INNOVATIVE COMPUTING LABORATORY 2019/20 REPORT

THE UNIVERSITY OF TENNESSEE KNOXVILLE

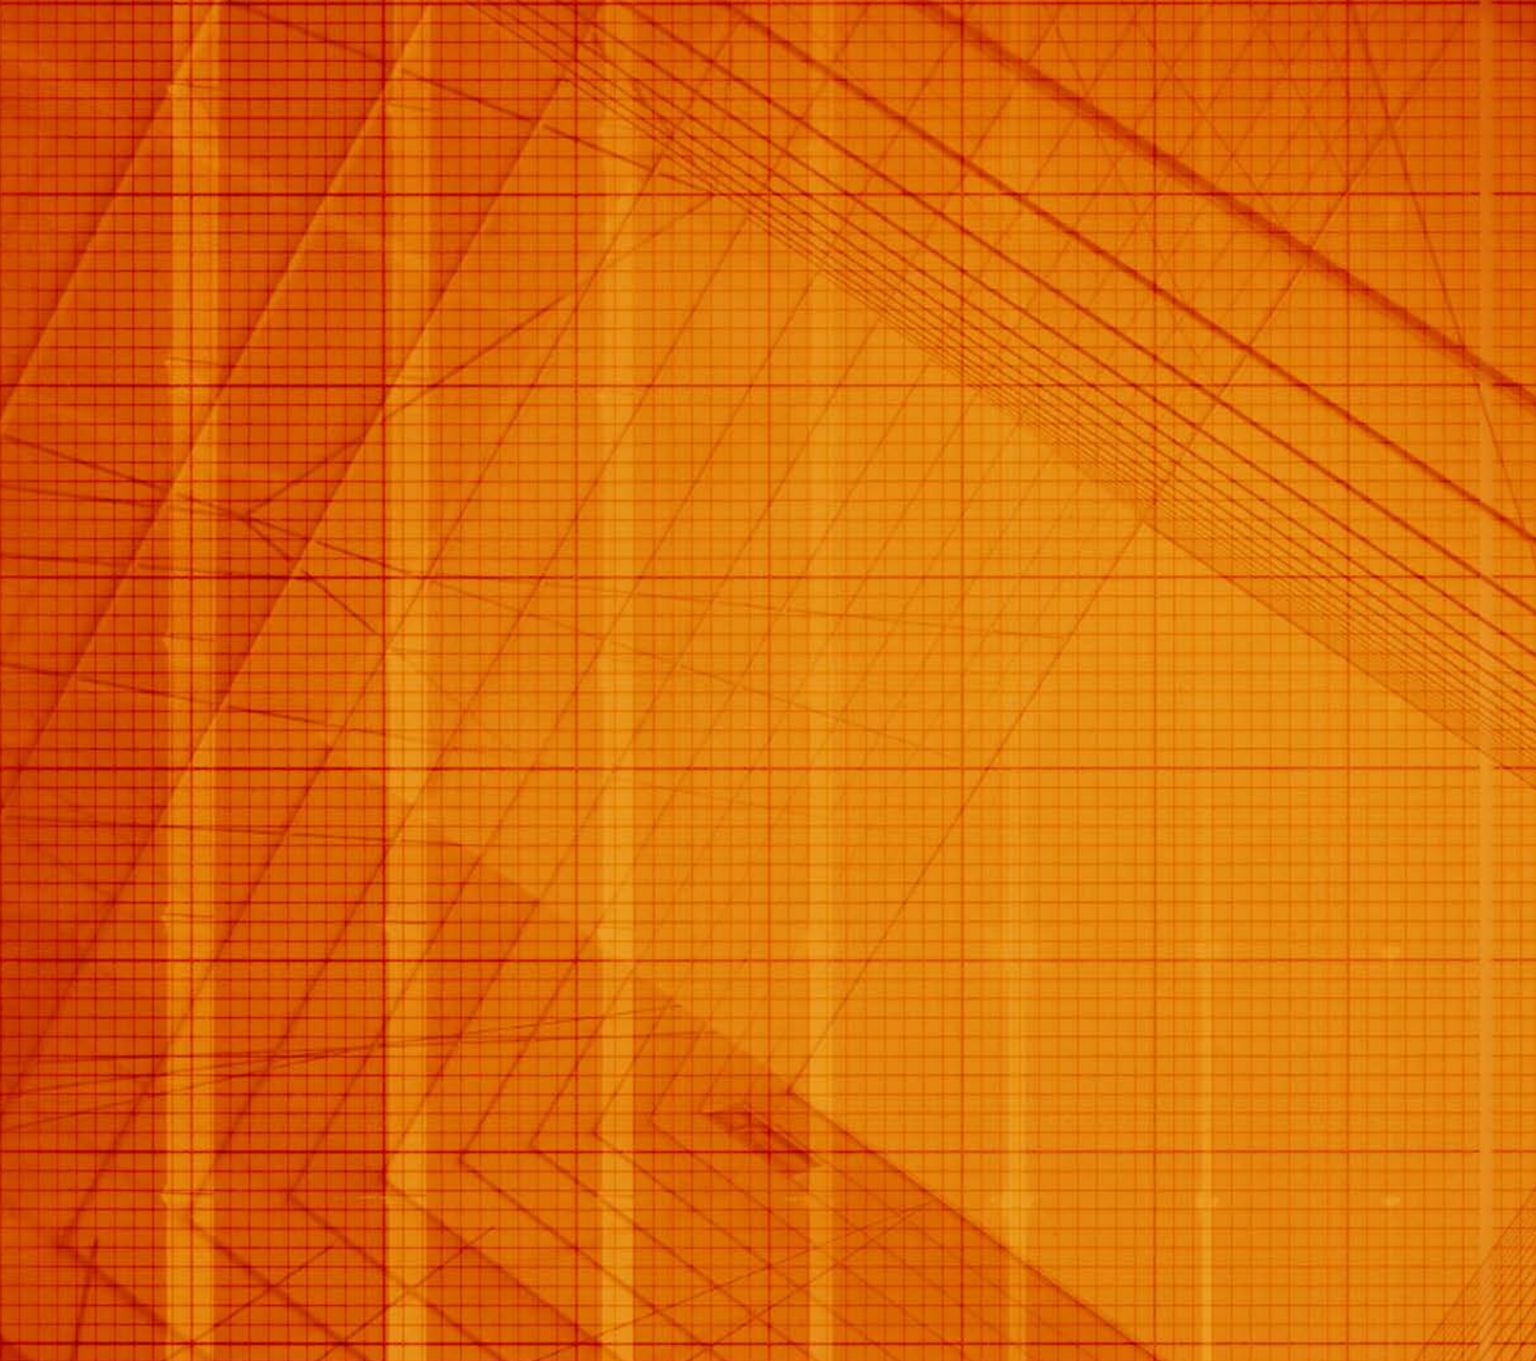

2019/20 ICL REPORT DESIGNED BY **David Rogers** EDITED BY **Sam Crawford**

INNOVATIVE COMPUTING LABORATORY 2019/20 REPORT

2 FROM THE DIRECTOR
4 INTRODUCTION
6 HISTORY
8 HIGHLIGHTS
10 RESEARCH
22 PUBLICATIONS
24 EVENTS
26 PARTNERSHIPS
28 LEADERSHIP
30 PEOPLE
36 ACKNOWLEDGMENTS

FROM THE DIRECTOR

Jack Dongarra
DIRECTOR, ICL

The Innovative Computing Laboratory's (ICL's) thirtieth year was a landmark for several reasons, not the least of which is that an independent computer science research group like ours—at a state university and with such longevity—is a very rare thing. More amazing still, we continue to grow. When we celebrated our twenty-fifth anniversary in 2014, I noted that the high-performance computing (HPC) community was just beginning its long climb to the creation of an exascale supercomputer. At that time, the Department of Energy (DOE) would not launch its billion dollar plus Exascale Computing Project (ECP) for another two years. Today, we are roughly halfway through ECP, ICL is leading five ECP projects and drawing significant funding from three others, and the overall project has just been approved to complete its mission over the next three years. Because of ECP, in 2019 we generated our largest-ever research spending total—more than five million dollars—and our group is the largest it has ever been: 54 people and still hiring.

But the impact of ECP goes beyond the numbers. Our focus on ECP's software development goals is making it possible for ICL to do something that its traditional research agenda never really permitted. Namely, it has given us a chance to revitalize the flagship software packages that have done so much to make us a world leader in scientific computing. With the SLATE project, our linear algebra group is already well along in its effort to completely reengineer, reimplement, and modernize the LAPACK/ScaLAPACK software we have been adding to for twenty-five years. And this year, with the launch of the PAPI++ effort, ECP funding is enabling our performance group to redevelop, from the ground up, a new C++ version of the twenty-year-old Performance API (PAPI) software that performance tool developers around the world rely upon as the foundation for their work.

Of course, while we keep our heads down to meet ECP's regime of development milestones, we cannot afford to let our opportunities for new and innovative research pass us by. In 2019, the distributed computing group led the charge on this front. Most notably, they won a four million dollar award from the National Science Foundation (NSF) to support their Ecosystem for Programming and Executing eXtreme Applications (EPEXA). EPEXA continues, for another four years, the group's important, ongoing collaboration with Robert Harrison at Stony Brook and his colleagues in the computational chemistry community.

There were also signs this year of promising opportunities for future research. Two such areas stand out. First, the push to achieve the highest possible performance in various application areas showed the need for a broader range of precision levels and the ability to vary precision levels within a single application. At the low end, the explosive growth of graphics processing hardware and artificial intelligence (AI) are driving the increased use of "half" precision, 16-bit floating point arithmetic. Our group has major research publications in this area dating back more than a decade. Second, building on the growing success of the High Performance Conjugate Gradients (HPCG) benchmark, this year we expanded our work in the field of scientific benchmarking by introducing HPL-AI—a benchmark for the field of deep learning. Deep learning is truly ushering in a new era in AI, and I expect that future results from the use of this benchmark will help the community understand the hardware changes needed to foster HPC and AI convergence.

As always, the wonderful accomplishments that are represented in this report depend most fundamentally on the high quality of the people who work here. There were bright prospects on that front in 2019 as well. Along with a group of very talented new graduate students, we were also able to add a large, outstanding group of new researchers to our staff, including Nuria Losada, Florent Lopez, Natalie Beams, Dalal Sukkari, Sebastian Cayrols, and Mohammed Al Farhan. Yet, we all know that having great people is as much about good luck as it is about smarts or hard work. For more than twenty-five years we had the good fortune to have an incredibly talented and consummate professional—Tracy Rafferty—leading our administrative staff. Her wide-ranging competence and creativity, her poise and comportment, her tremendous communication and organizational skills, and her dedication, reliability, and great sense of humor are familiar to everyone who worked at or collaborated with ICL during those years. Before she retired this year, she also managed to help facilitate another tremendous stroke of good luck, namely finding Joan Snodderly, another consummate and talented professional of extremely high caliber, to become the new Assistant Director.

As we begin our thirty-first year, the future seems very bright indeed. Our thirtieth reunion was both a terrific celebration of the longevity of ICL and a wonderful reminder of the creativity, hard work, and dedication of all the remarkable people who have passed through our group during those decades. Staying on the leading edge of scientific computing worldwide for all that time is a tremendous accomplishment. I certainly hope that others, both then and now, feel the same sense of pride and gratitude that I do for having the opportunity to have a share in the life of ICL with so many great people. Of course, I am also grateful to the many government, industry, and private sponsors whose funding has helped make that success possible, but I am especially grateful to the University of Tennessee administration and its Tickle College of Engineering for their strong and continuing support.

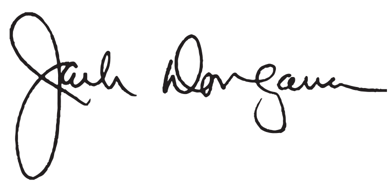

INTRODUCTION

Situated in the heart of the University of Tennessee campus and at the nexus of academia, government, and industry, ICL impacts the world as a leader in advanced scientific computing and HPC through research, education, and collaboration.

The unique challenges of today's computational research are characterized by large datasets and the need for greater performance, energy conservation, and resilience. ICL's cutting-edge efforts, which now span over 30 years, have evolved and expanded with the agility and focus required to address those challenges. ICL's work encompasses a solid understanding of the algorithms and libraries for multi-core, many-core, and heterogeneous computing, as well as performance evaluation and benchmarking for high-end computing. In addition, ICL's portfolio of expertise includes high-performance parallel and distributed computing, with keen attention to message passing and fault tolerance.

The tools and technologies that ICL designs, develops, and implements play a key role in supercomputing-based discoveries in areas like life sciences, climate science, earthquake prediction, energy exploration, combustion and turbulence, advanced materials science, drug design and more.

NUMERICAL LINEAR ALGEBRA

Numerical linear algebra algorithms and software form the backbone of many scientific applications in use today. With the ever-changing landscape of computer architectures, such as the massive increase in parallelism and the introduction of hybrid platforms utilizing both traditional CPUs as well as accelerators, these libraries must be revolutionized in order to achieve high performance and efficiency on these new hardware platforms. ICL has a long history of developing and standardizing these libraries in order to meet this demand, and we have multiple projects under development in this arena.

PERFORMANCE EVALUATION AND BENCHMARKING

Performance evaluation and benchmarking are vital to developing science and engineering applications that run efficiently in an HPC environment. ICL's performance evaluation tools enable programmers to see the correlation between the structure of source/object code and the efficiency of the mapping of that code to the underlying architecture. These correlations are important for performance tuning, compiler optimization, debugging, and finding and correcting performance bottlenecks. ICL's benchmark software is widely used to determine the performance profile of modern HPC machines and has come to play an essential role in the purchasing and management of major computing infrastructure by government and industry.

