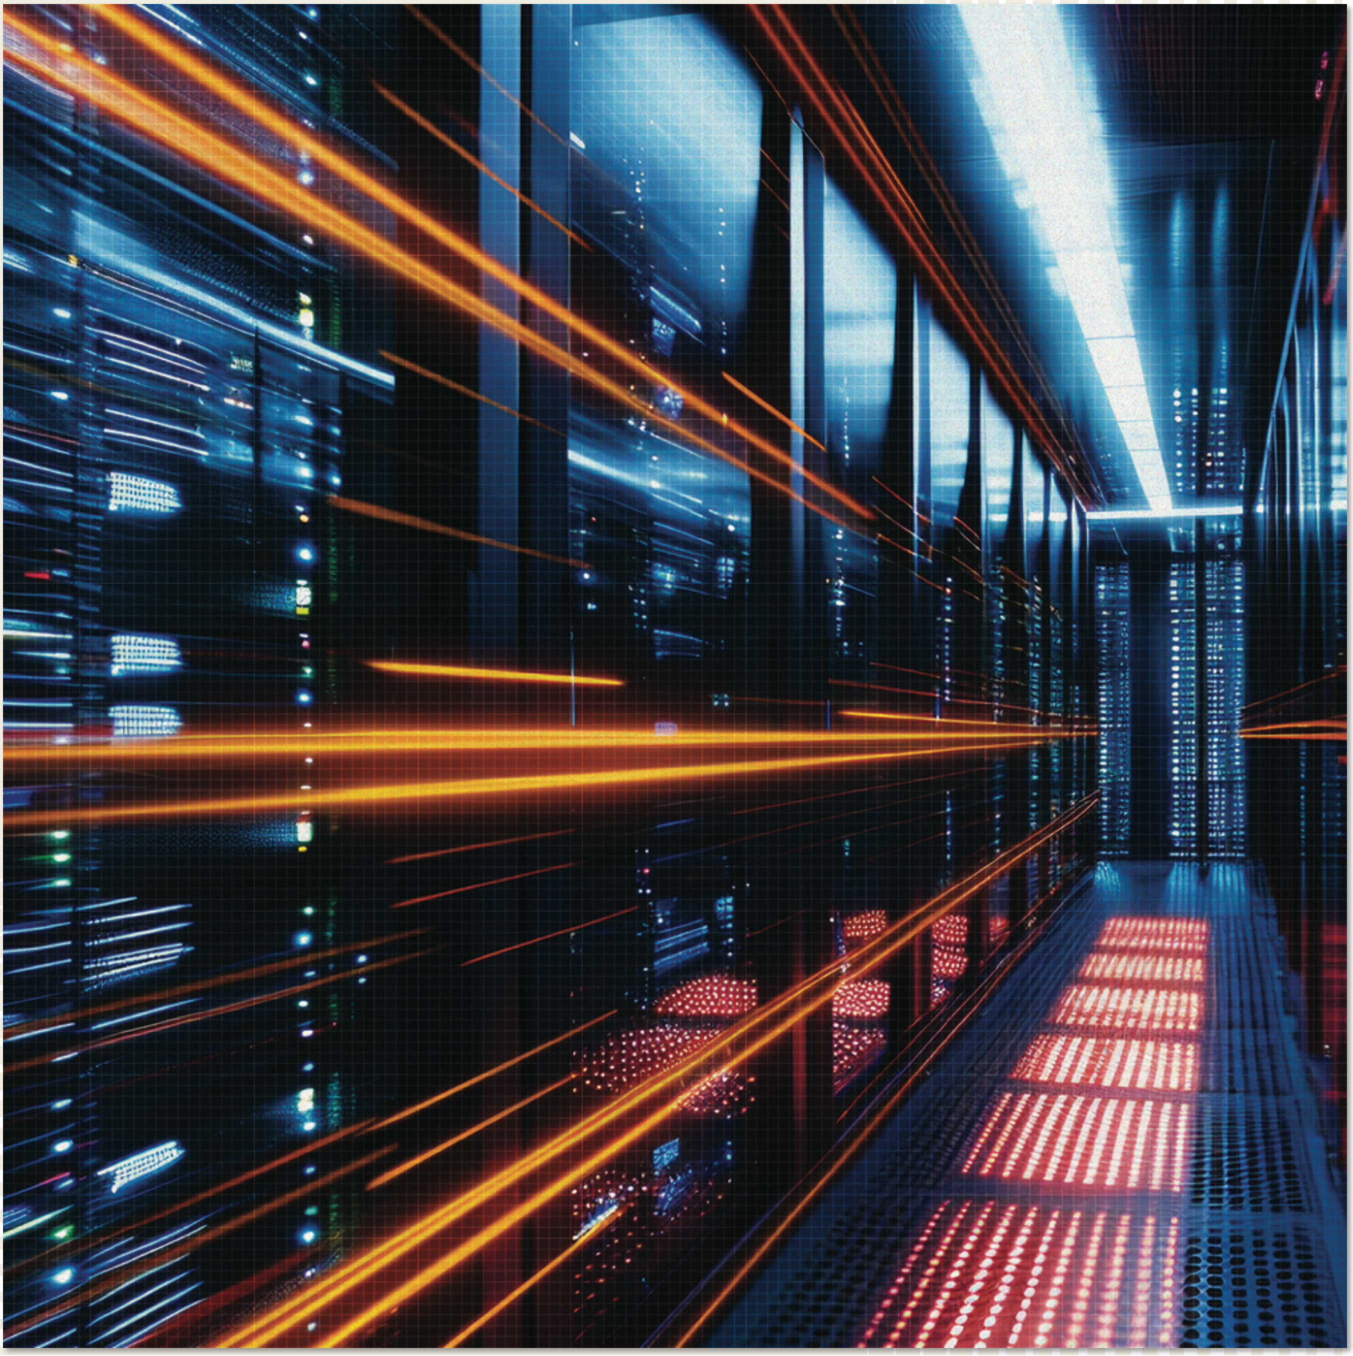

THE UNIVERSITY OF TENNESSEE KNOXVILLE

ICL 2023/24 REPORT

INNOVATIVE COMPUTING LABORATORY *2023/24 REPORT*

ICL 2023/24 REPORT

FROM THE DIRECTOR


Hartwig Anzt

REFLECTING ON this year, we recognize that 2023 will be remembered in ICL history as a year packed with success, change, new initiatives, and the culmination of 7 years of work on the Department of Energy's Exascale Computing Project (ECP). We could finally demonstrate (and debug!) our software on the first exascale machine, Frontier, at Oak Ridge National Laboratory.

ICL was fortunate to have had or contributed to 8 projects within ECP: SLATE developed distributed, GPU-accelerated dense linear algebra; Ginkgo developed GPU-accelerated sparse linear algebra; heFFTe provided distributed, GPU-accelerated 3D FFT; Exa-PAPI developed novel performance metrics and monitoring capabilities; Open MPI developed distributed communication; PaRSEC for distributed runtime scheduling; xSDK focused on math library interoperability and concentrated efforts on batched sparse and mixed precision functionality; and CEED served as a center for high-order finite-element tensor computations. At the end of ECP, several of ICL's software products and standardization efforts had become integral components of multiple high-profile ECP application projects, including plasma physics simulations in the WarpX and XGC projects, chemistry simulations in NWChemEX, molecular dynamics in LAMMPS, material science in MEUMAPPS-SS, particle simulations through CoPA, finite element simulations in MFEM, nonlinear solvers in SUNDIALS, sparse-direct solvers in STRUMPACK, and advancing the MPI communication standard. I am thankful for all ICL members contributing to ECP, and I am sure this gives ICL a headstart in future DOE initiatives.

We started 2023 with a well-attended winter reception that offered the opportunity to bid farewell to ICL's Assistant Director, Joan Snoderly. With Joan's retirement from ICL, we merged the ICL administrative support group and the ICL technical support group to become one team, the ICL operations group, reporting directly to the director.

With the end of ECP in sight, we organized an ICL directors' retreat in February to develop a strategy for post-ECP ICL. Team building activities helped us to identify challenges and opportunities ahead of us.

