

INNOVATIVE COMPUTING LABORATORY *2025/26 REPORT*

© 2026 Innovative Computing Laboratory. All rights reserved. Any trademarks or registered trademarks that appear in this document are the property of their respective owners.

EEO/TITLE IX/AA/SECTION 504 STATEMENT

All qualified applicants will receive equal consideration for employment and admission without regard to race, color, national origin, religion, sex, pregnancy, marital status, sexual orientation, gender identity, age, physical or mental disability, genetic information, veteran status, and parental status, or any other characteristic protected by federal or state law.

In accordance with the requirements of Title VI of the Civil Rights Act of 1964, Title IX of the Education Amendments of 1972, Section 504 of the Rehabilitation Act of 1973, and the Americans with Disabilities Act of 1990, the University of Tennessee affirmatively states that it does not discriminate on the basis of race, sex, or disability in its education programs and activities, and this policy extends to employment by the university.

Inquiries and charges of violation of Title VI (race, color, and national origin), Title IX (sex), Section 504 (disability), the ADA (disability), the Age Discrimination in Employment Act (age), sexual orientation, or veteran status should be directed to the Office of Equity and Diversity, 1840 Melrose Avenue, Knoxville, TN 37996-3560, telephone 865-974-2498. Requests for accommodation of a disability should be directed to the ADA Coordinator at the Office of Equity and Diversity.

The cover features a large orange textured rectangle on the left side, with several horizontal stripes of varying shades of orange extending from its bottom edge across the entire width of the page. The right side of the cover is a solid, lighter orange color. The title 'ICL 2025/26 REPORT' is centered in the lower half of the right side.

ICL 2025/26 REPORT

- 2 FROM THE DIRECTOR
- 4 2025 HIGHLIGHTS
- 12 RESEARCH PROFILES
- 22 PEOPLE
- 26 PARTNERSHIPS
- 27 ACKNOWLEDGMENTS

FROM THE DIRECTOR



In 2025, the Innovative Computing Laboratory (ICL) celebrated a year of scientific achievement, community milestones, and continued global leadership in high-performance computing (HPC). The laboratory marked its 35th anniversary with a dedicated workshop bringing together 38 speakers from across generations of ICL alumni and current members, reflecting the lasting impact of the ICL community. We also honored the retirement of longtime business manager Tracy Lee, whose dedication shaped the laboratory for many years, and welcomed Waldhys Rodoli as our new business manager, strengthening ICL's operational foundation for the future.

ICL continued its strong commitment to student success and workforce development. Daniel Barry completed his PhD advancing performance-counter analysis and measurement methodologies, while Max Melnichenko graduated with contributions to randomized numerical linear algebra and scalable algorithms. Dong Jun Wuon worked on his MSc thesis on standardized AMD GPU performance monitoring, strengthening ICL's research in portable performance analysis. Research Assistant Professors Aurélien Bouteiller, Anthony Danalis, and Mark Gates moved to positions at AMD, demonstrating ICL's strong impact in developing talent sought after by the HPC industry.

ICL students achieved notable recognition and industry engagement, including Tatiana Melnichenko's participation in the undergrad internship program SULI at ORNL advancing performance portability in Julia, and Dong Jun Wuon's work at Amazon on security and performance testing frameworks for Project Kuiper. Tatiana Melnichenko was also selected as an SC25 Best Poster Finalist in the ACM Student Research Competition, and ICL alum Rabab Alomairy received the SC25 Early Career Researchers Award for Excellence in HPC. In addition, Heike Jagode was named the inaugural recipient of the 2025 Tickle College of Engineering Research Award for Research Faculty.

ICL's research influence was highlighted internationally through the work of Jack Dongarra, who elevated national conversations on the future of computing through a widely recognized Software Engineering Daily podcast appearance and a co-authored *Science* article calling for a long-term U.S. strategy in HPC. Dongarra was further honored with the prestigious NEC C&C Prize. Also, ICL alumni Hatem Ltaief and Piotr Luszczek were appointed Editors-in-Chief of *ACM Transactions on Mathematical Software (TOMS)*, underscoring ICL's role in developing leaders who guide the global computational mathematics community.

ICL further advanced the global HPC software ecosystem through multiple PAPI 7.2.x releases in 2025, introducing support for AMD MI300 APUs and GPUs, emerging CPU architectures, heterogeneous P-core/E-core awareness, portable GPU performance presets, and Top-Down performance metrics—significantly simplifying cross-vendor, architecture-agnostic performance analysis for modern heterogeneous supercomputing platforms.

Throughout the year, ICL strengthened its role as a hub for collaboration and scientific exchange by hosting workshops on GPU numerical computing and Sparse BLAS standards and contributing to major international venues, including ISC 2025, SC25, the SIAM Annual Meeting, the NSF CSSI Meeting, the Heidelberg Laureate Forum, ILAS 2025, and SMC25.

Together, these accomplishments reflect ICL's continued commitment to advancing scientific innovation, developing future leaders, and building sustainable software and research communities that drive next-generation computing forward.

Heike Jagode
 Heike JAGODE
 INTERIM ICL DIRECTOR




ENABLING TECHNOLOGIES FOR HIGH-PERFORMANCE COMPUTING

SINCE 1989

The Innovative Computing Laboratory is a research center in the Tickle College of Engineering at the University of Tennessee. ICL has been a world leader in enabling technologies and software for scientific computing for more than 35 years. With deep expertise in numerical linear algebra, distributed computing, performance optimization, and benchmarking, ICL develops scalable software and methods for today's—and tomorrow's—computing architectures. These tools help accelerate research across areas like life sciences, climate modeling, earthquake prediction, energy exploration, and advanced materials. By combining foundational research with close collaboration among universities, national laboratories, and industry, ICL ensures its innovations translate into practical, high-impact solutions on the world's most powerful supercomputers.



2025 HIGHLIGHTS



ICL 35th Anniversary Workshop

*CELEBRATING 35 YEARS OF
INNOVATIVE COMPUTING*

In May 2025, ICL hosted a workshop to recognize its 35th anniversary. Held in the UT Student Union on the University of Tennessee campus, the event brought together ICL alumni and current members for two days of technical talks and conversation. The agenda included 38 speakers, two panel sessions, and a breakout discussion, with social events around Knoxville providing time to reconnect, network, and celebrate the milestone.

While the talks offered technical insight, the most lasting impression came from the people and connections that define ICL. Alumni shared updates on their current work and personal reflections on how their time at ICL helped shape their careers. Presenters described the "ICL experience" of being challenged early, learning by doing, collaborating across disciplines and cultures, and building software that lasts. They also reflected on mentorship, teamwork, and everyday problem-solving that shaped their careers and continue to influence how they lead projects today. In that sense, the workshop was both a forward-looking program and a homecoming, celebrating a shared history and the extended community that continues to carry ICL's spirit forward.


State-of-the-Art GPU Numerical Computing Workshop

IN HONOR OF STAN TOMOV



In April 2025, ICL hosted a two-day workshop honoring the life, legacy, and scientific contributions of Dr. Stanimire Tomov. Stan was an extraordinary researcher, a respected leader, and a cherished member of ICL for over two decades. He led the MAGMA project and served as a research director at ICL before joining NVIDIA in 2024. His technical brilliance, collaborative spirit, and deep kindness helped shape the culture and success of ICL.