DISTRIBUTED COMPUTING

Distributed computing is an integral part of the HPC landscape. As the number of cores, nodes, and other components in an HPC system continue to grow explosively, applications need runtime systems that can exploit all this parallelism. Moreover, the drastically lower meantime to failure of these components must be addressed with fault-tolerant software and hardware, and the escalating communication traffic that they generate must be addressed with smarter and more efficient message passing standards and practices. Distributed computing research at ICL has been a priority for two decades, and the lab has numerous projects in this arena under active development.

HISTORY

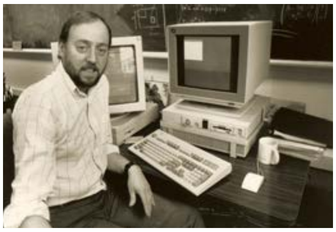

1989

The Level-3 **Basic Linear Algebra Subprograms (BLAS)** specification was developed to perform assorted matrix-multiplication and triangular-system computations.

◄ Jack Dongarra, 1989

1989

The **Parallel Virtual Machine (PVM)** was a software tool for parallel networking of computers designed to allow a network of heterogeneous Unix and Windows machines to be used as a single distributed parallel processor.

1992

Basic Linear Algebra Communication Subprograms (BLACS) were created to make linear algebra applications easier to program and more portable.

1995

Version 1.0 of the **Scalable LAPACK (ScaLAPACK)** library, which includes a subset of LAPACK routines redesigned for distributed memory multiple instruction, multiple data (MIMD) parallel computers, was released.

1997

Automatically Tuned Linear Algebra Software (ATLAS) was an instantiation of a new paradigm in high-performance library production and maintenance developed to enable software to keep pace with the incredible rate of hardware advancement inherent in Moore's Law.

1997

NetSolve (GridSolve) was a client-server system that enabled users to solve complex scientific problems using remote resources.

ATLAS & NetSolve Receive R&D 100 Awards ►

2002

Fault Tolerant MPI (FT-MPI) was an MPI plugin for HARNESS that provided support for fault-tolerant applications crucial for large, long-running simulations.

2003

HPC Challenge was developed for the Defense Advanced Research Projects Agency (DARPA) and consisted of four benchmarks: HPL, Streams, RandomAccess, and PTRANS.

2003

LAPACK for Clusters was developed in the framework of self-adapting numerical software to leverage the convenience of existing sequential environments bundled with the power and versatility of highly tuned parallel codes executed on clusters.

2006

Fault-Tolerant Linear Algebra (FT-LA) is a research effort to develop and implement algorithm-based fault tolerance in commonly used dense linear algebra kernels.

2010

Distributed Parallel Linear Algebra Software for Multi-core Architectures (DPLASMA) is a linear algebra package that enables sustained performance for distributed systems, where each node features multiple sockets of multi-core processors and, if applicable, accelerators like GPUs or Intel Xeon Phi.

2011

The **Parallel Ultra Light Systolic Array Runtime (PULSAR)** project developed a simple programming model for large-scale, distributed-memory machines with multi-core processors and hardware accelerators to automate multithreading, message passing, and multistream multi-GPU programming.

2012

The **Parallel Runtime Scheduling and Execution Controller (PaRSEC)** provides a generic framework for architecture-aware scheduling and management of microtasks on distributed, many-core heterogeneous architectures.

2012

User Level Failure Mitigation (ULFM) is a set of new interfaces for MPI that enables message passing programs to restore MPI functionality affected by process failures.

2014

Argo is an initiative to develop a new exascale operating system and runtime (OS/R) designed to support extreme-scale scientific computation.

2014

The **Rapid Python Deep Learning Infrastructure (RaPyDLI)** delivered productivity and performance to the deep learning community by combining high-level Python, C/C++, and Java environments with carefully designed libraries supporting GPU accelerators and Intel Xeon Phi coprocessors.

2015

PAPI-EX extends PAPI with measurement tools for changing hardware and software paradigms.

Data-driven Autotuning for Runtime Execution (DARE) provides application-level performance tuning capabilities to the end user.

◄ ICL 25th Anniversary

2016

The **Production-ready, Exascale-enabled Krylov Solvers for Exascale Computing (PEEKS)** project will explore the redesign of solvers and extend the DOE's Extreme-scale Algorithms and Solver Resilience (EASIR) project.

2016

The **Software for Linear Algebra Targeting Exascale (SLATE)** project will converge and consolidate previous ICL efforts with LAPACK and ScaLAPACK into a dense linear algebra library that will integrate seamlessly into the ECP ecosystem.

2017

The **Batched BLAS (BBLAS)** effort will create an API for numerical computing routines that process batches of either uniformly sized or varying-size matrices or vectors and will serve as a working forum for establishing this strategy as the next official BLAS standard.

2017

The **MAtrix, TEnsor, and Deep-learning Optimized Routines (MATEDOR)** team is performing the research required to define a standard interface for batched operations (BBLAS) and provide a performance-portable software library that demonstrates batching routines for a significant number of linear algebra kernels.

Prof. Jack Dongarra established ICL in 1989 when he received a dual appointment as a Distinguished Professor at the University of Tennessee, Knoxville (UTK) and as a Distinguished Scientist at Oak Ridge National Laboratory (ORNL). Thirty years later, ICL has grown into an internationally recognized research laboratory specializing in numerical linear algebra, distributed computing, and performance evaluation and benchmarking.

As we look back on the lab's body of work, which now spans three decades, it is important to remember the milestones that shaped the research and direction of ICL. To this end, we present the following projects and initiatives, all of which have special historical significance to ICL and our collaborators.

1992

Still developed today, the **Linear Algebra Package (LAPACK)** is a standard software library for numerical linear algebra.

1993

The **TOP500** was launched to improve and renew the Mannheim supercomputer statistics, which—at the time—had been in use for seven years.

1994

Version 1.0 of a standardized and portable message-passing system, called the **Message Passing Interface (MPI)**, was released.

PVM R&D 100 Award ►

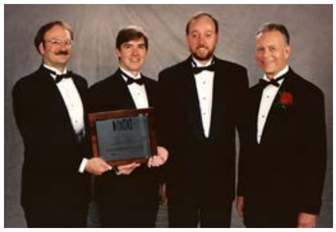

1999

The **Heterogeneous Adaptable Reconfigurable Networked SystemS (HARNESS)** was a pluggable, lightweight, heterogeneous, and distributed virtual machine environment.

1999

Still in active development, the **Performance Application Programming Interface (PAPI)** is a standardized, easy-to-use interface that provides access to hardware performance counters on most major processor platforms.

PAPI R&D 100 Award ►

2000

High-Performance LINPACK (HPL) Benchmark is a benchmark for distributed-memory computers that solves a (random) dense linear system in double-precision (64-bit) arithmetic.

2006

Four institutions merged efforts in the **Open Source Message Passing Interface (Open MPI)**: FT-MPI from UTK/ICL, LA-MPI from Los Alamos National Laboratory, and LAM/MPI from Indiana University, with contributions from the PACX-MPI team at the University of Stuttgart.

2008

Matrix Algebra on GPU and Multi-core Architectures (MAGMA) is a linear algebra library that enables applications to exploit the power of heterogeneous systems of multi-core CPUs and multiple GPUs or coprocessors.

2008

Parallel Linear Algebra Software for Multi-core Architectures (PLASMA) is a dense linear algebra package designed to deliver the highest possible performance from a system of multiple sockets of multi-core CPUs.

2009

The **International Exascale Software Project (IESP)** brought together representatives of the global HPC community to plan and create a new software infrastructure for the extreme-scale systems that represent the future of computational science.

2013

The **Big Data and Extreme-scale Computing (BDEC)** workshop was initiated to map out and account for the ways in which the major issues associated with big data intersect with national (and international) plans being laid out for achieving exascale computing.

BDEC Charleston ►

2013

The **Bench-testing Environment for Automated Software Tuning (BEAST)** project enables writing of tunable high-performance kernels by unleashing the power of heuristic autotuning.

2013

The **High Performance Conjugate Gradients (HPCG)** benchmark is designed to measure performance that is representative of modern HPC capability by simulating patterns commonly found in real science and engineering applications.

2015

The **SparseKaffe** project establishes fast and efficient sparse direct methods for platforms with multi-core processors with one or more accelerators.

2015

The **Task-based Environment for Scientific Simulation at Extreme Scale (TESSE)** uses an application-driven design to create a general-purpose software framework focused on programmer productivity and portable performance for scientific applications on massively parallel hybrid systems.

2016

ICL won seven awards through DOE's Exascale Computing Project (ECP) during the fall of 2016 and is the lead institution on four of these projects:

The **Distributed Tasking for Exascale (DTE)** project will extend the capabilities of the PaRSEC framework.

2016

The **Exascale Performance Application Programming Interface (Exa-PAPI)** project builds on PAPI-EX and extends it with performance counter monitoring capabilities for new and advanced ECP hardware and software technologies.

2018

The goal of **BDEC2**, a follow-on to BDEC and IESP, is to stage six international workshops to enable research communities in a wide range of disciplines to converge on a common platform in order to meet the daunting challenges of achieving exascale computing in the wake of a surging "data tsunami."

2018

The main objective of the ECP **Fast Fourier Transform (ECP-FFT)** project is to design and implement a fast and robust 2-D and 3-D FFT library that targets large-scale heterogeneous systems with multi-core processors and hardware accelerators and to do so as a co-design activity with other ECP application developers.

2019

ICL acquired three new awards in 2019: the **Ecosystem for Programming and Executing eXtreme Applications (EPEXA)**, the **Scalable Run Time for Highly Parallel, Heterogeneous Systems (ScaRT)**, and the **Development of Exascale Software for Heterogeneous and Interfacial Catalysis (DESC)** project.

◄ ICL 30th Anniversary

HIGHLIGHTS

CELEBRATING 30 YEARS

On August 7–9, 2019, ICL celebrated 30 years of conducting innovative research in HPC and cyberinfrastructure. Over the past three decades, ICL has grown from a handful of people in Ayres Hall to over 50 researchers and staff—not to mention the hundreds of ICL alum. The "Thirty Years of Innovative Computing" workshop gathered around 100 current and former ICL members from all over the globe—including attendees from as far away as Japan and Saudi Arabia—for two days of talks and camaraderie in celebration of an organization that has always been a hub of leadership and talent and has striven for excellence for over 30 years.

BEST PAPER FINALIST

Research Assistant Professor Piotr Luszczek and his coauthors were Best Paper Finalists for their entry in the IEEE High Performance Extreme Computing Conference (HPEC 2019) held in Waltham, MA in September of 2019.

The paper, "Increasing Accuracy of Iterative Refinement in Limited Floating-Point Arithmetic on Half-Precision Accelerators," discusses how the emergence of machine learning, which leverages accelerator-based hardware and lower or mixed-precision arithmetic rather than traditional CPUs and full or double precision, is changing the landscape for linear algebra operations in these types of applications.

This lower or mixed-precision approach provides improved performance over traditional methods, but it suffers from a loss of accuracy in the solution. Luszczek et al. proposed a new formulation of LU factorization, called "signed square root LU," which produces more numerically balanced L and U factors that directly address the limited range of the low-precision storage formats—making it possible to recover much of the accuracy in the system solution that would otherwise be lost.

Luszczek, P., I. Yamazaki, and J. Dongarra, "**Increasing Accuracy of Iterative Refinement in Limited Floating-Point Arithmetic on Half-Precision Accelerators**," *IEEE High Performance Extreme Computing Conference (HPEC 2019)*, Best Paper Finalist, Waltham, MA, IEEE, September 2019.

HIGH-PERFORMANCE HARDWARE

In September of 2019, ICL acquired an NVIDIA DGX-1 heterogeneous compute node that is more capable than any other single-node system currently on campus. The DGX-1 comes equipped with two Intel Xeon Broadwell CPUs in a dual socket configuration and is outfitted with eight NVIDIA V100 GPUs and 512 GB of system memory.

The award stems from a longstanding relationship with NVIDIA, including ICL's designation as an NVIDIA CUDA Center of Excellence. It also comes on the heels of another landmark award to UTK in the form of a new Power9-based compute cluster from IBM that uses nodes similar to those in ORNL's Summit supercomputer.

THE ROYAL SOCIETY

FOREIGN MEMBER OF THE ROYAL SOCIETY

On April 16, 2019, Jack Dongarra was elected as a Foreign Member of the Royal Society for his contributions to mathematics and computational science and engineering. The Royal Society is the oldest scientific academy in continuous existence, dating back to 1663, and members include Isaac Newton, Charles Darwin, and Alan Turing—among many other distinguished scientists.

RESEARCH SPOTLIGHT AWARD

In March of 2019, Interim Chancellor Wayne Davis, on behalf of UTK's Office of Research and Engagement (ORE), presented Jack Dongarra with a Research Spotlight Award for ICL's work in high performance and scientific computing research. The ORE awards are a recognition of comprehensive research enterprises—including activities related to funding, mentorship, creative achievement, community engagement, and responsible conduct of research.

OPENMP

In November of 2019, UTK/ICL joined the OpenMP Architecture Review Board (ARB), a group of leading hardware and software vendors and research organizations that are creating the standard for the most popular shared-memory parallel programming model in use today. With a decades-long tradition of being at the forefront of HPC development, UTK will contribute to OpenMP's goals of ease-of-use, portability, and efficiency in parallel programing on current and upcoming HPC systems.