Later in the spring, we held an ICL/ORNL workshop to strengthen that partnership and create new collaborations. The summer brought the traditional ICL retreat at Park Vista. A new format based on breakout sessions initiated numerous new initiatives and demonstrated ICL's adaptability to new research fields.

Fall 2023 brought a long list of visitors to ICL, including Veronica Montanaro, Andy Adelman, Sonali Mayani, Brieuc Nicolas, Lisa Gaedke-Merzhäeuser, Thomas Grützmacher, You Wu, Terry Cojean, Tobias Ribizel, Mike Tsai, Albert Njoroge Kahira, and Yves Robert. International research exchange has always been a hallmark of ICL, and I am thankful for all the visitors and ICL members who helped make the research visits successful.

The Sparse BLAS Workshop just before SC23 started a new initiative at ICL to define a standard for sparse linear algebra operations. Attendees from academic institutions, national laboratories, and major industry partners, including ARM, NVIDIA, AMD, Intel, and MathWorks, discussed different approaches and held hackathons to develop a reference implementation.

The end of ECP was a bittersweet moment. We celebrate all the work accomplished over seven years. Unfortunately, with the end of ECP, ICL also shrank, and we said goodbye to many people. As has been true throughout the history of ICL, people come to obtain training as graduate students, post-docs, and research scientists and then go out into industry and academia to advance their careers and contribute to cutting-edge HPC. Many ICL alumni continue to work in the same field and help to ensure that the ideas developed here spread and thus have a broad impact. We look forward to continuing collaborations with the ICL alumni.

The end of the year is also the end for me in the office of the ICL Director. Although I am very grateful for the position, trust, and support that has been shown to me, it is with a heavy heart that I have decided to take on a new role at the Technical University of Munich in Germany, beginning in 2024. I am sure that ICL is well-positioned to thrive and grow over the next few years, and I will continue to contribute to ICL's success as much as it is still in my power!

Hartwig Anzt

Hartwig ANZT
ICL DIRECTOR, 2022–2023

Mark Gates

WE ARE sad to say goodbye to Hartwig Anzt, who directed ICL for the past 18 months through the completion of ECP. We will continue collaborating with Hartwig through the Ginkgo project, which has seen an uptake in several DOE applications and math libraries, the sparse BLAS standardization, and undoubtedly other efforts.

Starting in 2024, I am taking on the role of Interim ICL Director during the transition period until the next permanent ICL Director. I came to ICL in 2011 as a post-doc after graduating from the University of Illinois, Urbana-Champaign, never dreaming that I would spend over 12 years here and have the opportunity to lead ICL. It has been an honor to work with the many people at ICL.

Looking forward to the coming year, we will continue to sustain and leverage our successes in ECP, foster collaborations with applications using ICL software, and look for new opportunities. Artificial Intelligence (AI) dominates the current market, and ICL is well-positioned to contribute to AI software through our linear algebra libraries and performance analysis utilities. Our Surrogates project is advancing the state-of-the-art in neural networks. We have several newly funded projects using AI for computational science, including AI for Environmental Remediation, Biochar Enhanced Ecosystems for Energy, Critical Minerals, and the newly awarded Consortium for Nuclear Forensics (CNF).

Our recent involvement in AI applications and computational science, building upon our historical leadership in HPC, uniquely positions us to lead advances in many areas of science and engineering.