In tribute to his contributions and spirit, State-of-the-Art GPU Numerical Computing brought together colleagues, collaborators, and former students for two days of technical talks and shared memories. The workshop featured presentations on cutting-edge advances in numerical computing, along with reflections on Stan’s influence, mentorship, and friendship.



Spring 2025 Sparse BLAS Workshop

TOWARD A NEW SPARSE BLAS STANDARD

The Sparse BLAS Workshop 2025, the third in a series of community workshops led by Hartwig Anzt, was held February 26–28 at the University of Tennessee, Knoxville, bringing together researchers and developers from academia, national labs, and industry. Hosted by ICL, the meeting advanced a shared effort to define a standardized, hardware-portable interface for sparse linear algebra, supporting better performance and portability across diverse computing architectures. Breakout discussions and a hackathon-style session helped explore how existing libraries could be mapped to the proposed API.

2025 HIGHLIGHTS

AWARDS AND HONORS

Jack Dongarra Honored by NEC *A GLOBAL AWARD FOR ENDURING IMPACT IN HPC*

ICL founder Jack Dongarra was named a recipient of the 2025 NEC Computer and Communications (C&C) Prize, an award that recognizes outstanding contributions to research and development in the integration of computing and communications technologies and their impact on society. The 2025 C&C Prize honored two distinct groups: one recognized the team responsible for the invention, practical implementation, and global spread of the QR Code, and the other celebrated Jack’s decades of pioneering work in high-performance scientific computing, including foundational numerical libraries and performance benchmarking infrastructure that have shaped the field. The formal award ceremony took place on November 26 in Tokyo.



Heike Jagode Named Recipient of 2025 TCE Research Award for Research Faculty

ICL is proud to report that Heike Jagode was selected by Tickle College of Engineering as the inaugural recipient of the 2025 TCE Research Award for Research Faculty. Heike received the award at the TCE Awards Dinner on May 1 at Bridgewater Place.



Tatiana Melnichenko Earns SC25 Honors


ICL student Tatiana Melnichenko earned recognition at SC25 for her work with an ORNL team studying the Mojo programming language’s GPU capabilities for HPC science workloads. Their paper, “Mojo: MLIR-based Performance-Portable HPC Science Kernels on GPUs for the Python Ecosystem,” received a Best Paper Award at WACCPD 2025, and Tatiana’s companion poster was named a Best Poster Finalist in the ACM Student Research Competition.

2025 PHD GRADUATES



**Daniel
BARRY**

Dr. Daniel Barry completed his PhD under the supervision of Heike Jagode, in collaboration with Anthony Danalis. His dissertation, "Automated Classification and Verification of Performance Counters," developed methods to systematically organize and validate hardware performance counter data, supporting more effective performance analysis and optimization of scientific applications on HPC systems. Following graduation, Daniel has remained at ICL as a Research Scientist, continuing his contributions to the lab's projects and collaborations.



**Max
MELNICHENKO**

Dr. Max Melnichenko completed his PhD under the guidance of Jack Dongarra with a dissertation titled "Responsible Recklessness: Why High Performance Computing Must Embrace Randomized Numerical Linear Algebra." Max also credits the mentorship of Piotr Luszczek, who served as an informal advisor throughout his graduate work. After graduating, Max accepted a postdoctoral appointment with the International Computer Science Institute (ICSI) and the UC Berkeley Department of Statistics, beginning in January 2026, where he will work under the supervision of Michael Mahoney.

2025 HIGHLIGHTS

2025 ICL SEMINAR SERIES


Friday lunch talks are a long-standing tradition at ICL that offer opportunities to foster collaborations and bring new ideas to the group. ICL's 2025 Seminar Series featured a distinguished group of invited speakers from UTK, ORNL, and other institutions around the globe.



JAN 24
Stephen MARZ
EECS



FEB 7
Tabitha SAMUEL
NICS



FEB 14
Yang ZHANG
EECS



FEB 21
Chiang-Heng CHIEN
Brown University



FEB 21
Benjamin KIMIA
Brown University




MAR 14
Pedro VALERO LARA
ORNL




APR 4
Max MELNICHENKO
ICL




APR 11
Gerald RAGGHIANI
ICL



MAY 2
Daniel BARRY
ICL



MAY 9
Rabab ALOMAIRY
MIT



JUL 18
Zhuowei GU
Saint Louis University



SEP 12
Sameh ABDULAH
KAUST




SEP 19
Terry MOORE
Retired ICL Alum



OCT 3
Martin SWANY
Indiana University



OCT 17
Jake KING
Flat Lux



OCT 24
Antonios KOUKAS
Illinois Tech



NOV 7
Jeremy THOMPSON
CU Boulder




NOV 14
Hatem LTAIEF
KAUST



DEC 5
Piotr LUSZCZEK
ICL / MIT



DEC 12
Azzam HAIDAR
NVIDIA



DEC 19
Treece BURGESS
ICL

RECENT PUBLICATIONS



Dongarra, J.
Durable Engines of Discovery
Communications of the ACM, September 2025



READ THE ARTICLE AT
tiny.utk.edu/acm-discovery

In this *Communications of the ACM* article, ICL founder Jack Dongarra highlighted how long-term U.S. federal research support has helped fuel major advances in computing and scientific discovery—and why sustaining that investment remains critical for future innovation.

Chien, C-H., A. Abdelfattah, and B. Kimia
Accelerating Homotopy Continuation with GPUs: Application to Trifocal Pose Estimation

2025 *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Milano, Italy, IEEE, July 2025

Deelman, E., J. Dongarra, B. Hendrickson, A. Randles, D. Reed, E. Seidel, and K. Yelick
High-performance computing at a crossroads

Science, vol. 387, issue 6736, pp. 829 - 831, February 2025

Dongarra, J., J. Gunnels, H. Bayraktar, A. Haidar, and D. Ernst

Accelerating Supercomputing: AI-Hardware-Driven Innovation for Speed and Efficiency

2025 *IEEE High Performance Extreme Computing Conference (HPEC)*, Wakefield, MA, USA, IEEE, October 2025

Abdelfattah, A., J. Dongarra, M. Fasi, M. Mikaitis, and F. Tisseur

Analysis of Floating-Point Matrix Multiplication Computed via Integer Arithmetic

arXiv, June 2025

Barry, D.
Automated Classification and Verification of Performance Counters
 December 2025

Petit, Q., C. Li, N. Emad, and J. Dongarra
Efficient Embedding Initialization via Dominant Eigenvector Projections
SC Workshops '25: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis, St Louis, MO USA, ACM, November 2025

Luo, L., K. Bochenina, T. M. Abuhay, N. Dorzhu, G. Kampis, S. Kovalchuk, V. Krzhizhanovskaya, M. Paszyński, C. de Mulatier, J. Dongarra, et al.