RESEARCH

What originally began over 30 years ago as in-depth investigations of the numerical libraries that encode the use of linear algebra in software has grown into an extensive research portfolio. ICL has evolved and expanded our research agenda to accommodate the heterogeneous computing revolution and focus on algorithms and libraries for multi-core and hybrid computing. As we have gained a solid understanding of the challenges presented in these domains, we have further expanded our scope to include work in performance evaluation and benchmarking for high-end computers, as well as work in high-performance parallel and distributed computing, with efforts focused on message passing and fault tolerance.

In the fall of 2016, ICL won an array of seven awards from the DOE's Exascale Computing Project (ECP). In doing so, ICL earned a place among an elite set of researchers from DOE laboratories who will create the software infrastructure for the nation's first exascale machines. On the following pages, we provide brief summaries of some of our efforts in these areas.

Distributed Tasking for Exascale

The Distributed Tasking for Exascale (DTE) project will extend the capabilities of ICL's Parallel Runtime and Execution Controller (PaRSEC) project—a generic framework for architecture-aware scheduling and management of microtasks on distributed, many-core, heterogeneous architectures. The PaRSEC environment also provides a runtime component for dynamically executing tasks on heterogeneous distributed systems along with a productivity toolbox and development framework that supports multiple domain-specific languages (DSLs) and extensions and tools for debugging, trace collection, and analysis.

PaRSEC also enables fast prototyping DSLs to express the dependencies between tasks and provides a stable, scalable, and efficient distributed runtime so they can run on any execution platform at any scale. The underlying dataflow paradigm attacks both sides of the exascale challenge: managing extreme-scale parallelism and maintaining the performance portability of the code. The DTE award is a vital extension and continuation of this effort and will ensure that PaRSEC meets the critical needs of ECP application communities in terms of scalability, interoperability, and productivity.

FIND OUT MORE AT **http://icl.utk.edu/dte/**

Exascale Performance Application Programming Interface

The Exascale Performance Application Programming Interface (Exa-PAPI) project is developing a new C++ Performance API (PAPI++) software package from the ground up that offers a standard interface and methodology for using low-level performance counters in CPUs, GPUs, on/off-chip memory, interconnects, and the I/O system—including energy/power management. PAPI++ is building upon classic-PAPI functionality and strengthening its path to exascale with a more efficient and flexible software design—a design that takes advantage of C++'s object-oriented nature but preserves the low-overhead monitoring of performance counters and adds a vast testing suite.

In addition to providing hardware counter–based information, a standardizing layer for monitoring software-defined events (SDE), which exposes the internal behavior of runtime systems and libraries (e.g., communication and math libraries) to the applications, is being incorporated. As a result, the notion of performance events is broadened from strictly hardware-related events to also include software-based information. Enabling monitoring of both hardware and software events provides more flexibility to developers when capturing performance information.

FIND OUT MORE AT **http://icl.utk.edu/exa-papi/**

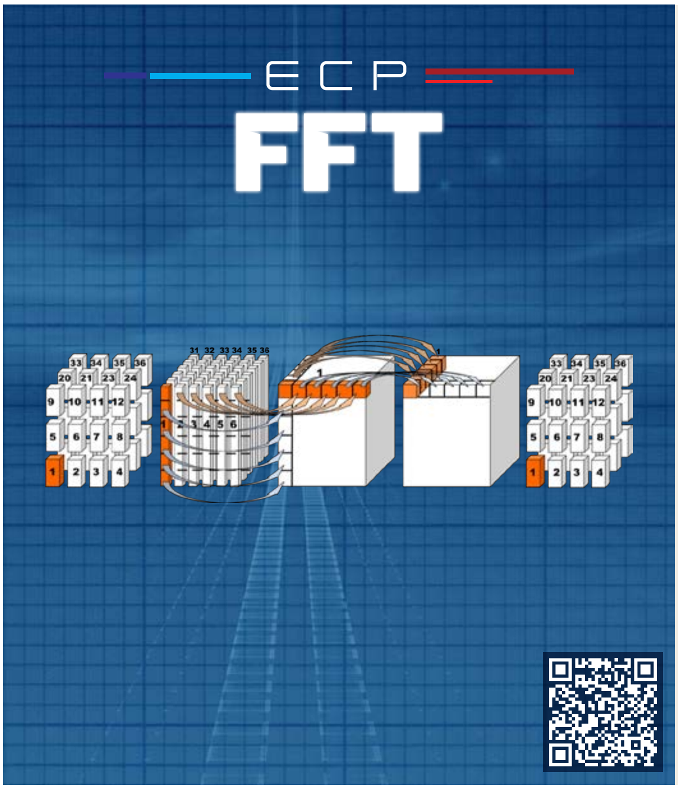

Fast Fourier Transform

The fast Fourier transform (FFT) is used in many domain applications—including molecular dynamics, spectrum estimation, fast convolution and correlation, signal modulation, and wireless multimedia applications, but current state-of-the-art FFT libraries are not scalable on large heterogeneous machines with many nodes.

The main objective of the FFT-ECP project is to design and implement a fast and robust 2-D and 3-D FFT library that targets large-scale heterogeneous systems with multi-core processors and hardware accelerators and to do so as a co-design activity with other ECP application developers. The work involves studying and analyzing current FFT software from vendors and open-source developers in order to understand, design, and develop a 3-D FFT-ECP library that could benefit from these existing optimized FFT kernels or will rely on new optimized kernels developed under this framework. Version 0.1 of the Highly Efficient FFTs for Exascale (heFFTe) library was released in October 2019, marking the first release for ICL's FFT-ECP effort.

FIND OUT MORE AT **http://icl.utk.edu/fft/**

Production-ready, Exascale-Enabled Krylov Solvers for Exascale Computing

Many large-scale scientific applications rely heavily on preconditioned iterative solvers for large linear systems. For these solvers to efficiently exploit extreme-scale hardware, both the solver algorithms and the implementations must be redesigned to address challenges like extreme concurrency, complex memory hierarchies, costly data movement, heterogeneous node architectures, and the increasing adoption of low-precision processor technology.

The PEEKS effort aims to tackle these challenges and advance the capabilities of the ECP software stack by making the new scalable algorithms accessible within the Trilinos and Ginkgo software ecosystems. Targeting exascale-enabled Krylov solvers, incomplete factorization routines, and parallel preconditioning techniques will ensure successful delivery of scalable Krylov solvers in robust, production-quality software that can be relied on by ECP applications.

FIND OUT MORE AT **http://icl.utk.edu/peeks/**

Software for Linear Algebra Targeting Exascale

For decades, ICL has applied algorithmic and technological innovations to the process of pioneering, implementing, and disseminating dense linear algebra software—including the Linear Algebra PACKage (LAPACK) and Scalable Linear Algebra PACKage (ScaLAPACK) libraries. The Software for Linear Algebra Targeting Exascale (SLATE) project will converge and consolidate that software into a dense linear algebra library that will integrate seamlessly into the ECP ecosystem.

For context, ScaLAPACK was first released in 1995, some 24 years ago. In the past two decades, HPC has witnessed tectonic shifts in the hardware technology, followed by paradigm shifts in the software technology, and a plethora of algorithmic innovations in scientific computing. At the same time, no viable replacement for ScaLAPACK emerged. SLATE is meant to be this replacement, boasting superior performance and scalability in the modern, heterogeneous, distributed-memory environments of HPC.

FIND OUT MORE AT **http://icl.utk.edu/slate/**

CEED

The Lawrence Livermore National Laboratory (LLNL)–led Center for Efficient Exascale Discretizations (CEED) co-design effort will develop next-generation discretization software and algorithms—which deliver a significant performance gain over conventional low-order methods—to enable a wide range of DOE and National Nuclear Security Administration (NNSA) applications to run efficiently on future exascale hardware. CEED is a research partnership involving 30+ computational scientists from two DOE labs and five universities, including UTK.

For UTK's part, ICL will be instrumental in identifying, developing, and optimizing tensor contractions that are essential building blocks for these kinds of DOE/NNSA applications. The ICL team will also play an integral role in co-designing application programming interfaces (APIs) with the LLNL scientists, external partners, and vendors, and will deliver a high-performance tensor contractions package through the Matrix Algebra on GPU and Multicore Architectures (MAGMA) library.

FIND OUT MORE AT **http://ceed.exascaleproject.org/**

OMPI-X

The Open MPI for Exascale (OMPI-X) project focuses on preparing the Message Passing Interface (MPI) standard—and its implementation in Open MPI—for exascale through improvements in scalability, capability, and resilience. Since its inception, the MPI standard has become ubiquitous in high-performance parallel computational science and engineering, and Open MPI is a widely used, high-quality, open-source implementation of the MPI standard. Despite their history and popularity, however, neither Open MPI nor the MPI standard itself is currently ready for the changes in hardware and software that will accompany exascale computing.

To mitigate this concern, OMPI-X will address a broad spectrum of issues in both the standard and the implementation by ensuring runtime interoperability for MPI+X and beyond, extending the MPI standard to better support coming exascale architectures, improving Open MPI scalability and performance, supporting more dynamic execution environments, enhancing resilience in MPI and Open MPI, evaluating MPI tools interfaces, and maintaining quality assurance.

FIND OUT MORE AT **http://www.icl.utk.edu/research/ompi-x/**

xSDK4ECP

The Extreme-Scale Scientific Software Development Kit for the Exascale Computing Project (xSDK4ECP) is a collaboration between Argonne National Laboratory (ANL), ICL, Lawrence Berkeley National Laboratory (LBNL), LLNL, Sandia National Laboratories (SNL), and the University of California at Berkeley (UCB). The project aims to enable seamless integration and combined use of diverse, independently developed software packages for ECP applications. Currently, this includes a wide range of high-quality software libraries and solver packages that address the strategic requirements of DOE's Office of Science.

ICL's MAGMA project was integrated into the xSDK 0.3 release, and ICL's Parallel Linear Algebra Software for Multi-core Architectures (PLASMA) package was added to xSDK 0.4. To ensure the consistency of naming conventions, runtime behavior, and installation procedure, xSDK informs the project development process by providing requirements and guidelines that are influential throughout the software development phase. xSDK lightens the burden on system administrators and application developers, because each xSDK package provides a Spack installation script that can be invoked independently or through the installation of the xSDK's Spack package—the default installation procedure since the 0.3 release.

FIND OUT MORE AT **https://xsdk.info/ecp/**

AsyncIS

The Asynchronous Iterative Solvers for Extreme-Scale Computing (AsyncIS) project aims to explore more efficient numerical algorithms by decreasing their overhead. AsyncIS does this by replacing the outer Krylov subspace solver with an asynchronous optimized Schwarz method, thereby removing the global synchronization and bulk synchronous operations typically used in numerical codes.

AsyncIS, a DOE-funded collaboration between Georgia Tech, UTK, Temple University, and SNL, also focuses on the development and optimization of asynchronous preconditioners (i.e., preconditioners that are generated and/or applied in an asynchronous fashion). The novel preconditioning algorithms that provide fine-grained parallelism enable preconditioned Krylov solvers to run efficiently on large-scale distributed systems and many-core accelerators like GPUs.

FIND OUT MORE AT **http://www.icl.utk.edu/research/asyncis/**

BATCHED BLAS

The Batched Basic Linear Algebra Subprograms (BBLAS) effort, an international collaboration between INRIA, Rutherford Appleton Laboratory, Umeå University, the University of Manchester, and UTK, will create an API for numerical computing routines that process batches of either uniformly sized or varying-size matrices or vectors. This will go beyond the original Basic Linear Algebra Subprogram (BLAS) standard by specifying a programming interface for modern scientific applications, which produce large numbers of small matrices at once.

Individually, the small sizes of the inputs obviate the potential benefits of using BLAS but are a perfect fit for BBLAS. The BBLAS project will also serve as a working forum for establishing the consensus for the next official standard that will serve the scientific community and ensure support from hardware vendors.

FIND OUT MORE AT **http://icl.utk.edu/bblas/**

BONSAI

The goal of the Benchtuning OpeN Software Autotuning Infrastructure (BONSAI) project is to develop a software infrastructure for using scalable, parallel hybrid systems to carry out large, concurrent autotuning sweeps in order to dramatically accelerate the optimization process of computational kernels for GPU accelerators and many-core coprocessors.

Recent developments include a distributed benchmarking engine—already released and currently under active development—capable of scaling to tens of thousands of nodes. This enables benchmarking millions of kernel configurations for different problem sizes and many input datasets while collecting hundreds of performance metrics—including time, energy consumption, cache misses, and memory bandwidth.

FIND OUT MORE AT **http://icl.utk.edu/bonsai/**

CAARES

The Cross-layer Application-Aware Resilience at Extreme Scale (CAARES) project, a collaborative effort between ICL, Rutgers University, and Stony Brook, aims to provide a theoretical foundation for multi-level fault management techniques and provide a clear understanding of existing obstacles that could obstruct generic and efficient approaches for fault management at scale. This effort is vital for large-scale science, because, as extreme-scale computational power enables new and important discoveries across all science domains, the current understanding of fault rates is casting a grim shadow and revealing a future where failures are not exceptions but are the norm.

By studying combinations of fault tolerance techniques instead of studying them in isolation from each other, CAARES seizes the opportunity to identify moldable techniques at the frontier of known approaches and highlight a composition of methodologies that inherit their individual benefits but do not exhibit their drawbacks, leading to the development of resilience techniques able to bridge the gap between fault tolerance ergonomics and efficiency.

