundamentally new approaches not only to algorithmic design, including integrated fault-tolerance strategies, but also to debugging and large-scale performance automation that changing computing system architectures would demand. In addition, the team sought to perform research to significantly increase the scientific output obtained from LLNL's large-scale computing resources by improving application scientists' productivity and system utilization.

One of the key technologies that matured as part of this LDRD effort was SCR, the first multilevel checkpoint/restart system, which helps HPC jobs run more efficiently and recover more work upon failure. HPC systems are growing more powerful by using more compute nodes. As the number of components grows, the system meantime between failure correspondingly drops, and applications must checkpoint frequently to ensure valuable work is not lost in the case of a failure. However, at scale, the cost in time and bandwidth of checkpointing to a parallel file system becomes prohibitive. A solution to this problem is multilevel checkpointing.

With SCR, checkpoint files are stored locally, at other computer nodes and in the parallel file system. The frequency with which checkpoints saved at each level is proportional to the speed at which they can be saved at that level, thereby vastly decreasing the cost of saving valuable restart information. Today, we expect future computers to be built with burst buffer technologies, an architectural trend fostered, in part, by SCR's multilevel strategy and which will provide a key hardware component for SCR checkpointing. SCR provides the application interface and has thus significantly impacted the evolution of

system architectures. This approach will be widely deployed on next-generation systems. The team developed a Markov model that proves that the approach will extend the applicability of checkpoint/restart to Sierra (anticipated delivery beginning in 2017) and to exascale systems.

Another technology that was matured as part of the ExaCT SI project was STAT, or the Stack Trace Analysis Tool. This software tool implements a lightweight debugging technique that radically improves the process of debugging MPI jobs on large-scale systems. Rather than exploring each process in the job, the tool identifies a small subset in which to understand the programming error. The LDRD work focused on improving the scaling of the STAT tool and on allowing it to work with threaded applications, a critical development for current and next-generation architectures. Since its development, STAT has been applied to MPI jobs with more than six million processes and has reduced debugging times from a week or more to 30 seconds or less. It is currently deployed on all LLNL systems and is also distributed by Cray as a part of the standard development environment on their systems. STAT has been essential in solving many of the most difficult problems of the Laboratory's code teams and has saved them countless hours in quickly identifying the root causes of program crashes and hangs. STAT offered a true paradigm shift for debugging, and an R&D 100 award acknowledged it as a technological breakthrough.

Both STAT and SCR are now a critical part of the software ecosystem on LLNL's largest supercomputers to improve application scientists' productivity. They are included in the requests for proposals that LLNL makes for the largest supercomputers that it purchases. These tools will be essential for exascale computing and are supported by DOE's ECP and by various computer vendors (e.g., Intel and Cray) as part of their HPC software offerings.

Essential to the success of the team and a testament to the innovative nature of the work was the recruitment and eventual conversion to full-time staff of postdoctoral researchers Todd Gamblin, Barry Rountree, Kathryn Mohror, and Olga Pearce, as well as three excellent summer students who returned to LLNL for postdoctoral appointments—Thomas Scogland (now a full-time staff member), Aniruhda Maranthe, and Hormozd Gahvari, who became LLNL's first Fernbach Fellow (named postdoctoral fellowship in the computing sciences).

Mathematical Techniques for Data Mining Analysis
PI: Van Emden Henson

LDRD investments dating back to 2010 and led by computational mathematician Van Emden Henson, have focused on using linear algebra, graph theory, and computer science to find latent relationships hidden within massive amounts of unexplored data. Many important problems are modeled using graphs, or collections of vertices and edges connecting the vertices. In the cyber security and intelligence domains, people or computers are often represented by the vertices, while relationships connecting them, such as text messages, phone conversations, e-mail exchanges, or data transmissions, are represented by edges. Graphs also represent social networks and the web; in business, they are crucial to transportation, shipping, and recommender systems; in biology, they

represent protein interactions and genome expressions. Graphs can grow to extraordinary sizes, sometimes $O(10^{10})$ vertices, often evolve rapidly in real time, and typically have properties that make analysis extremely difficult.

Graph algorithms enable the discovery of relationships among vertices such as communication patterns, data flow, connectivity, communities, and associations. The graphs can be represented with matrices and then linear algebra can be used to model and analyze the graphs. Spectral graph algorithms are one method of analysis that has wide-ranging applications, including ranking the importance of vertices, graph partitioning, or community identification; computing "distances" on the graph; and determining the ease of communication among vertices.

LDRD has funded two completed and one ongoing project in this area. The first focused on scalable methods for computing the eigenvalues and eigenvectors (eigenpairs) of matrices representing large-scale graphs and to deliver theory and tools for spectral analysis of such graphs, especially in three major applications areas: 1) partitioning, community-finding, and clustering, 2) distances and paths, 3) vertex importance and ranking. Later, in a second project, Henson and his team expanded the factorization techniques to yield new, robust capabilities aimed at more complicated datasets, namely rectangular matrices and nonsymmetrical square matrices. The results of these factorizations allow researchers to discover relationships in the data, such as hierarchies, groups, and communities. Both projects produced theory and tools that have been turned over to sponsored projects for further research and development into deployable tools.

A current LDRD project, led by team member Geoff Sanders, explores a novel approach to 'graph transit with memory,' a method for finding rich relationships that would be too expensive for standard linear algebra.