Evolution of the computational science community: The dynamics of topics and collaborations in 24 years of ICCS and JoCS publications

Journal of Computational Science, vol. 89, July 2025

Dongarra, J., and P. Luszczek

HPL-MxP benchmark: Mixed-precision algorithms, iterative refinement, and scalable data generation

The International Journal of High Performance Computing Applications, September 2025

Schuchart, J., A. Bouteiller, T. Herault, E. Valeev, G. Bosilca, and R. J. Harrison

Constraints and Mutexflows for Scalable Block-Sparse Matrix Multiplication Using Template Task Graphs

SN Computer Science, vol. 7, issue 2, January 2026

Tahmid, T., and H. Jagode
PAPI Support for Specialized AI Architectures

SC25: 10th International Parallel Data Systems Workshop (PDSW 2025), St. Louis, MO, IEEE, November 2025

Tahmid, T., M. Gates, P. Luszczek, and C. D. Schuman

SpikeRL: A Scalable and Energy-efficient Framework for Deep Spiking Reinforcement Learning

2025 *International Conference on Neuromorphic Systems (ICONS)*, Seattle, WA, USA, IEEE, January 2026

Tsai, Y-H., M. Bode, and H. Anzt
What Will the Grace Hopper-Powered Jupiter Supercomputer Bring for Sparse Linear Algebra

SCA/HPCAsia 2026: Supercomputing Asia and International Conference on High Performance Computing in Asia Pacific Region, Osaka, Japan, ACM, January 2026

2025 HIGHLIGHTS

SC25 HPC IGNITES

SC25 brought the global HPC community together in St. Louis, Missouri (November 16–21), drawing more than 16,500 attendees and a record 559 exhibitors. Under this year’s theme, “HPC Ignites,” the conference highlighted how high-performance computing sparks breakthroughs across research, industry, and society. A contingent of research staff and students represented ICL at the event and hosted the University of Tennessee booth in partnership with NICS and AI-Tenn. Throughout the week, ICL staff, students, and alumni contributed in technical sessions, alumni networking, and community dialogues.



Mixed-Precision HPC Tutorial at ISC 2025

On June 13, ICL researchers Jack Dongarra, Piotr Luszczek, and Hartwig Anzt hosted a half-day tutorial, “Modern Mixed-Precision Methods: Hardware Perspectives, Algorithms, Kernels, and Solvers,” that exposed attendees to the rapidly expanding landscape of mixed- and multi-precision methods in modern scientific computing. Framed around the growing cross-pollination between HPC and AI, the session highlighted recent algorithmic progress in exploiting multiple precisions for increased efficiency in performance, communication bandwidth, and storage overhead on modern hardware.



Responsibly Reckless Algorithms

Speaking to an international audience that included 33 Laureates and more than 600 young scientists at the 74th Lindau Nobel Laureate Meeting in Lindau, Germany, ICL founder Jack Dongarra delivered his featured talk titled “An Overview of High Performance Computing and Responsibly Reckless Algorithms.”



VIEW JACK'S TALK AT
tiny.utk.edu/lindau-dongarra

CONFERENCES ATTENDED



JAN 8 VIRTUAL
MPI Forum



FEB 11-13 EUGENE, OR
AMD Tools Workshop 2025



FEB 11-13 JUELICH, GERMANY
ExaNLA Workshop



FEB 19-21 ST. LOUIS, MO
WAMTA 2025



MAR 10-15 SINGAPORE
ACM ASEAN HPC School



MAR 10-13 SINGAPORE
SupercomputingAsia 2025



JUN 10-13 HAMBURG, GERMANY
ISC 2025




JUN 23-27 KAOHSIUNG, TAIWAN
ILAS2025
The 26th Conference of the
International Linear Algebra Society



JUN 29 - JUL 5 KAOHSIUNG, TAIWAN
Lindau Nobel Laureate Meeting



JUL 7-9 SINGAPORE
ICCS 2025
International Conference on
Computational Science



JUL 28 - AUG 1 MONTRÉAL, QUÉBEC
SIAM Annual Meeting



JUL 29-31 VIRTUAL
ACM ICONS 2025
International Conference on
Neuromorphic Systems



JUL 29-31 DENVER, CO
2025 NSF CSSI Meeting



AUG 4-5 ST. CHARLES, IL
ATPESC 2025
Argonne Training Program on
Extreme-Scale Computing



AUG 25-28 CHATTANOOGA, TN
SMC25
Smoky Mountains Computational
Sciences and Engineering Conference



SEP 13-19 HEIDELBERG, GERMANY
Heidelberg Laureate Forum



OCT 25-27 SHANGHAI, CHINA
World Laureates Forum



OCT 26-31 ABU DHABI, UAE
ADIA Labs Symposium




NOV 16-21 ST. LOUIS, MO
SC25



DEC 3-5 BEN GUERIR, MOROCCO
SupercomputingAfrica 2025

RESEARCH PROFILES



NUMERICAL LINEAR ALGEBRA

Numerical linear algebra lies at the heart of scientific computing, forming the mathematical foundation for simulations, optimizations, and machine learning applications across a wide range of disciplines. The high-performance libraries that implement these algorithms must evolve in tandem with rapidly changing hardware architectures to fully harness modern computing resources. As processors become more heterogeneous—integrating GPUs, FPGAs, and specialized accelerators—traditional approaches to dense and sparse linear algebra must be re-engineered to maximize efficiency, scalability, and portability.

BALLISTIC

Basic ALgebra Libraries for Sustainable Technology with Interdisciplinary Collaboration (BALLISTIC) is an NSF-funded effort to create new software components capable of running at every level of the hardware hierarchy by delivering seamless access to the most up-to-date algorithms, numerics, and performance via familiar Linear Algebra PACKage (LAPACK) and Scalable Linear Algebra PACKage (ScaLAPACK) interfaces. BALLISTIC makes advanced algorithms, numerics, and performance capabilities available through new interface extensions and by providing a well-engineered conduit for channeling new developments to science and engineering applications that depend on high-performance, high-quality linear algebra libraries.

Scientific software libraries have long provided a large and growing resource for high-quality, reusable software components upon which applications from science and engineering can be rapidly constructed. The BALLISTIC project will introduce tools to simplify the transition to the next generation of extreme-scale computer architecture.

FIND OUT MORE AT
icl.utk.edu/ballistic

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 simultaneously produce large numbers of small matrices.

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
icl.utk.edu/bblas

Batched Sparse LA

Batched sparse linear algebra solvers form the new frontier for algorithmic development and performance engineering. Many applications require simultaneous solutions of small linear systems of structurally sparse equations. In order to move towards higher hardware utilization, it is essential to provide these applications with appropriate interfaces to efficient implementations of batched sparse solvers running on modern hardware accelerators.

The Batched Sparse Linear Algebra effort developed batched kernels for small tensor operations for unassembled matrix-free iterative solvers, batched solvers for partially assembled operators, and batched solvers with support for various sparse formats. Participants included representatives from ECP's numerical libraries, scientific applications, and hardware vendors. The project milestones were sets of interfaces to batched sparse linear algebra solvers running on hardware accelerators for use in ECP libraries and applications. Other noteworthy outcomes included the development of sparse batched kernels, solvers/preconditioners, and the creation of interoperability in xSDK libraries with sparse and dense batched functions to benefit ECP applications.

FOREST

FOREST (Fostering an Open-Source Runtime Eco-System for the Task-based system PaRSEC) is supported by a \$300,000 award from the National Science Foundation's Pathways to Enable Open-Source Ecosystems (POSE) program. Led by the Innovative Computing Laboratory at the University of Tennessee, in collaboration with Stony Brook University and Saint Louis University, the one-year project (July 2025–June 2026) focuses on transitioning PaRSEC into a sustainable, community-driven open-source ecosystem. As the PaRSEC developer base has grown more geographically and institutionally distributed, FOREST establishes new governance, contribution, and collaboration models to ensure long-term viability and broader participation.

The initiative strengthens the technical and social infrastructure needed to expand PaRSEC's user and contributor communities across academia, national laboratories, and industry. By formalizing community practices, improving onboarding and documentation, and defining durable development processes, FOREST aims to position PaRSEC as a sustainable, high-performance runtime platform. In doing so, the project supports continued research in distributed task-based systems while enabling greater flexibility, portability, and performance for scientific applications targeting next-generation HPC architectures.