FIND OUT MORE AT **http://www.icl.utk.edu/research/caares/**

CORES

The Convex Optimization for Real-time Embedded Systems (CORES) project aims to develop highly efficient, real-time convex optimization algorithms and toolsets for solving important engineering problems on hierarchical and heterogeneous embedded system architectures. Though recent advances in optimization solvers have enabled the solution of optimization problems on low-cost embedded systems, the size of the problems that can be solved in real time is still limited.

The CORES project, a collaboration between ICL and Michigan Technological University, works to address this limitation. The ICL team's main responsibility is the design and development of higher-performance, structure-aware linear solvers that would enable us to solve, in real time, the convex optimization problems that have significantly higher performance—and are orders of magnitude greater in size—compared to current state-of-the-art solvers.

FIND OUT MORE AT **http://www.icl.utk.edu/research/cores/**

DESC

The Development of Exascale Software for Heterogeneous and Interfacial Catalysis (DESC) project focuses on understanding the relationship between algorithms and hardware platforms and how to jointly optimize the software and hardware to achieve efficient implementations for applications in materials science, chemistry, and physics. DESC is a joint effort between ICL/UTK, DOE's Ames Laboratory, EP Analytics, Inc., Georgia Tech, Old Dominion University, and Virginia Tech and is funded by the DOE Computational Chemical Sciences project.

ICL's contribution focuses on expressing GAMESS computational chemistry algorithms in the form of a dataflow graph and subsequently mapping the DAG representation of the kernels to the hardware platforms. This representation allows for capturing the essential properties of the algorithms (e.g., data dependencies) and computation at extreme scale by utilizing the hardware components (e.g., CPU or GPU) best suited for the type of computational task under consideration. The dataflow-based form of these algorithms makes them compatible with next-generation task scheduling systems like PaRSEC, StarPU, and Legion.

FIND OUT MORE AT **https://www.icl.utk.edu/research/desc**

DPLASMA

The Distributed Parallel Linear Algebra Software for Multi-core Architectures (DPLASMA) package is the leading implementation of a dense linear algebra package for distributed heterogeneous systems. It is designed to deliver sustained performance for distributed systems, where each node features multiple sockets of multi-core processors and, if available, accelerators like GPUs or Intel Xeon Phi coprocessors. DPLASMA achieves this objective by deploying PLASMA algorithms on distributed-memory systems using the state-of-the-art PaRSEC runtime.

In addition to traditional ScaLAPACK data distribution, DPLASMA provides interfaces for users to expose arbitrary data distributions. The algorithms operate transparently on local data or introduce implicit communications to resolve dependencies, thereby removing the burden of initial data reshuffle and providing the user with a novel approach to address load balance.

FIND OUT MORE AT **http://icl.utk.edu/dplasma/**

EPEXA

A collaborative project involving Virginia Tech, Stony Brook, and ICL/UTK, the Ecosystem for Programming and Executing eXtreme Applications (EPEXA), aims to create a software framework that implements high-performance methods for irregular and dynamic computations that are poorly supported by current programming paradigms. Employing science-driven codesign, the EPEXA team will harden a successful research prototype into an accessible, production-quality programming model that will leverage domain-specific languages (DSLs) to improve accessibility and accelerate the adoption of high-performance tools for computer scientists and domain scientists.

The project bridges the so-called “valley of death” between a successful proof of concept and an implementation with enough quality, performance, and community support to motivate application scientists and other researchers to adopt it and push for its community use. Specifically, the new powerful data-flow programming model and associated parallel runtime directly address multiple challenges faced by scientists as they leverage rapidly changing computer technologies—including current, massively parallel, hybrid, and many-core systems.

FIND OUT MORE AT **http://www.icl.utk.edu/research/epexa**

Evolve

Evolve, a collaborative effort between ICL and the University of Houston, expands the capabilities of Open MPI to support the NSF’s critical software-infrastructure missions. Core challenges include: extending the software to scale to 10,000–100,000 processes; ensuring support for accelerators; enabling highly asynchronous execution of communication and I/O operations; and ensuring resilience. Part of the effort involves careful consideration of modifications to the MPI specification to account for the emerging needs of application developers on future extreme-scale systems.

So far, Evolve efforts have involved exploratory research for improving different performance aspects of the Open MPI library. Notably, this has led to an efficiency improvement in multi-threaded programs using MPI in combination with other thread-based programming models (e.g., OpenMP). A novel collective communication framework with event-based programming and data dependencies was investigated, and it demonstrated a clear advantage in terms of aggregate bandwidth in heterogeneous (shared memory + network) systems. Support for MPI resilience following the User-Level Failure Mitigation (ULFM) fault-tolerance proposal was released based on the latest Open MPI version and will soon be fully integrated into Open MPI.

FIND OUT MORE AT **https://www.icl.utk.edu/research/evolve**

HPCG

The High Performance Conjugate Gradients (HPCG) benchmark is designed to measure performance that is representative of modern scientific applications. It does so by exercising the computational and communication patterns commonly found in real science and engineering codes, which are often based on sparse iterative solvers. HPCG exhibits the same irregular accesses to memory and fine-grain recursive computations that dominate large-scale scientific workloads used to simulate complex physical phenomena.

The HPCG 3.1 reference code was released in March of 2019. In addition to bug fixes, this release positioned HPCG to even better represent modern partial differential equation (PDE) solvers and made it easier to run HPCG on production supercomputing installations. The reference version is accompanied by binary or source code releases from AMD, ARM, Intel, and NVIDIA, which are carefully optimized for the vendors' respective hardware platforms. The current HPCG performance list was released at SC19 and now features 160 supercomputing entries. HPCG rankings have also been tracked by TOP500.org since June of 2017.

FIND OUT MORE AT **http://www.hpcg-benchmark.org/**

HPL

The High Performance LINPACK (HPL) benchmark solves a dense linear system in double precision (64-bit) arithmetic on distributed-memory computers. HPL is written in a portable ANSI C and requires an MPI implementation and either BLAS or the Vector Signal and Image Processing Library (VSIPL). HPL is often one of the first programs to run on large HPC installations, producing a result that can be submitted to the TOP500 list of the world's fastest supercomputers.

The major focus of HPL 2.3, released in 2018, was to improve the accuracy of reported benchmark results and ensure easier configuration and building on modern HPC platforms. HPL now features more detailed reporting of the solution's scaled residual and of the achieved performance number. Another addition is a software configuration tool based on Autotools and the removal of deprecated MPI functions. The LINPACK app for iOS achieved over 8 GFLOP/s on the iPhone X. For the November 2019 TOP500 list, an optimized version of the HPL code achieved nearly 150 PFLOP/s on the Summit supercomputer at ORNL.

FIND OUT MORE AT **http://icl.utk.edu/hpl/**

HPL-AI

The HPL-AI benchmark seeks to highlight the convergence of HPC and artificial intelligence (AI) workloads based on machine learning (ML) and deep learning (DL) by solving a system of linear equations using novel, mixed-precision algorithms that exploit modern hardware. While traditional HPC focuses on simulation runs for modeling phenomena in a variety of scientific disciplines, the mathematical models that drive these computations require, for the most part, 64-bit accuracy. On the other hand, the ML/DL methods that fuel advances in AI can achieve the desired results at 32-bit or even lower floating-point precisions. This lesser demand for accuracy fueled a resurgence of interest in new hardware platforms that deliver a mix of unprecedented performance levels and energy savings to achieve the classification and recognition fidelity afforded by higher-accuracy formats.

HPL-AI strives to unite these two realms by connecting its solver formulation to the decades-old HPL framework of benchmarking supercomputing installations. So far, ORNL's Summit is the only machine to be benchmarked at scale with HPL-AI, and it achieved 445 PFLOP/s in mixed precision. This is nearly triple the 148 PFLOP/s that Summit achieved on the standard (double-precision) HPL benchmark used for the TOP500.

FIND OUT MORE AT **https://icl.bitbucket.io/hpl-ai/**

LAPACK AND ScaLAPACK

The Linear Algebra PACKage (LAPACK) and Scalable LAPACK (ScaLAPACK) are widely used libraries for efficiently solving dense linear algebra problems. ICL has been a major contributor to the development and maintenance of these two packages since their inception. LAPACK is sequential, relies on the BLAS library, and benefits from the multi-core BLAS library. ScaLAPACK is parallel, distributed, and relies on the BLAS, LAPACK, MPI, and BLACS libraries.

LAPACK 3.9.0, released in November 2019, adds a QR-preconditioned QR SVD method and an LAPACK Householder reconstruction routine. Since 2011, LAPACK has included LAPACKE, a native C interface for LAPACK developed in collaboration with Intel, which provides NAN check and automatic workspace allocation. ScaLAPACK 2.1.0, which includes a new robust ScaLAPACK routine for computing the QR factorization with column pivoting along with improved accuracy of the Frobenius norm, was released in November 2019.

FIND OUT MORE AT **http://www.netlib.org/lapack/** AND **http://www.netlib.org/scalapack/**

MAGMA

Matrix Algebra on GPU and Multi-core Architectures (MAGMA) is a collection of next-generation linear algebra libraries for heterogeneous computing. The MAGMA package supports interfaces for current linear algebra packages and standards (e.g., LAPACK and BLAS) to enable computational scientists to easily port any linear algebra–reliant software components to heterogeneous computing systems. MAGMA enables applications to fully exploit the power of current hybrid systems of many-core CPUs and multi-GPUs/coprocessors to deliver the fastest possible time to accurate solution within given energy constraints.

MAGMA 2.5.2 features LAPACK-compliant routines for multi-core CPUs enhanced with NVIDIA GPUs (including the Volta V100). MAGMA now includes more than 400 routines that cover one-sided dense matrix factorizations and solvers, two-sided factorizations, and eigen/singular-value problem solvers, as well as a subset of highly optimized BLAS for GPUs. A MagmaDNN package has been added and further enhanced to provide high-performance data analytics, including functionalities for machine learning applications that use MAGMA as their computational back end. The MAGMA Sparse and MAGMA Batched packages have been included since MAGMA 1.6.

FIND OUT MORE AT **http://icl.utk.edu/magma/**

MATEDOR

The MAtrix, TEnsor, and Deep-learning Optimized Routines (MATEDOR) project will perform the research required to define a standard interface for batched operations and provide a performance-portable software library that demonstrates batching routines for a significant number of kernels. This research is critical, given that the performance opportunities inherent in solving many small batched matrices often yield more than a 10× speedup over the current classical approaches.

Working closely with affected application communities, along with ICL's Batched BLAS initiative, MATEDOR will define modular, optimizable, and language-agnostic interfaces that can work seamlessly with a compiler. This modularity will provide application, compiler, and runtime system developers with the option to use a single call to a routine from the new batch operation standard and would allow the entire linear algebra community to collectively attack a wide range of small matrix or tensor problems.

FIND OUT MORE AT **http://www.icl.utk.edu/research/matedor/**

Open MPI

The Open MPI Project is an open-source Message Passing Interface (MPI) implementation developed and maintained by a consortium of academic, research, and industry partners. MPI primarily addresses the message-passing parallel programming model, in which data is moved from the address space of one process to that of another process through cooperative operations on each process. Open MPI integrates technologies and resources from several other projects (e.g., HARNESS/FT-MPI, LA-MPI, LAM/MPI, and PACX-MPI) in order to build the best MPI library available.

A completely new MPI 3.2–compliant implementation, Open MPI offers advantages for system and software vendors, application developers, and computer science researchers. ICL's efforts in the context of Open MPI have significantly improved its scalability, performance on many-core environments, and architecture-aware capabilities—such as adaptive shared memory behaviors and dynamic collective selection—making it ready for next-generation exascale challenges.

FIND OUT MORE AT **https://www.open-mpi.org/**

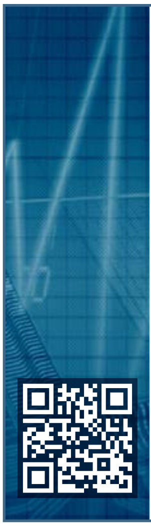

PAPI

The Performance Application Programming Interface (PAPI) supplies a consistent interface and methodology for collecting performance counter information from various hardware and software components, including most major CPUs, GPUs and accelerators, interconnects, I/O systems, and power interfaces, as well as virtual cloud environments. Industry liaisons with AMD, Cray, Intel, IBM, NVIDIA, and others ensure seamless integration of PAPI with new architectures at or near their release. As the PAPI component architecture becomes more populated, performance tools that interface with PAPI automatically inherit the ability to measure these new data sources.

In 2019, ICL, together with the University of Maine, worked on PAPI-ex to build support for performance counters available in the latest generations of CPUs and GPUs, develop support for system-wide hardware performance counter monitoring, and strengthen the sampling interface in PAPI. As of November 2019, PAPI-ex also incorporates a counter analysis toolkit (CAT) designed to improve the understanding of low-level hardware events. Since 2016, the PAPI effort has been bolstered through the ECP Exa-PAPI project.

FIND OUT MORE AT **http://icl.utk.edu/papi/**

PaRSEC

The Parallel Runtime Scheduling and Execution Controller (PaRSEC) is a generic framework for architecture-aware scheduling and management of microtasks on distributed many-core heterogeneous architectures. Applications considered are expressed as a DAG of tasks with edges designating the data dependencies. DAGs are represented in a compact, problem-size independent format that can be queried to discover data dependencies in a totally distributed fashion—a drastic shift from today's programming models, which are based on sequential flow of execution.