These projects have significant applications, particularly in Lawrence Livermore's Global Security organization. These solutions to revealing opaque relationships in the data have built a solid foundation for approaching similar mathematical issues in the future, with highly valuable application to practical problems. The projects have had a major impact on the Laboratory's ability to process and analyze crucial data.

In terms of program growth, the eigensolver LDRD spurred an external sponsor to fund $8M for related research, split between Lawrence Livermore and Sandia. The Laboratory received an additional $250K from a different sponsor and established relationships that would not otherwise have existed. In addition, the team's results have been rigorous enough to challenge widely held beliefs in the data science community.

This research has served as an effective pipeline of students, visitors, and postdoctoral researchers. Postdocs that have been supported by these LDRD investments and have since converted to full-time staff include Geoff Sanders and Christine Klymko. Emily Diana was supported as a summer student and is now a full-time employee.

Data-Centric Computing
PI: Maya Gokhale

Big data—requiring the continuous processing and analysis of sensor, experimental, and

simulation data—is a dominant Laboratory challenge requiring scalable, flexible architectures to match a wide range of applications and budgets. LDRD investments address a critical mission need for data-centric computing and benefit data science applications for both informatics and simulation data analysis, in support of the Laboratory's core competency in HPC, simulation, and data science.

In anticipation of the unfolding revolution in memory technology, an LDRD project led by Maya Gokale pursued research to devise system software and memory hardware to optimize memory access for both data science and scientific applications. Key targets were capacity—near seamless out-of-core execution to large data sets spanning dynamic random-access memory (DRAM) and nonvolatile random-access memory (NVRAM)—and irregular, cache-unfriendly accesses, particularly those using indirect access through index arrays and other noncontiguous access patterns.

The team quantitatively prototyped and evaluated a data-centric compute node memory architecture. The large unified memory/storage subsystem combines DRAM and NVRAM, giving individual compute nodes access to multi-terabyte-scale memory at a fraction of the cost and power of present day DRAM-only memory subsystems. The unified memory is accessed through the data-intensive memory map (DI-MMAP) Linux operating system software module developed in the project that delivered up to fourfold speed compared to the standard Linux mmap implementation on a metagenomics application. The software module is deployed on Lawrence Livermore's Catalyst cluster as part of the standard software environment.

Using field-programmable gate arrays, the team developed prototype active memory and storage controllers that can accelerate memory-bound data-intensive applications by up to 3.4× while simultaneously reducing energy up to 2.6×. The active controller logic operates within the memory or storage subsystem, offloading application-specific computation and data access from the main central processing unit. The active memory controller hardware blocks were the basis of a record of invention and subsequent patent application. The work also won the 2015 Best Paper award at the International Symposium on Memory Systems in Washington, D.C.

To validate the unified active memory concept, a massively parallel, throughput-oriented data-centric graph traversal framework was developed and demonstrated its superior performance. The framework was designed to be scalable and has been demonstrated from a single NVRAM-augmented central processing unit up to the largest supercomputer with 131,000 central processing units. The viability and efficiency of these novel architecture solutions and the graph traversal framework were proven in the June 2015 international Graph 500 competition, when the team achieved rank 3 with the Catalyst cluster, using 300 nodes and Flash storage on a Scale 40 problem; all the other Scale 40 entries used at least 49,000 nodes.

Not only has the nonvolatile memory research supported data-intensive computing in fields as diverse as social network analysis and bioinformatics, it has also boosted exascale computing preparation and influenced pre-exascale procurements. A streamline-tracing scientific data-analysis application that uses the unified memory-access Linux module has enabled novel locality-preserving algorithms well-suited to future exascale

memory systems. The DI-MMAP module is used by LMAT, an open-source software tool in use in the bioinformatics community.

The LDRD-supported persistent memory research has resulted in program growth in several ways, including Advanced Scientific Computing Research (ASCR) awards, funding from NoLoSS (Node Local Storage System), FoX (Fault Oblivious Extreme Scale Execution Environment), Argo ECP (with Argonne), and Memory Hierarchy ECP (with Los Alamos). The team received two Defense Advanced Research Projects Agency (DARPA) awards for Storage Intensive Computing Systems as a DARPA seedling and under UHPC with Georgia Tech Research Institute, as well as two UC Fee awards for data intensive supercomputing and RRAM-based TCAM. In addition, two new projects were funded by an external sponsor to conduct research on graph traversal and data rearrangement in active memory.

Several postdoctoral researchers and technical staff joined the Lawrence Livermore workforce to contribute to this innovative effort, including Sasha Ames (partial funding), Brian Van Essen, Roger Pearce, Scott Lloyd, and Chris Macaraeg.

Conclusion

These seven projects and the excellent teams behind them have stimulated exploration at the forefront of computational science and technology and delivered long-term rewards for the Laboratory and the nation. Along the way, the teams frequently earned recognition through awards and papers published in peer-reviewed journals, building the Laboratory's reputation by providing additional visibility. The projects also helped attract, train, and retain new technical staff.

The projects detailed here are a small sample of an impressive compendium of HPC-related projects that have received LDRD funding in the past 25 years. They were selected to represent the breadth, quality, and impact of scientific work that is conducted under the banner of LDRD.

The HPC advances and scientific innovations from LDRD have provided multiple benefits to Lawrence Livermore programs and stakeholders, consistent with Congressional intent and the Laboratory's science strategy. The research results have led to requisite computational tools to ensure NNSA and LLNL programs meet their enormous potential for advancing national security and to keep Livermore at the forefront of HPC, simulation, and data science.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-TR-717318