CETOP

ICL is participating in the Center for Edge of Tokamak Optimization (CETOP), a SciDAC effort led by the Princeton Plasma Physics Laboratory (PPPL). In addition to PPPL and UTK, the project brings together researchers from universities (Columbia, Utah State), national labs (Lawrence Livermore, Oak Ridge, Brookhaven, Argonne, National Renewable Energy Lab), and industry (Fiat Lux, General Atomics). CETOP aims to advance the current understanding of transient edge-localized modes (ELMs) in burning plasmas. ELMs are a key risk factor in tokamak fusion reactors, and controlling them is vital to creating a tokamak capable of deployment in a practical fusion power plant. Harnessing the power of HPC, CETOP will create new high-fidelity-based ELM stability maps and also perform stability analyses for some non-ELMing regimes. An AI/ML thrust will extrapolate from databases to guide ELM-free design optimization. ICL's main contribution to CETOP is improving GPU-capable preconditioning of solvers in the plasma physics codes using the Ginkgo library.

FIND OUT MORE AT
icl.utk.edu/research/cetop

Ginkgo

In the Ginkgo project, we develop high performance numerical linear algebra functionality reflecting the parallelism of modern HPC platforms. The focus is on GPU-accelerated systems, and Ginkgo can currently be used on AMD GPUs, Intel GPUs, and NVIDIA GPUs using backends written in the respective vendor languages. Ginkgo features a variety of Krylov solvers, sparse direct solvers, sophisticated preconditioners exposing fine-grain parallelism (including incomplete factorizations, incomplete sparse approximate inverses, and algebraic multigrid technology), mixed-precision algorithms, and preconditioned batched iterative solvers.

Ginkgo is implemented in modern C++ and is used by several popular simulation frameworks, including MFEM, SUNDIALS, deal.ii, HyTeg, openCARP, XGC, NEK5000, and OpenFOAM. Scalability of the Ginkgo backend to up to thousands of GPUs has been demonstrated on the Frontier and Perlmutter supercomputers. The Ginkgo library is open source under the 3-clause BSD license.

FIND OUT MORE AT
ginkgo-project.github.io

RESEARCH PROFILES

heFFTe

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 previous FFT libraries were not scalable on large heterogeneous machines with many nodes.

The main objective of the ECP FFT project was to design and develop the Highly Efficient FFTs for Exascale (heFFTe) library that provides fast and robust multidimensional FFTs for large-scale heterogeneous systems with multi-core processors and hardware accelerators. HeFFTe collects and leverages existing FFT capabilities while building a sustainable FFT library that minimizes data movements, optimizes MPI communications, overlaps computations with communications, and autotunes performance on various architectures and large-scale platforms. The current heFFTe version 2.4.1 release achieves very good scalability on exascale systems and performance close to 90% of the roofline peak.

FIND OUT MORE AT
icl.utk.edu/fft

MAGMA

Matrix Algebra on GPU and Multi-core Architectures (MAGMA) is a collection of next-generation linear algebra libraries for heterogeneous computing. MAGMA 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 component to heterogeneous computing systems. MAGMA enables applications to fully exploit the power of current hybrid systems of many-core CPUs and GPUs to deliver the fastest possible time to accurate solutions within given energy constraints.

MAGMA features LAPACK-compliant routines for multi-core CPUs enhanced with NVIDIA or AMD GPUs; there is also now pre-release support for most routines on Intel GPUs. MAGMA 2.9.0 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.

FIND OUT MORE AT
icl.utk.edu/magma

LAPACK/ScaLAPACK

Linear Algebra PACKage (LAPACK) and Scalable LAPACK (ScaLAPACK) are foundational libraries for solving dense linear algebra problems. ICL has contributed to the development and long-term maintenance of both packages, which underpin a vast ecosystem of scientific and engineering applications. LAPACK targets single-node systems and relies on highly optimized, multi-core Basic Linear Algebra Subprograms (BLAS) for performance, while ScaLAPACK extends this functionality to parallel, distributed-memory environments using BLAS, LAPACK, MPI, and BLACS.

LAPACK is written in Fortran and provides routines for solving systems of linear equations, least-squares problems, eigenvalue and singular value problems, along with associated matrix factorizations such as LU, Cholesky, QR, and SVD. Its algorithms are designed around block operations to efficiently exploit modern memory hierarchies and achieve high performance across diverse architectures. The project has been sponsored in part by MathWorks and Intel for many years.

Mixed-Precision Numerical Computing

With the rapidly expanding landscape of mixed- and multi-precision methods, the ongoing cross-pollination between HPC and machine learning or generative artificial intelligence is leading to intelligent computational steering of large-scale simulations. As these disparate scientific fields share hardware platforms, exploiting their wide range of computational modes has led to a proliferation of multiple representations of floating-point data. Taking full advantage of these developments is this project's main goal.

Against this backdrop, ICL's high-performance libraries (also those produced by internet-scale companies, hardware vendors, national laboratories, and academic institutions) spearhead the recent algorithmic progress in exploiting multiple precisions for increased efficiency in achieved performance, required communication, or optimized storage needs. The techniques used in this effort employ numerical representations such as limited-precision formats, quantized integers, and modular-precision ecosystems, among others. Note that the lossless or lossy compression approaches can independently benefit HPC codes as their algorithmic and accuracy advances are developed independently to the progress in mixed-precision methods.

PESA

Emerging HPC systems require innovations in existing infrastructure to deliver the best performance for science domains. The MPI 4.0 standard has also brought forward new opportunities for co-designing applications. These include partitioned point-to-point and collective operations and neighborhood collectives. Recent developments in GPU-based compression techniques provide an attractive option to optimize communication. With these advances, there is a critical need to update the commonly used tools and libraries that form the basis for the NSF's HPC cyberinfrastructure. The Performance Engineering Scientific Applications (PESA) with MVAPICH and TAU using Emerging Communication Primitives project undertakes this challenge and pursues new performance engineering avenues – by exploiting a co-design approach using the MPI_T tooling API – in the MVAPICH2 and TAU libraries with scientific applications. PESA focuses on two popular HPC applications: Anelastic Wave Propagation (AWP-ODC) and Highly efficient FFTs for Exascale (heFFTe). AWP-ODC is a highly scalable parallel finite-difference application with point-to-point operations that enables 3D earthquake calculations. HeFFTe, dominated by collective operations, is a massively parallel application that provides a scalable and efficient implementation of the widely used Fast Fourier Transform (FFT) operations.

FIND OUT MORE AT

icl.utk.edu/research/pesa

SLATE

For decades, ICL has applied algorithmic and technological innovations to pioneer, implement, and disseminate dense linear algebra software—including the LAPACK and ScaLAPACK libraries. The Software for Linear Algebra Targeting Exascale (SLATE) project represents the next generation of ICL's dense linear algebra software, designed from the ground up using modern C++ and targeting both CPU-only and GPU-accelerated systems within the exascale computing ecosystem.

Originally released in 1995, ScaLAPACK predates many of the architectural shifts that define modern HPC. Over the past three decades, the emergence of heterogeneous nodes, deep memory hierarchies, and massive parallelism has driven the need for a new design. SLATE was conceived as a successor to ScaLAPACK, introducing advanced algorithms such as communication-avoiding factorizations, GPU-aware MPI communication, and flexible tile-based data layouts. Developed through the Exascale Computing Project, SLATE emphasizes performance, scalability, and software sustainability, while maintaining portability across AMD, Intel, and NVIDIA GPU architectures in today's distributed-memory environments.