PaRSEC orchestrates the execution of an algorithm on a particular set of resources, assigns computational threads to the cores, overlaps communications and computations, and uses a dynamic, fully distributed scheduler. PaRSEC includes a set of tools to generate the DAGs and integrate them into legacy codes, a runtime library to schedule the microtasks on heterogeneous resources, and tools to evaluate and visualize the efficiency of the scheduling. Many dense and sparse linear algebra extensions have been implemented, as well as chemistry and seismology applications, which produced significant speedup in production codes.

FIND OUT MORE AT **http://icl.utk.edu/parsec/**

PLASMA

The Parallel Linear Algebra Software for Multi-core Architectures (PLASMA) implements a set of fundamental linear algebra routines using the latest updates to the Open Multi-Processing (OpenMP) standard. PLASMA includes, among others, routines for solving linear systems of equations, linear least square problems, parallel BLAS, and parallel matrix norms.

Over the last decade, PLASMA has been used on a variety of systems using Intel CPUs and coprocessors, IBM POWER processors, and ARM processors. As a research vehicle, PLASMA is an example of modern design for new dense linear algebra algorithms. This has paved the way for distributed-memory developments, including the new ECP SLATE project, which will ultimately deliver these capabilities at exascale. At the same time, PLASMA benefits from the continuous evolution of the OpenMP standard that now includes off-load functionality and enables porting to hardware accelerators.

FIND OUT MORE AT **https://bitbucket.org/icl/plasma/**

PULSE

The PAPI Unifying Layer for Software-defined Events (PULSE) project focuses on enabling cross-layer and integrated monitoring of whole application performance by extending PAPI with the capability to expose performance metrics from key software components found in the HPC software stack. Up to this point, the abstraction and standardization layer provided by PAPI has been limited to profiling information generated by hardware only. Information about the behavior of the underlying software stack had to be acquired either through low-level binary instrumentation or through custom APIs.

To overcome this shortfall, PULSE is extending the abstraction and unification layer that PAPI has provided for hardware events to also encompass software events. On one end, PULSE offers a standard, well-defined and well-documented API that high-level profiling software can utilize to acquire performance information about the libraries used by an application and present it to the application developers. On the other end, it provides standard APIs that library and runtime writers can utilize to communicate information about the behavior of their software to higher software layers.

FIND OUT MORE AT **https://www.icl.utk.edu/research/pulse/**

ScaRT

The Scalable Run Time for Highly Parallel, Heterogeneous Systems (ScaRT) project aims to increase the scientific throughput of existing and future cyberinfrastructure platforms by reducing communication overheads; improving the match between modern, parallel-computing frameworks and the applications running on top; and by better matching the functionality of the underlying communication library to the capabilities of modern communication adapters.

To this end, SCaRT brings together a multidisciplinary team to (1) design and implement a communication library with new communication primitives; (2) accelerate multiple task-based runtimes (e.g., Legion and PaRSEC) and communication libraries (e.g., MPI and GasNET); (3) port key components to a programmable NIC; and (4) deliver improvements and extensions to mainstream communication libraries to provide the new functionality.

FIND OUT MORE AT **https://www.icl.utk.edu/research/scart/**

SMURFS

The Simulation and Modeling for Understanding Resilience and Faults at Scale (SMURFS) project seeks to acquire the predictive understanding of the complex interactions of a given application, a given real or hypothetical hardware and software environment, and a given fault-tolerance strategy at extreme scale.

SMURFS is characterized by two facets: (1) medium and fine-grained predictive capabilities and (2) coarse-grained fault tolerance strategy selection. Accordingly, ICL plans to design, develop, and validate new analytical and system component models that use semi-detailed software and hardware specifications to predict application performance in terms of time to solution and energy consumption. Also, based on a comprehensive set of studies using several application benchmarks, proxies, full applications, and several different fault tolerance strategies, ICL will gather valuable insights about application behavior at scale.

FIND OUT MORE AT **https://www.icl.utk.edu/research/smurfs/**

TOP500

With three decades of tracking the progress of supercomputing, the TOP500 list continues to provide a reliable historical record of supercomputers around the world. The list clearly lays out critical HPC metrics across all 500 machines and draws a rich picture of the state of the art in terms of performance, energy consumption, and power efficiency. The TOP500 now features an HPCG ranking, which measures a machine's performance using irregular accesses to memory and fine-grain recursive computations—factors that dominate real-world, large-scale scientific workloads.

In November 2019, the 54th TOP500 list was unveiled at SC19 in Denver, CO. The United States remained on top with ORNL's Summit machine. Summit updated its HPL benchmark result for the November 2019 list and achieved 148.6 PFLOP/s (vs. 122.3 PFLOP/s in June 2018).

FIND OUT MORE AT **https://www.top500.org/**

ULFM

User Level Failure Mitigation (ULFM) is a set of new interfaces for MPI that enables message passing applications to restore MPI functionality affected by process failures. The MPI implementation is spared the expense of internally taking protective and corrective automatic actions against failures. Instead, it can prevent any fault-related deadlock situation by reporting operations wherein the completions were rendered impossible by failures.

Using the constructs defined by ULFM, applications and libraries drive the recovery of the MPI state. Consistency issues resulting from failures are addressed according to an application's needs, and the recovery actions are limited to the necessary MPI communication objects. A wide range of application types and middlewares are already building on top of ULFM to deliver scalable and user-friendly fault tolerance. Notable recent additions include the CoArray Fortran language and SAP databases. ULFM software is available in recent versions of both MPICH and Open MPI.

FIND OUT MORE AT **http://fault-tolerance.org/**

PUBLICATIONS

Evidence of our research and our contributions to the HPC community might be best exemplified by the numerous publications we produce every year. Here is a listing of our most recent papers, including journal articles, book chapters, and conference proceedings. Many of these are available for download from our website.

Abdelfattah, A., S. Tomov, and J. Dongarra, **"Towards Half-Precision Computation for Complex Matrices: A Case Study for Mixed Precision Solvers on GPUs,"** *ScalA19: 10th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems*, Denver, CO, IEEE, November 2019.

Anzt, H., G. Flegar, T. Gruetzmacher, and E. S. Quintana-Ortí, **"Toward a Modular Precision Ecosystem for High-Performance Computing,"** *The International Journal of High Performance Computing Applications*, vol. 33, issue 6, pp. 1069–1078, November 2019.

Anzt, H., T. Cojean, and E. Kuhn, **"Towards a New Peer Review Concept for Scientific Computing ensuring Technical Quality, Software Sustainability, and Result Reproducibility,"** *Proceedings in Applied Mathematics and Mechanics*, vol. 19, issue 1, November 2019.

Aupy, G., A. Benoit, B. Goglin, L. Pottier, and Y. Robert, **"Co-Scheduling HPC Workloads on Cache-Partitioned CMP Platforms,"** *International Journal of High Performance Computing Applications*, vol. 33, issue 6, pp. 1221–1239, November 2019.

Ayala, A., S. Tomov, X. Luo, H. Shaiek, A. Haidar, G. Bosilca, and J. Dongarra, **"Impacts of Multi-GPU MPI Collective Communications on Large FFT Computation,"** *Workshop on Exascale MPI (ExaMPI) at SC19*, Denver, CO, November 2019.

Benoit, A., T. Herault, V. Le Fèvre, and Y. Robert, **"Replication is More Efficient Than You Think,"** *The IEEE/ACM Conference on High Performance Computing Networking, Storage and Analysis (SC19)*, Denver, CO, ACM Press, November 2019.

Cao, Q., Y. Pei, T. Herault, K. Akbudak, A. Mikhalev, G. Bosilca, H. Ltaief, D. Keyes, and J. Dongarra, **"Performance Analysis of Tile Low-Rank Cholesky Factorization Using PaRSEC Instrumentation Tools,"** *Workshop on Programming and Performance Visualization Tools (ProTools 19) at SC19*, Denver, CO, ACM, November 2019.

Gates, M., J. Kurzak, A. Charara, A. YarKhan, and J. Dongarra, **"SLATE: Design of a Modern Distributed and Accelerated Linear Algebra Library,"** *International Conference for High Performance Computing, Networking, Storage and Analysis (SC19)*, Denver, CO, ACM, November 2019.

Han, L., V. Le Fèvre, L-C. Canon, Y. Robert, and F. Vivien, **"A Generic Approach to Scheduling and Checkpointing Workflows,"** *International Journal of High Performance Computing Applications*, vol. 33, issue 6, pp. 1255–1274, November 2019.

Herault, T., Y. Robert, G. Bosilca, and J. Dongarra, **"Generic Matrix Multiplication for Multi-GPU Accelerated Distributed-Memory Platforms over PaRSEC,"** *ScalA'19: 10th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems*, Denver, CO, IEEE, November 2019.

Jagode, H., A. Danalis, H. Anzt, and J. Dongarra, **"PAPI Software-Defined Events for in-Depth Performance Analysis,"** *The International Journal of High Performance Computing Applications*, vol. 33, issue 6, pp. 1113–1127, Nov. 2019.

Pei, Y., G. Bosilca, I. Yamazaki, A. Ida, and J. Dongarra, **"Evaluation of Programming Models to Address Load Imbalance on Distributed Multi-Core CPUs: A Case Study with Block Low-Rank Factorization,"** *PAW-ATM Workshop at SC19*, Denver, CO, ACM, November 2019.

Ribizel, T., and H. Anzt, **"Parallel Selection on GPUs,"** *Parallel Computing*, November 2019.

Altintas, I., K. Marcus, V. Vural, S. Purawat, D. Crawl, G. Antoniu, A. Costan, O. Marcu, P. Balaprakash, R. Cao, et al., **"A Collection of White Papers from the BDEC2 Workshop in San Diego, CA,"** *Innovative Computing Laboratory Technical Report*, no. ICL-UT-19-13: University of Tennessee, October 2019.

Bujanovic, Z., and Z. Drmac, **"New Robust ScaLAPACK Routine for Computing the QR Factorization with Column Pivoting,"** *LAPACK Working Note*, no. LAWN 296, ICL-UT-19-14: University of Tennessee, October 2019.

Tomov, S., A. Abdelfattah, V. Barra, N. Beams, J. Brown, J-S. Camier, V. Dobrev, J. Dongarra, Y. Dudouit, P. Fischer, et al., **"CEED ECP Milestone Report: Performance Tuning of CEED Software and 1st and 2nd Wave Apps,"** Zenodo, October 2019.

Tomov, S., A. Haidar, A. Ayala, H. Shaiek, and J. Dongarra, **"FFT-ECP Implementation Optimizations and Features Phase,"** *Innovative Computing Laboratory Technical Report*, no. ICL-UT-19-12: University of Tennessee, October 2019.

Abdelfattah, A., S. Tomov, and J. Dongarra, **"Progressive Optimization of Batched LU Factorization on GPUs,"** *IEEE High Performance Extreme Computing Conference (HPEC'19)*, Waltham, MA, IEEE, September 2019.

Beck, M., T. Moore, N. French, E. Kissel, and M. Swany, **"Data Logistics: Toolkit and Applications,"** *5th EAI International Conference on Smart Objects and Technologies for Social Good*, Valencia, Spain, September 2019.

Davis, J., T. Gao, S. Chandrasekaran, H. Jagode, A. Danalis, P. Balaji, J. Dongarra, and M. Taufer, **"Characterization of Power Usage and Performance in Data-Intensive Applications using MapReduce over MPI,"** *2019 International Conference on Parallel Computing (ParCo2019)*, Prague, Czech Republic, September 2019.

Gamblin, T., P. Beckman, K. Keahey, K. Sato, M. Kondo, and G. Balazs, **"BDEC2 Platform White Paper,"** *Innovative Computing Laboratory Technical Report*, no. ICL-UT-19-11: University of Tennessee, September 2019.

Gao, Y., L-C. Canon, Y. Robert, and F. Vivien, **"Scheduling Independent Stochastic Tasks on Heterogeneous Cloud Platforms,"** *IEEE Cluster 2019*, Albuquerque, New Mexico, IEEE Computer Society Press, September 2019.

Gates, M., A. Charara, J. Kurzak, D. Sukkari, A. YarKhan, and J. Dongarra, **"SLATE Working Note 13: Implementing Singular Value and Symmetric Eigenvalue Solvers,"** *SLATE Working Notes*, no. 13, ICL-UT-19-07: Innovative Computing Laboratory, University of Tennessee, September 2019.

Luszczek, P., I. Yamazaki, and J. Dongarra, **"Increasing Accuracy of Iterative Refinement in Limited Floating-Point Arithmetic on Half-Precision Accelerators,"** *IEEE High Performance Extreme Computing Conference (HPEC 2019)*, Best Paper Finalist, Waltham, MA, IEEE, September 2019.

Patinyasakdikul, T., D. Eberius, G. Bosilca, and N. Hjelm, **"Give MPI Threading a Fair Chance: A Study of Multithreaded MPI Designs,"** *IEEE Cluster*, Albuquerque, NM, IEEE, September 2019.

YarKhan, A., J. Kurzak, A. Abdelfattah, and J. Dongarra, **"An Empirical View of SLATE Algorithms on Scalable Hybrid System,"** *Innovative Computing Laboratory Technical Report*, no. ICL-UT-19-08: University of Tennessee, Knoxville, September 2019.

Zhong, D., A. Bouteiller, X. Luo, and G. Bosilca, **"Runtime Level Failure Detection and Propagation in HPC Systems,"** *European MPI Users' Group Meeting (EuroMPI '19)*, Zürich, Switzerland, ACM, September 2019.

Kurzak, J., Y. Tsai, M. Gates, A. Abdelfattah, and J. Dongarra, **"Massively Parallel Automated Software Tuning,"** *48th International Conference on Parallel Processing (ICPP 2019)*, Kyoto, Japan, ACM Press, August 2019.

Kurzak, J., M. Gates, A. Charara, A. YarKhan, I. Yamazaki, and J. Dongarra, **"Linear Systems Solvers for Distributed-Memory Machines with GPU Accelerators,"** *Euro-Par 2019: Parallel Processing*, vol. 11725: Springer, pp. 495–506, August 2019.

Yamazaki, I., E. Chow, A. Bouteiller, and J. Dongarra, **"Performance of Asynchronous Optimized Schwarz with One-sided Communication,"** *Parallel Computing*, vol. 86, pp. 66–81, August 2019.

Yamazaki, I., A. Ida, R. Yokota, and J. Dongarra, **"Distributed-Memory Lattice H-Matrix Factorization,"** *The International Journal of High Performance Computing Applications*, vol. 33, issue 5, pp. 1046–1063, August 2019.

Betancourt, F., K. Wong, E. Asemota, Q. Marshall, D. Nichols, and S. Tomov, **"OpenDIEL: A Parallel Workflow Engine and DataAnalytics Framework,"** *Practice and Experience in Advanced Research Computing (PEARC '19)*, Chicago, IL, ACM, July 2019.

Jagode, H., A. Danalis, and J. Dongarra, **"What it Takes to keep PAPI Instrumental for the HPC Community,"** *1st Workshop on Sustainable Scientific Software (CW3S19)*, Collegeville, Minnesota, July 2019.

Le Fèvre, V., T. Herault, Y. Robert, A. Bouteiller, A. Hori, G. Bosilca, and J. Dongarra, **"Comparing the Performance of Rigid, Moldable, and Grid-Shaped Applications on Failure-Prone HPC Platforms,"** *Parallel Computing*, vol. 85, pp. 1–12, July 2019.

Lopez, M. G., W. Joubert, V. Larrea, O. Hernandez, A. Haidar, S. Tomov, and J. Dongarra, **"Evaluation of Directive-Based Performance Portable Programming Models,"** *International Journal of High Performance Computing and Networking*, vol. 14, issue 2, pp. 165–182, July 2019.

Nichols, D., K. Wong, S. Tomov, L. Ng, S. Chen, and A. Gessinger, **"MagmaDNN: Accelerated Deep Learning Using MAGMA,"** *Practice and Experience in Advanced Research Computing (PEARC '19)*, Chicago, IL, ACM, July 2019.

Anzt, H., Y. Chen Chen, T. Cojean, J. Dongarra, G. Flegar, P. Nayak, E. S. Quintana-Ortí, Y. M. Tsai, and W. Wang, **"Towards Continuous Benchmarking,"** *Platform for Advanced Scientific Computing Conference (PASC 2019)*, Zürich, Switzerland, ACM Press, June 2019.

Canon, L-C., A K W. Chang, Y. Robert, and F. Vivien, **"Scheduling Independent Stochastic Tasks under Deadline and Budget Constraints,"** *International Journal of High Performance Computing Applications*, June 2019.

Dongarra, J., M. Gates, A. Haidar, J. Kurzak, P. Luszczek, P. Wu, I. Yamazaki, A. YarKhan, M. Abalenkovs, N. Bagherpour, et al., **"PLASMA: Parallel Linear Algebra Software for Multicore Using OpenMP,"** *ACM Transactions on Mathematical Software*, vol. 45, issue 2, June 2019.

Kurzak, J., M. Gates, A. Charara, A. YarKhan, and J. Dongarra, **"Least Squares Solvers for Distributed-Memory Machines with GPU Accelerators,"** *ACM International Conference on Supercomputing (ICS '19)*, Phoenix, Arizona, ACM, pp. 117–126, June 2019.

Kurzak, J., M. Gates, A. Charara, A. YarKhan, and J. Dongarra, **"SLATE Working Note 12: Implementing Matrix Inversions,"** *SLATE Working Notes*, no. 12, ICL-UT-19-04: Innovative Computing Laboratory, University of Tennessee, June 2019.

Nichols, D., N-S. Tomov, F. Betancourt, S. Tomov, K. Wong, and J. Dongarra, **"MagmaDNN: Towards High-Performance Data Analytics and Machine Learning for Data-Driven Scientific Computing,"** *ISC High Performance*, Frankfurt, Germany, Springer International Publishing, June 2019.

Wong, K., S. Tomov, and J. Dongarra, **"Hands-on Research and Training in High-Performance Data Sciences, Data Analytics, and Machine Learning for Emerging Environments,"** *ISC High Performance*, Frankfurt, Germany, Springer International Publishing, June 2019.

Abdelfattah, A., S. Tomov, and J. Dongarra, **"Fast Batched Matrix Multiplication for Small Sizes using Half Precision Arithmetic on GPUs,"** *33rd IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Rio de Janeiro, Brazil, IEEE, May 2019.

Antoniu, G., A. Costan, O. Marcu, M. S. Pérez, N. Stojanovic, R. M. Badia, M. Vázquez, S. Girona, M. Beck, T. Moore, et al., **"A Collection of White Papers from the BDEC2 Workshop in Poznan, Poland,"** *Innovative Computing Laboratory Technical Report*, no. ICL-UT-19-10: University of Tennessee, Knoxville, May 2019.

Anzt, H., T. Ribizel, G. Flegar, E. Chow, and J. Dongarra, **"ParILUT – A Parallel Threshold ILU for GPUs,"** *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Rio de Janeiro, Brazil, IEEE, May 2019.

Anzt, H., and G. Flegar, **"Are we Doing the Right Thing? – A Critical Analysis of the Academic HPC Community,"** *2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, Rio de Janeiro, Brazil, IEEE, May 2019.

Aupy, G., A. Gainaru, V. Honoré, P. Raghavan, Y. Robert, and H. Sun, **"Reservation Strategies for Stochastic Jobs,"** *33rd IEEE International Parallel and Distributed Processing Symposium (IPDPS 2019)*, Rio de Janeiro, Brazil, IEEE Computer Society Press, May 2019.

Bai, Z., J. Dongarra, D. Lu, and I. Yamazaki, **"Matrix Powers Kernels for Thick-Restart Lanczos with Explicit External Deflation,"** *International Parallel and Distributed Processing Symposium (IPDPS)*, Rio de Janeiro, Brazil, IEEE, May 2019.

Danalis, A., H. Jagode, T. Herault, P. Luszczek, and J. Dongarra, **"Software-Defined Events through PAPI,"** *2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, Rio de Janeiro, Brazil, IEEE, May 2019.

Kaya, O., and Y. Robert, **"Computing Dense Tensor Decompositions with Optimal Dimension Trees,"** *Algorithmica*, vol. 81, issue 5, pp. 2092–2121, May 2019.

Ribizel, T., and H. Anzt, **"Approximate and Exact Selection on GPUs,"** *2019 IEEE International Parallel and Distributed Processing Symposium Workshops*, Rio de Janeiro, Brazil, IEEE, May 2019.

Brown, J., A. Abdelfattah, V. Barra, V. Dobrev, Y. Dudouit, P. Fischer, T. Kolev, D. Medina, M. Min, T. Ratnayaka, et al., **"CEED ECP Milestone Report: Public Release of CEED 2.0,"** Zenodo, April 2019.

Charara, A., J. Dongarra, M. Gates, J. Kurzak, and A. YarKhan, **"SLATE Mixed Precision Performance Report,"** *Innovative Computing Laboratory Technical Report*, no. ICL-UT-19-03: University of Tennessee, April 2019.

Tomov, S., A. Haidar, A. Ayala, D. Schultz, and J. Dongarra, **"Design and Implementation for FFT-ECP on Distributed Accelerated Systems,"** *Innovative Computing Laboratory Technical Report*, no. ICL-UT-19-05: University of Tennessee, April 2019.

Anzt, H., J. Dongarra, G. Flegar, N. J. Higham, and E. S. Quintana-Ortí, **"Adaptive Precision in Block-Jacobi Preconditioning for Iterative Sparse Linear System Solvers,"** *Concurrency and Computation: Practice and Experience*, vol. 31, no. 6, pp. e4460, March 2019.

Dongarra, J., S. Gottlieb, and W. T. Kramer, **"Race to Exascale,"** *Computing in Science and Engineering*, vol. 21, issue 1, pp. 4–5, March 2019.

Badia, R. M., M. Beck, F. Bodin, T. Boku, F. Cappello, A. Choudhary, C. Costa, E. Deelman, N. Ferrier, K. Fujisawa, et al., **"A Collection of Presentations from the BDEC2 Workshop in Kobe, Japan,"** *Innovative Computing Laboratory Technical Report*, no. ICL-UT-19-09: University of Tennessee, Knoxville, February 2019.

Danalis, A., H. Jagode, H. Hanumantharayappa, S. Ragate, and J. Dongarra, **"Counter Inspection Toolkit: Making Sense out of Hardware Performance Events,"** *11th International Workshop on Parallel Tools for High Performance Computing*, Dresden, Germany, Cham, Switzerland: Springer, February 2019.

Losada, N., G. Bosilca, A. Bouteiller, P. González, and M. J. Martín, **"Local Rollback for Resilient MPI Applications with Application-Level Checkpointing and Message Logging,"** *Future Generation Computer Systems*, vol. 91, pp. 450–464, February 2019.

Charara, A., M. Gates, J. Kurzak, and J. Dongarra, **"SLATE Developers' Guide,"** *SLATE Working Notes*, no. 11, ICL-UT-19-02: Innovative Computing Laboratory, University of Tennessee, January 2019.

Gates, M., A. Charara, J. Kurzak, and J. Dongarra, **"SLATE Users' Guide,"** *SLATE Working Notes*, no. 10, ICL-UT-19-01: Innovative Computing Laboratory, University of Tennessee, January 2019.

Gruetzmacher, T., T. Cojean, G. Flegar, F. Göbel, and H. Anzt, **"A Customized Precision Format Based on Mantissa Segmentation for Accelerating Sparse Linear Algebra,"** *Concurrency and Computation: Practice and Experience*, vol. 40319, issue 262, January 2019.

Masliah, I., A. Abdelfattah, A. Haidar, S. Tomov, M. Baboulin, J. Falcou, and J. Dongarra, **"Algorithms and Optimization Techniques for High-Performance Matrix-Matrix Multiplications of Very Small Matrices,"** *Parallel Computing*, vol. 81, pp. 1–21, January 2019.

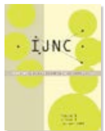

Benoit, A., A. Cavelan, F. M. Ciorba, V. Le Fèvre, and Y. Robert, **"Combining Checkpointing and Replication for Reliable Execution of Linear Workflows with Fail-Stop and Silent Errors,"** *International Journal of Networking and Computing*, vol. 9, no. 1, pp. 2–27.

Herault, T., Y. Robert, A. Bouteiller, D. Arnold, K. Ferreira, G. Bosilca, and J. Dongarra, **"Checkpointing Strategies for Shared High-Performance Computing Platforms,"** *International Journal of Networking and Computing*, vol. 9, no. 1, pp. 28–52, 2019.

Hori, A., Y. Tsujita, A. Shimada, K. Yoshinaga, N. Mitaro, G. Fukazawa, M. Sato, G. Bosilca, A. Bouteiller, and T. Herault, **"System Software for Many-Core and Multi-core Architectures,"** *Advanced Software Technologies for Post-Peta Scale Computing: The Japanese Post-Peta CREST Research Project*, Singapore, Springer Singapore, pp. 59–75, 2019.

FIND MORE AT **www.icl.utk.edu/publications**

EVENTS

JANUARY 14–18 HOUSTON, TX
2019 ECP Annual Meeting

JANUARY 29–30 ROCKVILLE, MD
The DOE ASCR Applied Mathematics PI meeting

FEBRUARY 19–21 KOBE, JAPAN
BDEC2

FEBRUARY 25–MARCH 1 SPOKANE, WA
SIAM Conference on Computational Science and Engineering

MARCH 4–8 CHATTANOOGA, TN
MPI Forum

MARCH 11–15 SANTA FE, NM
GPU Hackathon Santa Fe

MARCH 14–15 SAN FRANCISCO, CA
Future Information and Communication Conference

MARCH 18–21 SAN JOSE, CA
GPU Technology Conference (GTC 2019)

MARCH 25–26 OAK RIDGE, TN
7th ADAC Workshop

APRIL 9–12 KNOXVILLE, TN
VI-HPS Tuning Workshop

APRIL 11–14 HUNTSVILLE, AL
IEEE SouthEastCon 2019

APRIL 14–18 KNOXVILLE, TN
JLESC

MAY 14–16 POZNAŃ, POLAND
BDEC2

MAY 20–24 RIO DE JANEIRO, BRAZIL
IPDPS

MAY 21–23 OAK RIDGE, TN
2019 OLCF User Meeting

MAY 28–31 CHICAGO, IL
MPI Forum

MAY 29–30 CHICAGO, IL
DoE/MEXT

JUNE 3–7 CAMBRIDGE, MA
MIT GPU Hackathon

JUNE 16–20 FRANKFURT, GERMANY
ISC High Performance 2019

JUNE 26–28 BORDEAUX, FRANCE
14th Scheduling for Large Scale Systems Workshop

JUNE 26–28 PHOENIX, AZ
ICS'19

JULY 7–12 KOBE, JAPAN
International HPC Summer School

JULY 10–11 EUGENE, OR
1st Workshop on NSF and DOE HPC Tools

JULY 22–25 COLLEGEVILLE, MN
Collegeville Workshop on Sustainable Scientific Software