FIND OUT MORE AT

icl.utk.edu/slate

PLASMA

Parallel Linear Algebra Software for Multicores and Accelerators (PLASMA) implements a set of fundamental linear algebra routines using the latest updates to the OpenMP standard, including advanced tasking and GPU offload pragmas. PLASMA includes, among others, LAPACK-equivalent routines for solving linear systems of equations, linear least square problems, parallel BLAS, and parallel matrix norms.

Over the past decades, PLASMA has been used on various systems using Intel CPUs and coprocessors as well as AMD, IBM POWER, and ARM processors. As a research vehicle, PLASMA is an example of a modern design for new dense linear algebra algorithms. At the same time, PLASMA benefits from the continuous evolution of the OpenMP standard, which now includes offloading functionality and enables porting to hardware accelerators. The latest PLASMA release, version 25.5.27 from May 2025, added computation of eigenvalues and eigenvectors for symmetric/Hermitian algorithms using 2-stage algorithm, code generator and CMake integration, developer-oriented CI/CD improvements, and comprehensive Python-based tester as well as integration with external projects such as the ECP-sponsored Spack package manager.

FIND OUT MORE AT

github.com/icl-utk-edu/plasma

Sparse BLAS


This project is a testament to the undeniable importance of sparse matrix codes for scientific simulations and engineering applications across a wide range of disciplines. Although a Sparse BLAS standard was proposed decades ago, it has seen limited updates and adoption, preventing it from becoming a widely used reference for portable, high-performance sparse linear algebra on modern hardware. As architectures have grown increasingly heterogeneous, this gap has become more pronounced, motivating renewed community-driven efforts to modernize the interface.

Led by Hartwig Anzt, the Sparse BLAS initiative brings together scientists and software engineers from academia, national laboratories, and industry to define a next-generation interface for sparse linear algebra functionality. Over more than two years of weekly virtual meetings and multiple hybrid workshops, participants converged on a shared blueprint that emphasizes performance portability, extensibility, and vendor interoperability. The next phase of the project focuses on developing a reference implementation and enabling existing vendor libraries to interoperate through the new Sparse BLAS API, laying the foundation for a sustainable, standards-based ecosystem for sparse numerical computing.

FIND OUT MORE AT

icl.utk.edu/workshops/sparseblas2025

RESEARCH PROFILES



PERFORMANCE ANALYSIS

Maximizing the efficiency of HPC applications requires not only high-performance algorithms but also a deep understanding of how these algorithms interact with complex system architectures. Performance analysis plays a crucial role in this process by providing developers with the insights needed to optimize execution, identify inefficiencies, and adapt to evolving hardware and software environments. As modern supercomputers incorporate heterogeneous computing elements and deeper memory hierarchies, tools for analyzing performance must evolve to capture and interpret increasingly intricate performance behaviors.

MINCER

MINCER aims to provide comprehensive insights into distributed heterogeneous computing environments, enhancing research and education in platform reliability, efficiency, and security through three thrusts: (1) Improved AI system performance and utilization via advanced performance and power monitoring. (2) Unique network-related measurements to support research in anomaly detection, resilience, and energy efficiency. (3) Integration with platforms like Chameleon and the Open Science Grid to enhance understanding of cloud and edge computing, as well as the relationship between data transfer and computation in large-scale distributed computing.

MINCER offers open-source software to strengthen research in computing and networking, fostering broad community engagement in nationally significant systems-related research. Integration with research platforms such as ACCESS, Chameleon, and the Open Science Grid expands their utility for systems research and education. Through partnerships with minority-serving institutions and hands-on mentoring, MINCER introduces researchers and students to its capabilities, broadening participation and fostering inclusivity in systems research.

FIND OUT MORE AT

icl.utk.edu/research/mincer

PAPI

The Performance Application Programming Interface (PAPI) offers a universal interface and methodology for gathering performance counter information from diverse hardware and software components. This includes major CPUs, GPUs, accelerators, interconnects, I/O systems, power interfaces, and even virtual cloud environments. Collaborations with industry leaders like AMD, Cray/HPE, IBM, Intel, NVIDIA, and others ensure seamless integration of PAPI with new architectures as they are introduced or come close to release. As the PAPI component architecture expands, 3rd-party performance tools interfacing with PAPI gain the capability to measure data from these emerging architectures.

In 2025, PAPI released version 7.2.0b2 as a beta release and version 7.2.0 as the official major release. PAPI 7.2.0 introduces two new components: (1) `rocp_sdk`, which supports AMD GPUs and APUs via the ROCprofiler-SDK interface, and (2) `topdown`, which provides proper support for Intel Top-Down metrics. PAPI 7.2.0 also introduces preset events for non-CPU devices, starting with CUDA events. In addition, the component code has been extended to include a statistics qualifier (e.g., for CUDA events), offering more concise and functional output in the `papi_native_avail` utility.

FIND OUT MORE AT

icl.utk.edu/papi

SPADE

ICL is breaking new ground with SPADE, an NSF-funded "Frameworks" project titled "Scalable Performance and Accuracy analysis for Distributed and Extreme-scale systems (SPADE)." ICL leads this collaborative project with Heike Jagode as the PI and Anthony Danalis as co-PI. The project also involves partnerships with the University of Maine and the University of Texas, El Paso. Spanning four years, the project started on September 15, 2023, with a total budget of \$3.5M (ICL's share is \$2.1M).

The SPADE project is dedicated to enhancing monitoring, optimization, evaluation, and decision-making functions for extreme-scale systems, catering to the needs of the HPC and scientific applications communities. As HPC resources evolve towards extreme scale, there's a growing necessity for integrated frameworks to tackle performance and reliability issues. Through extending support for heterogeneity and scalability across various computing platforms, and employing the established PAPI performance monitoring library, SPADE aims to provide the necessary software and APIs to effectively address the demands of scientific and machine learning applications while exploring new accuracy versus performance trade-offs with low-precision floating-point types.



VIEW POSTER AT
tiny.utk.edu/spade-poster-2025



Jagode, H., S. Moore, V. Weaver, A. Danalis, and C. Lauter
SPADE: Scalable Performance and Accuracy analysis for Distributed and Extreme-scale systems
 NSF CSSI Meeting, July 2025

STEP

The Software Tools Ecosystem Project (STEP) is a DOE ASCR-funded initiative, with ICL's Heike Jagode as one of the co-PIs. As part of its efforts, STEP organized a series of three town hall meetings throughout the summer of 2023. The primary objective of these meetings has been to formulate a strategic action plan for DOE/ASCR, with a specific focus on the long-term stewardship and advancement of the HPC tools ecosystem.

As part of the STEP initiative, the PAPI project will receive five years of funding totaling \$1.375M. The project aims to develop monitoring capabilities for emerging hardware technologies and create innovative abstractions for software-defined events. PAPI follows sustainable software practices and collaborates with stakeholders to meet the specific needs of the HPC community.