JULY 22–26 KNOXVILLE, TN
OLCF/ECP OpenMP Hackathon

JULY 28–AUGUST 1 TAHOE CITY, CA
2019 Scalable Tools Workshop

JULY 28–AUGUST 9 ST. CHARLES, IL
Argonne Training Program on Extreme-Scale Computing

JULY 29 OAK RIDGE, TN
AI Expo

AUGUST 5–8 MONTEREY, CA
Monterey Data Conference

AUGUST 5–8 KYOTO, JAPAN
ICPP 2019

AUGUST 6–8 BLACKSBURG, VA
CEED Third Annual Meeting

AUGUST 26–30 GÖTTINGEN, GERMANY
Euro-Par 2019

SEPTEMBER 2–6 CAMBERRA, AUSTRALIA
Challenges in High Performance Computing Workshop

SEPTEMBER 2–3 DRESDEN, GERMANY
13th Parallel Tools Workshop

SEPTEMBER 4–6 SANTA FE, NM
2019 ECP CoPA All-Hands Meeting

SEPTEMBER 5–7 ZÜRICH, SWITZERLAND
MPI Forum

SEPTEMBER 10–13 ZÜRICH, SWITZERLAND
EuroMPI

SEPTEMBER 12–13 ALEXANDRIA, VA
SPEC HPC2020

SEPTEMBER 17–19 LEMONT, IL
Aurora Programming Workshop

SEPTEMBER 23–27 UPTON, NY
GPU Hackathon Brookhaven

SEPTEMBER 24–26 BOSTON, MA
IEEE HPEC 2019

SEPTEMBER 30–OCTOBER 4 EGMOND AAN ZEE, NETHERLANDS
ENUMATH 2019

OCTOBER 15–17 SAN DIEGO, CA
BDEC2

NOVEMBER 16–22 DENVER, CO
International Conference for High Performance Computing Networking, Storage, and Analysis 2019 (SC19)

The International Conference for High Performance Computing Networking, Storage, and Analysis (SC) is a staple of ICL's November itinerary. SC is vital to the growth and evolution of HPC in the United States, because it is the only US event that elicits substantial participation from all segments of the HPC community—including hundreds of users, developers, vendors, research institutions, and representatives of government funding agencies. Such a talent-rich gathering enables participants to discuss challenges, share innovations, and coordinate relationships and collaborations with some of the best minds in scientific and high-performance computing.

SC19 was held in Denver, Colorado on November 17–22. This year, five computational science research centers from the University of Tennessee—the Bredesen Center, the Global Computing Laboratory, the Innovative Computing Laboratory, the Joint Institute for Computational Sciences, and the SimCenter—represented the university by anchoring the University of Tennessee booth. As usual, ICL had a significant presence at SC, with faculty, research staff, and students giving talks, presenting papers, and leading "Birds of a Feather" sessions.

PARTNERSHIPS

ICL fosters relationships with many academic institutions and research centers and has proactively built enduring partnerships with HPC vendors and industry leaders in the United States and abroad. In this section, we recognize many of those partners and collaborators, most of whom we continue to work with today.

GOVERNMENT AND ACADEMIC

GCLab
Dr. Michela Taufer

Global Computing Laboratory

In June 2018, Prof. Michela Taufer joined UTK's Department of Electrical Engineering and Computer Science and relocated the Global Computing Laboratory to the Min H. Kao Electrical Engineering and Computer Science Building. The Global Computing Laboratory focuses on various aspects of HPC and scientific computing—including computational chemistry and chemical engineering, pharmaceutical sciences, seismology, and mathematics.

- Argonne National Laboratory
- Ames Laboratory
- Emory University
- Georgia Tech
- Indiana University
- Iowa State University
- JICS Joint Institute for Computational Sciences
- Berkeley Lab
- Lawrence Livermore National Laboratory
- Los Alamos National Laboratory
- Michigan Technological University
- NCSA
- NICS
- Northwestern University
- Oak Ridge National Laboratory
- Old Dominion University
- Pacific Northwest National Laboratory
- Rice
- Rutgers, The State University of New Jersey
- SDSC San Diego Supercomputer Center
- Sandia National Laboratories
- Stony Brook University
- Texas A&M University
- Berkeley, University of California
- UC Riverside
- The University of Chicago
- University of Colorado Denver
- University of Florida
- University of Houston
- Illinois, University of Illinois at Urbana-Champaign
- The University of Maine
- The University of New Mexico
- USC University of Southern California
- The University of Utah
- University of Washington
- Virginia Tech

INDUSTRY

- AMD
- Atos
- Cray, The Supercomputer Company
- Engility
- EPAnalytics
- IBM
- IERUS Technologies
- Intel
- MathWorks
- nag
- NVIDIA

INTERNATIONAL

Barcelona Supercomputing Center
BARCELONA, SPAIN

Central Institute for Applied Mathematics
JÜLICH, GERMANY

Doshisha University
KYOTO, JAPAN

École Normale Supérieure de Lyon
LYON, FRANCE

École Polytechnique Fédérale de Lausanne
LAUSANNE, SWITZERLAND

ETH Zürich
ZÜRICH, SWITZERLAND

European Centre for Research and Advanced Training in Scientific Computing
TOULOUSE, FRANCE

European Exascale Software Initiative
EUROPEAN UNION

Forschungszentrum Jülich
JÜLICH, GERMANY

High Performance Computing Center Stuttgart STUTTGART, GERMANY

INRIA FRANCE

Karlsruhe Institute of Technology
KARLSRUHE, GERMANY

Kasetsart University
BANGKOK, THAILAND

King Abdullah University of Science and Technology THUWAL, SAUDI ARABIA

Laboratoire d'Informatique de Paris 6 (LIP6) PARIS, FRANCE

Moscow State University
MOSCOW, RUSSIA

National Institute of Advanced Industrial Science and Technology TSUKUBA, JAPAN

Parallel and HPC Application Software Exchange TSUKUBA, JAPAN

Prometeus GmbH
MANNHEIM, GERMANY

Regionales Rechenzentrum Erlangen (RRZE) ERLANGEN, GERMANY

RIKEN
WAKŌ, JAPAN

Rutherford Appleton Laboratory
OXFORD, ENGLAND

Soongsil University
SEOUL, SOUTH KOREA

Technische Universität Wien
VIENNA, AUSTRIA

Technische Universität Dresden
DRESDEN, GERMANY

Tokyo Institute of Technology
TOKYO, JAPAN

Umeå University
UMEÅ, SWEDEN

Université Claude Bernard Lyon 1
LYON, FRANCE

University of Bordeaux
BORDEAUX, FRANCE

University of Cape Town
CAPE TOWN, SOUTH AFRICA

University of Manchester
MANCHESTER, ENGLAND

University of Paris-Sud
PARIS, FRANCE

University of Picardie Jules Verne
AMIENS, FRANCE

University of Tsukuba
TSUKUBA, JAPAN

LEADERSHIP

While leading-edge research and high-impact software are hallmarks of ICL's mission, the lab also cultivates a strong core of leadership. To this end, ICL actively engages the HPC and computational research communities through impactful efforts like those described below.

THE UNIVERSITY OF TENNESSEE KNOXVILLE

CENTER FOR INFORMATION TECHNOLOGY RESEARCH

FIND OUT MORE AT
http://citr.cs.utk.edu/

CENTER FOR INFORMATION TECHNOLOGY RESEARCH

The Center for Information Technology Research (CITR) was established in 2001 to drive the growth and development of leading-edge information technology research at UTK. CITR's primary objective is to develop a thriving, well-funded community in basic and applied information technology research to help the university capitalize on the rich supply of opportunities that now exist in this area. As part of this goal, CITR staff members currently provide primary administrative and technical support for ICL, helping maintain the lab's status as a world leader in high-performance and scientific computing. CITR has also provided secondary support for other UTK centers.

IGMCS

FIND OUT MORE AT
http://igmcs.utk.edu/

INTERDISCIPLINARY GRADUATE MINOR IN COMPUTATIONAL SCIENCE

Addressing the need for a new educational strategy in computational science, CITR worked with faculty and administrators from several departments and colleges in 2007 to help establish a new university-wide program that supports advanced degree concentrations in this critical new field across the curricula. Under the Interdisciplinary Graduate Minor in Computational Science (IGMCS), students pursuing advanced degrees in a variety of fields of science and engineering are able to extend their education with special courses of study that teach them both the fundamentals and the latest ideas and techniques from this new era of information-intensive research. The IGMCS curriculum, requirements, and policies are governed by a program committee composed of faculty members from participating IGMCS academic units and departments.

BDEC2

BIG DATA AND EXTREME-SCALE COMPUTING 2

FIND OUT MORE AT
http://www.exascale.org/

BIG DATA AND EXTREME-SCALE COMPUTING 2

In the past decade, the United States, the European Union, Japan, and China have each moved aggressively to develop their own plans for achieving exascale computing in the wake of a surging "data tsunami." Focusing on scientific research and building on the previous International Exascale Software Project (IESP) and Big Data and Extreme-scale Computing (BDEC) efforts, the goal of the follow-on BDEC2 project is to stage six international workshops to enable research communities in a wide range of disciplines to converge on a common platform in order to meet the daunting challenges of computing in the era of exascale and big data.

In 2019, ICL was instrumental in organizing and staging three BDEC2 workshops in Kobe, Japan; Poznań, Poland; and San Diego, CA. Along with Jack Dongarra, several members of ICL's CITR staff—including Terry Moore, Tracy Rafferty, Joan Snoderly, David Rogers, and Sam Crawford—played essential roles in making these BDEC2 workshops a major success. The next workshop, in which ICL is also heavily involved, is planned for March 2020 in Porto, Portugal.

JLESC

FIND OUT MORE AT
https://jlesc.github.io/

JOINT LABORATORY FOR EXTREME SCALE COMPUTING

ICL is now an Associate Partner of the Joint Laboratory for Extreme Scale Computing (JLESC). JLESC, founded in 2009 by the French Institute for Research in Computer Science and Automation (INRIA) and the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign, is an international, virtual organization that aims to enhance the ability of member organizations and investigators to overcome software challenges found in extreme-scale, high-performance computers.

JLESC engages computer scientists, engineers, application scientists, and industry leaders to ensure that the research facilitated by the joint laboratory addresses science and engineering's most critical needs and takes advantage of the continuing evolution of computing technologies. Other partners include Argonne National Laboratory, the Barcelona Supercomputing Center, the Jülich Supercomputing Center, and the RIKEN Center for Computational Science.

PEOPLE

Ahmad Abdelfattah
RESEARCH SCIENTIST II

Mohammed Al Farhan
POST DOCTORAL RESEARCH ASSOCIATE

STAFF AND STUDENTS

As the HPC landscape continues to evolve rapidly, remaining at the forefront of discovery requires great vision and skill. To address this evolution and to remain a leader in innovation, we have assembled a staff of top researchers from all around the world who apply a variety of novel and unique approaches to the challenges and problems inherent in world-class scientific computing.

As part of an engineering college at a public research university, we have a responsibility to combine exemplary teaching with cutting-edge research. As such, we regularly employ bright and motivated graduate and undergraduate students. We have been, and will continue to be, very proactive in securing internships and assistantships for highly motivated and hardworking students.

Hartwig Anzt
CONSULTANT

Alan Ayala
POST DOCTORAL RESEARCH ASSOCIATE

Daniel Barry
GRADUATE RESEARCH ASSISTANT

John Batson
GRADUATE RESEARCH ASSISTANT

In 2019, two of ICL's Graduate Research Assistants earned their PhDs from UTK under the guidance and mentorship of Prof. Dongarra and ICL's research scientists. Congratulations to:

Reazul Hoque
PhD in Computer Science, December 2019

Thananon Patinyasakdikul
PhD in Computer Science, December 2019

Natalie Beams
RESEARCH SCIENTIST I

George Bosilca
RESEARCH ASSISTANT PROFESSOR

Aurelien Bouteiller
RESEARCH DIRECTOR

Cade Brown
UNDERGRADUATE RESEARCH ASSISTANT

ICL/GCLab Joint Retreat
FALL 2019

Qinglei Cao
GRADUATE RESEARCH ASSISTANT

Earl Carr
PROGRAM ADMINISTRATOR

Tony Castaldo
RESEARCH SCIENTIST II

Sebastien Cayrols
POST DOCTORAL RESEARCH ASSOCIATE

Ali Charara
RESEARCH SCIENTIST II

Sam Crawford
TECHNICAL SERVICES LEADER

Anthony Danalis
RESEARCH DIRECTOR

Jack Dongarra
UNIVERSITY DISTINGUISHED PROFESSOR

David Eberius
GRADUATE RESEARCH ASSISTANT

Teresa Finchum
ADMINISTRATIVE SPECIALIST II

Jamie Finney
SOFTWARE ENGINEER

Mark Gates
RESEARCH DIRECTOR

Damien Genet
RESEARCH SCIENTIST I

Hidehiko Hasegawa
VISITING SCHOLAR

Thomas Herault
RESEARCH DIRECTOR

Heike Jagode
RESEARCH ASSISTANT PROFESSOR

Julie Langou
RESEARCH LEADER

STAFF AND STUDENTS

CONTINUED

Tracy Lee
FINANCIAL SPECIALIST I

Jiali Li
GRADUATE RESEARCH ASSISTANT

Yicheng Li
GRADUATE RESEARCH ASSISTANT

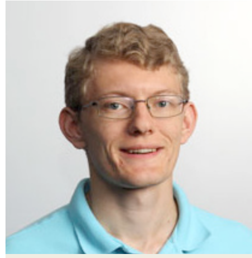

Neil Lindquist
GRADUATE RESEARCH ASSISTANT

Florent Lopez
POST DOCTORAL RESEARCH ASSOCIATE

Nuria Losada
POST DOCTORAL RESEARCH ASSOCIATE

Xi Luo
GRADUATE RESEARCH ASSISTANT

Piotr Luszczek
RESEARCH ASSISTANT PROFESSOR

Terry Moore
ASSOCIATE DIRECTOR

Phil Mucci
CONSULTANT

Yu Pei
GRADUATE RESEARCH ASSISTANT

Gerald Ragghianti
RESEARCH LEADER

Yves Robert
VISITING SCHOLAR

David Rogers
IT SPECIALIST III

Daniel Schultz
GRADUATE RESEARCH ASSISTANT

Leighanne Sisk
ADMINISTRATIVE SPECIALIST I

Joan Snoderly
ASSISTANT DIRECTOR

Dalal Sukkari
POST DOCTORAL RESEARCH ASSOCIATE

Stanimire Tomov
RESEARCH ASSISTANT PROFESSOR

Yaohung Tsai
GRADUATE RESEARCH ASSISTANT

Aaron Welch
NON-UT STUDENT ASSISTANT

Frank Winkler
CONSULTANT

Asim YarKhan
RESEARCH SCIENTIST II

Dong Zhong
GRADUATE RESEARCH ASSISTANT

VISITORS

ICL has a long-standing tradition of hosting visitors from all over the world. Some stay only briefly to give insightful seminars or presentations, while others remain with us for as long as a year to collaborate, teach, and learn. Our connection to these researchers enables us to leverage an immense array of intellectual resources and work with the best and brightest people in the HPC community.