FIND OUT MORE AT
ascr-step.org

cyPAPI

cyPAPI is a Python interface to the Performance Application Programming Interface (PAPI) that enables direct, low-overhead access to hardware performance counters from Python applications. cyPAPI is implemented using Cython as a thin, optimized wrapper around the PAPI interface, providing Pythonic abstractions while preserving the accuracy and efficiency of native PAPI measurements. cyPAPI supports modern CPUs and GPUs across vendors and architectures and is designed to integrate naturally into Python-based HPC, data analytics, and AI/ML workflows.

cyPAPI enables first-person, multi-threaded hardware profiling from within Python, operating both inside and outside the Global Interpreter Lock and aligning with emerging Python concurrency models. The project also exposes PAPI's Software Defined Events (SDE) infrastructure, allowing applications to define and collect custom performance metrics. By making hardware profiling more accessible, portable, and vendor-independent, cyPAPI lowers the barrier to performance analysis and energy-aware optimization for next-generation scientific and machine learning applications.

FIND OUT MORE AT
github.com/icl-utk-edu/cyPAPI

RESEARCH PROFILES



DISTRIBUTED COMPUTING

The increasing complexity of modern HPC systems, characterized by massive concurrency and multi-tiered memory hierarchies, presents a formidable challenge for software developers. Applications must efficiently coordinate computations across thousands or even millions of processing elements while minimizing communication overhead and ensuring resilience against failures. The scale and heterogeneity of these systems necessitate innovative approaches to workload distribution, fault tolerance, and dynamic resource management.

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. DPLASMA is designed to deliver sustained performance for distributed systems, where each node features multiple sockets of multi-core processors and, if available, accelerators like NVIDIA, AMD, or Intel GPUs. DPLASMA achieves this objective by deploying PLASMA tile-based 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
icl.utk.edu/dplasma

EPEXA

A collaborative project involving Virginia Tech, Stony Brook, and ICL, 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 co-design, the EPEXA team hardens a successful research prototype into an accessible, production-quality programming model that leverages 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 and push for its community use. Specifically, the new powerful data-flow programming model and associated parallel runtime directly address multiple challenges scientists face as they leverage rapidly changing computer technologies—including current massively parallel, hybrid, and many-core systems.

FIND OUT MORE AT
icl.utk.edu/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.

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. It demonstrated a clear advantage regarding 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 has been fully integrated into Open MPI.

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 directed acyclic graph (DAG) of tasks with edges designating the data dependencies. DAGs are represented in a compact format that can be queried to discover data dependencies in a distributed and scalable fashion—a drastic shift from today's programming models, which are based on the replicated 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 DAGs and integrate them into legacy codes, a runtime library to schedule 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. In recent years, PaRSEC has underpinned the deployment of multiple Gordon Bell Prize finalist runs, and in 2024, the PaRSEC-enabled entry received the ACM Gordon Bell Prize for Climate Modelling.

FIND OUT MORE AT
icl.utk.edu/parsec

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 one process's address space to another 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) to build the best MPI library available.

A completely new MPI 4.1-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
open-mpi.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 failures rendered the completions impossible.

Using the constructs defined by ULFM, applications and libraries drive the recovery of the parallel application and execution environment state. Consistency issues resulting from failures are addressed according to an application's needs, and the recovery actions are limited to MPI communication objects. Many application types and middleware are built on top of ULFM to deliver scalable fault tolerance. Notable additions include the CoArray Fortran language and SAP databases. ULFM software is available in recent versions of MPICH and Open MPI.

FIND OUT MORE AT
fault-tolerance.org

RESEARCH PROFILES



HPC BENCHMARKING

Benchmarking serves as a critical mechanism for evaluating and comparing the performance of computing systems. ICL has been a leader in the development of benchmark codes for scientific workloads, most famously with High-Performance Linpack (HPL). As HPC architectures continue to evolve with novel processors, accelerators, and interconnect technologies, ICL remains at the forefront of benchmarking research, refining methodologies to capture the true capabilities of modern and future supercomputing platforms.

HPCG

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

The HPCG 3.1 reference code was released in March 2019. This release positioned HPCG to represent modern PDE solvers better and made it easier to run HPCG on production supercomputing installations. The reference version is accompanied by multiple binary or source code releases from AMD, ARM, Intel, and NVIDIA, which are carefully optimized for these vendors' respective hardware platforms. The current HPCG performance list was released at SC25 and features over 100 entries across the supercomputing landscape that are also tracked by TOP500, giving a unique opportunity to compare the bookend performance levels of these machines.

FIND OUT MORE AT
hpcg-benchmark.org

HPL

The High Performance LINPACK (HPL) benchmark solves a dense linear system of equations, $Ax=b$, in double precision (64-bit arithmetic) on distributed-memory computers. Written in ANSI C, it requires MPI and BLAS implementations. HPL is often the first program run on large HPC systems, with carefully optimized versions available from major HPC vendors.

When HPL was updated in 2018 to version 2.3, the reporting of the benchmark results was improved both in terms of accuracy and ease of installation: configuring and building the code on newer HPC platforms was streamlined. Since then, the development of HPL has moved to GitHub. The benchmark now features more detailed reporting of the scaled residual of the linear system's solution as well as the achieved performance number. Additionally, the code base no longer uses deprecated MPI functions. Finally, HPL now allows users to configure and build the code using the GNU Autotools portability system.

FIND OUT MORE AT
icl.utk.edu/hpl



Representatives from Lawrence Livermore National Laboratory were presented with certificates for HPCG, HPL-MxP, and TOP500 at the TOP500 BoF at SC25. The El Capitan system ranked #1 across all 3 benchmark lists in November 2025.

HPL-MxP

The HPL-MxP benchmark seeks to highlight the emerging convergence of HPC and artificial intelligence workloads. While traditional HPC focused on simulation runs for modeling phenomena in physics, chemistry, biology, and so on, the mathematical models that drive these computations require, for the most part, 64-bit accuracy. On the other hand, the machine learning methods that fuel advances in AI achieve desired results at 32-bit and even lower floating-point precision formats. 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-MxP strives to unite these two realms by delivering a blend of modern algorithms and contemporary hardware while simultaneously connecting to the solver formulation of the decades-old HPL framework of benchmarking the largest supercomputing installations in the world.

FIND OUT MORE AT
hpl-mxp.org

TOP500

With over three decades of tracking the progress of high performance computing, the TOP500 lists continue to provide a reliable historical record of supercomputers worldwide. The lists lay out critical HPC metrics across all of its 500 machines and draw 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 machines' performance using irregular accesses to memory and fine-grain recursive computations—the very factors that dominate real-world, large-scale scientific workloads.


In November 2025, the 66th TOP500 list was unveiled at SC25. El Capitan at Lawrence Livermore National Laboratory retained the No. 1 position and posted an updated HPL performance of 1.809 exaflop/s. Frontier at Oak Ridge National Laboratory remained second at 1.353 exaflop/s, while Aurora at Argonne National Laboratory held third at 1.012 exaflop/s. The biggest development this year was the full debut of JUPITER Booster at the Jülich Supercomputing Centre in Germany, which reached 1.000 exaflop/s on HPL and entered the list at No. 4. With that result, JUPITER became the fourth exascale system on the TOP500 and the first exascale machine outside the United States.

TOP500 November 2025

| | | | |
|----------|--|----------------|-----------------------------|
| 1 | El Capitan DOE/NNSA/LLNL, USA | 1,809.0 | R _{MAX} PFLOP/S |
| 2 | Frontier DOE/SC/ORNL, USA | 1,353.0 | |
| 3 | Aurora DOE/SC/ANL, USA | 1,012.0 | |
| 4 | JUPITER Booster EuroHPC/FZJ, Germany | 1,000.0 | |
| 5 | Eagle Microsoft Azure, USA | 561.2 | |



FIND OUT MORE AT
top500.org



PEOPLE

ICL STAFF & STUDENTS DECEMBER 2025

ICL's achievements reflect the dedication and technical strength of its staff and students. Their research and collaborative efforts form the backbone of the lab's contributions to high-performance computing.



**Ahmad
ABDEFATTAH**
INTERIM ASSOCIATE DIRECTOR /
RESEARCH ASSISTANT
PROFESSOR



**Hartwig
ANZT**
RESEARCH ASSOCIATE
PROFESSOR




**Daniel
BARRY**
RESEARCH SCIENTIST



**Natalie
BEAMS**
RESEARCH ASSISTANT
PROFESSOR




**George
BOSILCA**
RESEARCH ASSOCIATE
PROFESSOR



**Treece
BURGESS**
RESEARCH ASSOCIATE II



**Giuseppe
CONGIU**
VISITING SCHOLAR



**Jack
DONGARRA**
RESEARCH PROFESSOR
EMERITUS



**Dayle
FLANIGAN**
FINANCIAL ASSISTANT




Heike JAGODE
INTERIM DIRECTOR /
RESEARCH ASSOCIATE
PROFESSOR



Tracy LEE
CONSULTANT



Piotr LUSZCZEK
RESEARCH ASSOCIATE
PROFESSOR



Maksim MELNICHENKO
GRADUATE RESEARCH
ASSISTANT



Tatiana MELNICHENKO
UT STUDENT ASSISTANT



Gerald RAGGHIANI
TECHNICAL SERVICES
LEADER



Waldhys RODOLI
BUSINESS MANAGER



David ROGERS
COMMUNICATIONS
SERVICES LEADER



Joseph SCHUCHART
VISITING SCHOLAR



Tokey TAHMID
RESEARCH ASSOCIATE



Tonya WIMBERLEY
FINANCIAL ASSOCIATE



Jun WOUN
GRADUATE RESEARCH
ASSISTANT

PEOPLE

ALUMNI

| | | | | |
|---------------------|------------------------------|--------------------------|------------------------|--------------------------|
| Maksims ABALENKOV | Laura BLACK | Matthias COLIN | Christoph FABIANEK | Amina GUERMOUCHE |
| Carolyn AEBISCHER | Susan BLACKFORD | Charles COLLINS | Graham FAGG | Jason GURLEY |
| Sudesh AGRAWAL | Wesley BLAND | Rephael CONGMON | Hadeer FARAHAT | Fritz GÖBEL |
| Bivek AGRAWAL | Kartheek BODANKI | Stephanie MOREAUD COOPER | Mathieu FAVERGE | Bilel HADRI |
| Emmanuel AGULLO | David BOLT | Tom CORTESE | Diana FAYAD | Hunter HAGEWOOD |
| Kadir AKBUDAK | Fernando BOND | Camille COTI | Shengzhog FENG | Azzam HAIDAR |
| Mohammed AL FARHAN | Aurelien BOUTEILLER | Jason COX | Don FIKE | Christian HALLOY |
| Rabab AL-OMAIRY | Carolyn BOWERS | Sam CRAWFORD | Salvatore FILIPPONE | Julian HALLOY |
| Jennifer ALLGEYER | Vincent BRIDONNEAU | David CRONK | Teresa FINCHUM | Sven HAMMARLING |
| Wes ALVARO | Barry BRITT | Javier CUENCA | Anna FINCHUM | J. Mike Hammond |
| Ed ANDERSON | Randy BROWN | Manoel CUNHA | Mike FINGER | Caleb HAN |
| Robert ANDERSON | Cade BROWN | Yuanshun (SHAUN) DAI | Jamie FINNEY | Hanumantharayappa |
| Daniel ANDRZEJEWSKI | Murray BROWNE | Anthony DANALIS | Markus FISCHER | Hidehiko HASEGAWA |
| Thara ANGSKUN | Cynthia BROWNE | Cricket DEANE | Clement FLINT | Satomi HASEGAWA |
| Papa ARKHURST | Bonnie BROWNE | Remi DELMAS | Len FREEMAN | Chris HASTINGS |
| Dorian ARNOLD | Giuseppe BRUNO | Sameer DESHMUKH | Xiaoquan FU | Blake HAUGEN |
| Rizwan ASHRAF | Antonin BUKOVSKY | Frederic DESPREZ | Erika FUENTES | David HENDERSON |
| Cedric AUGONNET | Greg BUNCH | Jun DING | Karl FUERLINGER | Greg HENRY |
| Alan AYALA | Alfredo BUTTARI | Jin DING | Megan FULLER | John HENRY |
| Marc BABOULIN | Pierluigi (PG) CAMBIE-FABRIS | Ying DING | Edgar GABRIEL | Thomas HERAULT |
| Matthew BACHSTEIN | Anthony CANINO | Martin DO | Lisa GAEDKE-MERZHAUSER | Julien HERRMANN |
| Paul BAGWELL | Domingo GIMENEZ CANOVAS | Simplice DONFACK | Lynn GANGWER | Holly HICKS |
| Zhaojun BAI | Chongxiao CAO | Leon DONG | Tracy GANGWER | Alexandra HICKS-HARDIMAN |
| Ashwin BALAKRISHNAN | Qinglei CAO | Tingxing DONG | Nathan GARNER | Sid HILL |
| Richard BARRETT | Earl CARR | Nick DONGARRA | Kelley GARNER | Mark HIPSHIRE |
| Alex BASSI | Henri CASANOVA | David DOOLIN | Tina GARRISON | Tomoyuki HIROYASU |
| John BATSON | Cedric CASTAGNEDE | Joe DORRIS | Adriana GARTIES | George HO |
| David BATTLE | Tony CASTALDO | Andrew DOWNEY | Mark GATES | Josh HOFFMAN |
| Micah BECK | Sebastien CAYROLS | Mary DRAKE | Peter GAULTNEY | Reazul HOQUE |
| Daniel BECKER | Ramkrishna CHAKRABARTY | Julio DRIGGS | Christoph GEILE | Jeff HORNER |
| Dulcenea BECKER | Sharon CHAMBERS | Brian DRUM | Jean PATRICK GELAS | Mitch HORTON |
| Adam BEGUELIN | Ali CHARARA | Peng DU | Boris GELFEND | Yan HUANG |
| Annamaria BENZONI | Zizhong CHEN | David EBERIUS | Damien GENET | Harry HUGHES |
| Tom BERRY | Jaeyoung CHOI | Eduardo ECHAVARRIA | Jonathan GETTLER | Aurelie HURAUULT |
| Vincent BERTHOUX | Wahid CHRABAKH | Victor EIJKHOUT | Scott GIBSON | Chris HURT |
| Scott BETTS | Eric CLARKSON | Brett ELLIS | Eric GREASER | Paul JACOBS |
| Nikhil BHATIA | Andy CLEARY | Shawn ERICSON | Stan GREEN | Emmanuel JEANNOT |
| Daniel BIELICH | Michelle CLINARD | Zachary EYLER-WALKER | Alice GREGORY | Weizhong JI |
| Noel BLACK | Vincent COHEN-ADDAD | Lisa EZZELL | Thomas GRÜTZMACHER | Yulu JIA |