Srinivas Aluru
GEORGIA TECH

Marc Baboulin
UNIVERSITY OF PARIS-SUD

Vinton Cerf
GOOGLE RESEARCH

Stephen Herbein
LAWRENCE LIVERMORE NATIONAL LABORATORY

Oscar Hernandez
OAK RIDGE NATIONAL LABORATORY

Torsten Hoefler
ETH ZÜRICH

Axel Huebl
LAWRENCE BERKELEY NATIONAL LABORATORY

Vipin Kumar
UNIVERSITY OF MINNESOTA

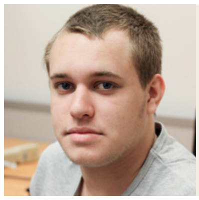

Valentin Le Fevre
ENS-LYON

Sticks Mabakane
COUNCIL FOR SCIENTIFIC AND INDUSTRIAL RESEARCH

Daniel Nichols
JOINT INSTITUTE FOR COMPUTATIONAL SCIENCES

Danny Rorabaugh
GLOBAL COMPUTING LABORATORY

Joseph Schuchart
HIGH PERFORMANCE COMPUTING CENTER, STUTTGART (HLRS)

Hejer Shaiek
ENSEEIHT

Victoria Stodden
UNIVERSITY OF ILLINOIS AT URBANA–CHAMPAIGN

Stephen Thomas
GLOBAL COMPUTING LABORATORY

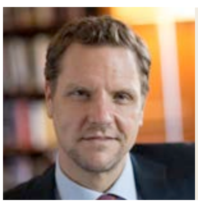

Michael Witmore
FOLGER SHAKESPEARE LIBRARY

ALUMNI

ICL has attracted many research scientists and students from a variety of backgrounds and academic disciplines. Many of these experts came to UTK specifically to work with Prof. Dongarra—beginning a long list of research talent to pass through ICL and move on to make exciting contributions at other institutions and organizations.

Maksims Abalenkovs
Carolyn Aebischer
Sudesh Agrawal
Bivek Agrawal
Emmanuel Agullo
Jennifer Allgeyer
Wes Alvaro
Ed Anderson
Daniel Andrzejewski
Thara Angskun
Papa Arkhurst
Dorian Arnold
Rizwan Ashraf
Cedric Augonnet
Marc Baboulin
Matthew Bachstein
Zhaojun Bai
Ashwin Balakrishnan
Richard Barrett
Alex Bassi
David Battle
Micah Beck
Dulceneia Becker
Daniel Becker
Adam Beguelin
Annamaria Benzoni
Tom Berry
Vincent Berthoux
Scott Betts
Nikhil Bhatia
Noel Black
Laura Black
Susan Blackford
Wesley Bland
Kartheek Bodanki
David Bolt
Fernando Bond
Carolyn Bowers
Vincent Bridonneau
Barry Britt
Randy Brown
Cynthia Browne
Murray Browne
Bonnie Browne
Giuseppe Bruno
Antonin Bukovsky
Greg Bunch
Alfredo Buttari
Anthony Canino
Domingo Gimenez Canovas
Chongxiao “Shawn” Cao
Henri Casanova
Cedric Castagnede
Ramkrishna Chakrabarty
Sharon Chambers
Zizhong Chen
Jaeyoung Choi
Wahid Chrabakh
Eric Clarkson
Andy Cleary
Michelle Clinard
Vincent Cohen-Addad
Matthias Colin
Charles Collins
Rephael Congmon
Stephanie Moreaud Cooper
Tom Cortese
Camille Coti
Jason Cox
David Cronk
Javier Cuenca
Manoel Cunha
Yuanshun (Shaun) Dai
Cricket Deane
Remi Delmas
Frederic Desprez
Ying Ding
Jun Ding
Jin Ding
Martin Do
Simplice Donfack
Leon Dong
Tingxing “Tim” Dong
Nick Dongarra
David Doolin
Joe Dorris
Andrew Downey
Mary Drake
Julio Driggs
Brian Drum
Peng Du
Eduardo Echavarria
Victor Eijkhout
Brett Ellis
Shawn Ericson
Zachary Eyler-Walker
Lisa Ezzell
Christoph Fabianek
Graham Fagg
Mathieu Faverge
Diana Fayad
Shengzhog Feng
Don Fike
Salvatore Filippone
Anna Finchum
Mike Finger
Markus Fischer
Len Freeman
Xiaoquan Fu
Erika Fuentes
Karl Fuerlinger
Megan Fuller
Edgar Gabriel
Lynn Gangwer
Tracy Gangwer
Nathan Garner
Kelley Garner
Tina Garrison
Adriana Garties
Peter Gaultney
Christoph Geile
Jean Patrick Gelas
Boris Gelfend
Jonathan Gettler
Scott Gibson
Eric Greaser
Stan Green
Alice Gregory
Amina Guermouche
Jason Gurley
Bilel Hadri
Hunter Hagewood
Azzam Haidar
Christian Halloy
Sven Hammarling
J. Mike Hammond
Hanumantharayappa
Satomi Hasegawa
Chris Hastings
Blake Haugen
David Henderson
Greg Henry
John Henry
Julien Herrmann
Holly Hicks
Alexandra Hicks-Hardiman
Sid Hill
Tomoyuki Hiroyasu
George Ho
Josh Hoffman
Reazul Hoque
Jeff Horner
Mitch Horton
Yan Huang
Harry Hughes
Aurelie Hurault
Chris Hurt
Paul Jacobs
Emmanuel Jeannot
Weizhong Ji
Yulu Jia
Weicheng Jiang
Song Jin
Patrick Johansson
Aral Johnson
Matt Johnson
Sean Jolly
Kim Jones
Jan Jones
Vijay Joshi
Khairul Kabir
Venkata Kakani
Ajay Kalhan
Balajee Kannan
Madhuri Kasam
Kiran Kumar Kasichayanula
Ajay Katta
David Katz
Joshua Kelly
Supriya Kilambi
Youngbae Kim
Myung Ho Kim
Jenya Kirshtein
Cindy Knisley

Michael Kolatis
Anara Kozhokanova
Moritz Kreutzer
Chandra Krintz
Tilman Kuestner
Jakub Kurzak
Krerkchai Kusolchu
Coire Kyle
Amanda Laake
Xavier Lacoste
Julien Langou
Jeff Larkin
Brian LaRose
Frank Lauer
Valentin Le Fevre
DongWoo Lee
Pierre Lemarinier
Todd Letsche
Sharon Lewis
Xiang Li
Weiran Li
Yinan Li
Chaoyang Liu
Kevin London
Matt Longley
Florent Lopez
Hatem Ltaief
Yuechao Lu
Daniel Lucio
Richard Luczak
Teng Ma
Sticks Mabakane
Robert Manchek
Gabriel Marin
Tushti Marwah
Theo Mary
Ian Masliah
Donald McCasland
Tyler McDaniel
Paul McMahan
Eric Meek
James Meyering
Jeremy Millar
Michelle Miller
Cindy Mitchell
Stuart Monty
Erik Moore
Keith Moore
Shirley Moore
Robert Morgan
Kishan Motheramgari
Steven Moulton
Daichi Mukunoki
Matthew Nabity
Shankar Narasimhaswami
Rajib Nath
Fernando Navarro
John Nelson
Donnie Newell
Peter Newton
Jonas Nilsson
Jakob Oestergaard
Caroline Papadopoulos
Leelinda Parker
Thibault Parpaite
Thananon Patinyasakdikul
Dilip Patlolla
Andy Pearson
Paul Peltz
Theresa Pepin
Antoine Petitet
Peter Pham
Gregoire Pichon
Vlado Pjesivac
Jelena Pjesivac-Grbovic
James S. Plank
Tim Poore
Roldan Pozo
Farzona Pulatova
Martin Quinson
Tammy Race
Tracy Rafferty
Sangamesh Ragate
James Ralph
Ganapathy Raman
Kamesh Ramani
Mei Ran
Arun Rattan
Sheri Reagan
Mike Reynolds
George Rhinehart
Jon Richardson
Stephen Richmond
Ken Roche
Andrew Rogers
Tom Rothrock
Tom Rowan
Narapat (Ohm) Saengpatsa
Kiran Sagi
Evelyn Sams
Ken Schwartz
Keith Seymour
Farial Shahnaz
Hejer Shaiek
Brian Sheely
Zhiao Shi
Sergei Shinkarev
Majed Sidani
Shilpa Singhal
Matt Skinner
Peter Soendergaard
Raffaele Solca
Gwang Son
Fengguang Song
Thomas Spencer
Jeffrey Steill
Erich Strohmaier
Xiaobai Sun
Martin Swany
Daisuke Takahashi
Judi Talley
Ronald Tam
Yuan Tang
Chunyan Tang
Yusuke Tanimura
Keita Teranishi
Dan Terpstra
Joe Thomas
John Thurman
Francoise Tisseur
Jude Toth
Bernard Tourancheau
Volodymyr Turchenko
Lauren Vaca
Phil Vaccaro
Sathish Vadhiyar
Robert van de Geijn
Chad Vawter
Eugene Vecharynski
Scott Venckus
Antoine Vernois
Reed Wade
Michael Walters
Robert Waltz
Mike Waltz
Xiaoyang Wang
Bob Ward
Jerzy Wasniewski
Vince Weaver
Scott Wells
David West
R. Clint Whaley
Jody Whisnant
James White
Scotti Whitmire
Susan Wo
Felix Wolf
Stephen Wood
Panruo Wu
Wei Wu
Jiayi Wu
Qiu Xia
Tinghua Xu
Ichitaro Yamazaki
Tao Yang
Erlin Yao
Kevin Ye
Jin Yi
Haihang You
Lamia Youseff
Brian Zachary
Dmitry Zaitsev
Omar Zenati
Yuanlei Zhang
Junlong Zhao
Yong Zheng
Min Zhou
Luke Zhou

ACKNOWLEDGMENTS

For over 30 years, our knowledge and hard work have earned the trust and support of many agencies and organizations that have funded and continue to fund our efforts. Without them, we simply would not be able to conduct leading-edge research.

The main source of support has been federal agencies that are charged with allocating public research funding, and we acknowledge the following agencies for supporting our efforts:

AFOSR AIR FORCE OFFICE OF SCIENTIFIC RESEARCH

U.S. DEPARTMENT OF ENERGY | Office of Science

U.S. DEPARTMENT OF ENERGY | NNSA National Nuclear Security Administration

NSF National Science Foundation

OAK RIDGE National Laboratory

In addition to the support of the federal government, we have solicited strong support from private industry, which has also played a significant role in our success and growth. We gratefully acknowledge the following vendors for their generosity and support:

AMD

Atos

ENGILITY

intel

MathWorks™

NVIDIA

UT-BATTELLE

Administrative and technical support for ICL is provided by the Center for Information Technology Research (CITR). CITR is supported in part by the University of Tennessee Office of Research and Engagement.

THE UNIVERSITY OF TENNESSEE KNOXVILLE

CENTER FOR INFORMATION TECHNOLOGY RESEARCH

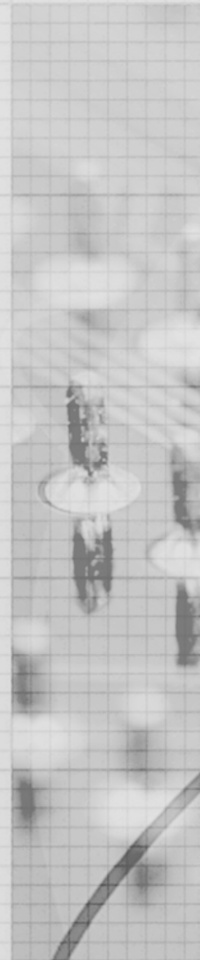

ICL INNOVATIVE COMPUTING LABORATORY

FIND OUT MORE AT **http://www.icl.utk.edu/**