| | | | | |
|---------------------------|------------------------|--------------------------|---------------------|-------------------|
| Weicheng JIANG | Yicheng LI | Phuong NGUYEN | Farial SHAHNAZ | Mike WALTZ |
| Song JIN | Jiali LI | Jonas NILSSON | Hejer SHAIK | Robert WALTZ |
| Patrick JOHANSSON | Neil LINDQUIST | Jakob OESTERGAARD | Daniel SHARP | Xiaoyang WANG |
| Aral JOHNSON | Chaoyang LIU | Anustuv PAL | Brian SHEELY | Bob WARD |
| Matt JOHNSON | Kevin LONDON | Caroline PAPADOPOULOS | Zhiao SHI | Jerzy WASNIEWSKI |
| Sean JOLLY | Matt LONGLEY | Leelinda PARKER | Sergei SHINKAREV | Vince WEAVER |
| Jan JONES | Florent LOPEZ | Thibault PARPAITE | Wissam SID LAKHDAR | Aaron WELCH |
| Kim JONES | Nuria LOSADA | Thananon PATINYASAKDIKUL | Majed SIDANI | Scott WELLS |
| Vijay JOSHI | Hatem LTAIEF | Dilip PATLOLLA | Shilpa SINGHAL | David WEST |
| Khairul KABIR | Yuechao LU | Andy PEARSON | Leighanne SISK | Jonah WESTON |
| Venkata KAKANI | Daniel LUCIO | Yu PEI | Matt SKINNER | R. CLINT WHALEY |
| Ajay KALHAN | Richard LUCZAK | Paul PELTZ | Joan SNODERLY | Jody WHISNANT |
| Balajee KANNAN | Xi LUO | Deborah PENCHOFF | Peter SOENDERGAARD | James WHITE |
| Madhuri KASAM | Teng MA | Theresa PEPIN | Raffaele SOLCA | Scotti WHITMIRE |
| Kiran KUMAR KASICHAYANULA | Sticks MABAKANE | Antoine PETITET | Gwang SON | Frank WINKLER |
| Ajay KATTA | Ethan MANASS | Peter PHAM | Fengguang SONG | Susan WO |
| David KATZ | Robert MANCHEK | Gregoire PICHON | Thomas SPENCER | Felix WOLF |
| Joshua KELLY | Gabriel MARIN | Vlado PJESIVAC | Jeffrey STEILL | Stephen WOOD |
| Supriya KILAMBI | Noah MARKS | Jelena PJESIVAC-GRBOVIC | Erich STROHMAIER | Eleigha WRANCHER |
| Youngbae KIM | Tushti MARWAH | James S. PLANK | Dalal SUKKARI | Jiayi WU |
| Myung HO KIM | Theo MARY | Tim POORE | Xiaobai SUN | Wei WU |
| Jenya KIRSHEIN | Ian MASLIAH | Roldan POZO | Martin SWANY | Panruo WU |
| Cindy KNISLEY | Donald MCCASLAND | Farzona PULATOVA | Daisuke TAKAHASHI | You WU |
| Hannah KNOCH | Tyler MCDANIEL | Martin QUINSON | Judi TALLEY | Qiu XIA |
| Michael KOLATIS | Paul MCMAHAN | Tammy RACE | Ronald TAM | Tinghua XU |
| Anara KOZHOKANOVA | Eric MEEK | Tracy RAFFERTY | Yuan TANG | Ichitaro YAMAZAKI |
| Moritz KREUTZER | James MEYERING | Sangamesh RAGATE | Chunyan TANG | Tao YANG |
| Chandra KRINTZ | Jeremy MILLAR | James RALPH | Yusuke TANIMURA | Erlin YAO |
| Tilman KUESTNER | Michelle MILLER | Ganapathy RAMAN | Keita TERANISHI | Asim YARKHAN |
| Jakub KURZAK | Daniel MISHLER | Kamesh RAMANI | Dan TERPSTRA | Kevin YE |
| Krerkchai KUSOLCHU | Cindy MITCHELL | Mei RAN | Joe THOMAS | Jin YI |
| Coire KYLE | Stuart MONTY | Arun RATTAN | John THURMAN | Haihang YOU |
| Brian LAROSE | Shirley MOORE | Sheri REAGAN | Francoise TISSEUR | Lamia YOUSEFF |
| Amanda LAAKE | Keith MOORE | Mike REYNOLDS | Stanimire TOMOV | Brian ZACHARY |
| Xavier LACOSTE | Terry MOORE | George RHINEHART | Jude TOTH | Dmitry ZAITSEV |
| Julien LANGOU | Erik MOORE | Jon RICHARDSON | Bernard TOURANCHEAU | Omar ZENATI |
| Julie LANGOU | Robert MORGAN | Stephen RICHMOND | Yaohung TSAI | Yuanlei ZHANG |
| Jeff LARKIN | Kishan MOTHERAMGARI | Yves ROBERT | Volodymyr TURCHENKO | Junlong ZHAO |
| Patrick HON MAN LAU | Steven MOULTON | Ken ROCHE | Lauren VACA | Yong ZHENG |
| Frank LAUER | Phil MUCCI | Andrew ROGERS | Phil VACCARO | Dong ZHONG |
| Valentin LE FEVRE | Daichi MUKUNOKI | Tom ROTHROCK | Sathish VADHIYAR | Luke ZHOU |
| DongWoo LEE | Matthew NABITY | Tom ROWAN | Robert VAN DE GEIJN | Min ZHOU |
| Pierre LEMARINIER | Shankar NARASIMHASWAMI | Narapat (OHM) SAENGPATSA | Chad VAWTER | |
| Todd LETSCHE | Rajib NATH | Kiran SAGI | Eugene VECHARYNSKI | |
| Sharon LEWIS | Fernando NAVARRO | Evelyn SAMS | Scott VENCKUS | |
| Xiang LI | John NELSON | Daniel SCHULTZ | Antoine VERNONIS | |
| Weiran LI | Donnie NEWELL | Ken SCHWARTZ | Reed WADE | |
| Yinan LI | Peter NEWTON | Keith SEYMOUR | Michael WALTERS | |

PARTNERSHIPS

INDUSTRY PARTNERSHIPS



GOVERNMENT & ACADEMIC PARTNERSHIPS



INTERNATIONAL COLLABORATIONS



Charles University
PRAGUE, CZECH REPUBLIC



Claude Bernard University Lyon 1
LYON, FRANCE



ENS de Lyon
LYON, FRANCE



Jülich Supercomputer Center
JÜLICH, GERMANY



INRIA
FRANCE



Karlsruhe Institute of Technology
KARLSRUHE, GERMANY



King Abdullah University of Science and Technology
THUWAL, SAUDI ARABIA



Prometeus GmbH
MANNHEIM, GERMANY



RIKEN
WAKŌ, JAPAN



Technical University of Munich
HEILBRONN, GERMANY



Umeå University
UMEÅ, SWEDEN



University of Bordeaux
BORDEAUX, FRANCE



University of Cape Town
CAPE TOWN, SOUTH AFRICA



University of Manchester
MANCHESTER, ENGLAND



University of Paris-Saclay
PARIS, FRANCE



University of Paris-Sud
PARIS, FRANCE



University of Tsukuba
TSUKUBA, JAPAN



Vienna University of Technology
VIENNA, AUSTRIA

ACKNOWLEDGMENTS

ICL's work has been made possible through the support of the federal agencies that are charged with allocating public research funding, and we thank the following agencies:



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





INNOVATIVE COMPUTING LABORATORY

FIND OUT MORE AT icl.utk.edu/