Mark Gates

Mark GATES
INTERIM ICL DIRECTOR

INTRODUCTION

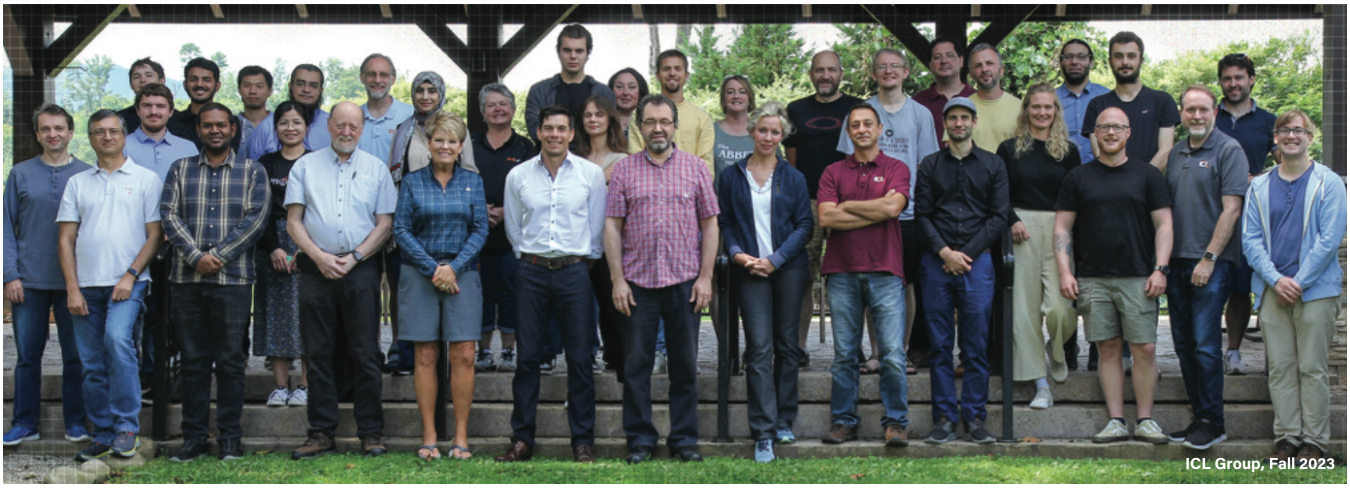
ICL Group, Fall 2023

ENABLING TECHNOLOGIES FOR HIGH-PERFORMANCE COMPUTING

SINCE 1989

THE INNOVATIVE COMPUTING LABORATORY (ICL) is a computer science research and development group in the Department of Electrical Engineering and Computer Science (EECS), situated in the heart of the University of Tennessee's Knoxville campus. ICL has a rich history that has spanned over 30 years of advancements in high-performance computing. These advancements have brought ever-increasing demands for parallelism, data management, energy efficiency, and resilience to the field.

ICL's work, which has evolved and expanded to address these challenges, 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-performance computing. In addition, ICL's portfolio of expertise includes high-performance parallel and distributed computing—with a focus on 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 such as life sciences, climate science, earthquake prediction, energy exploration, combustion and turbulence, advanced materials science, and drug design, among others.

FIND OUT MORE AT icl.utk.edu/

AREAS OF RESEARCH

THE NUMERICAL LINEAR algebra libraries that serve as the backbone of scientific computing applications also serve as the strong basis of ICL's research expertise. As hardware architectures and computing trends change over time, it is essential that these libraries also be improved and modernized to achieve high levels of performance and efficiency on new platforms. ICL has a long history of developing and standardizing these libraries in this way, and it has multiple projects under development in this arena.

AS THE NUMBER of cores, nodes, and other components in HPC systems continues to grow, applications require runtime systems that can exploit greater levels of parallelism. Moreover, as the number of components increases, it becomes more difficult to detect and recover from failures and protect the integrity of data. These challenges must be addressed with fault-tolerant software and hardware, and the escalating communication traffic that they generate necessitates smarter and more efficient message passing standards and practices.

ICL'S PERFORMANCE EVALUATION tools enable programmers to see correlations between the structure and performance of their code. These correlations are important for performance tuning, optimization, debugging, and finding and correcting performance bottlenecks. Similarly, 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 around the world.

THE PURSUIT OF the world's most pioneering research demands the integration of the most sophisticated computational expertise available. Over recent years, ICL has become more involved in the emerging field of computational science. This inherently interdisciplinary field combines high-performance computing expertise with domain-specific knowledge to tackle complex scientific challenges in a wide range of areas like climate modeling, biomedical research, and materials science.

HISTORY

JACK DONGARRA ESTABLISHED ICL in 1989 when he received a dual appointment as a Distinguished Professor at UTK and as a Distinguished Scientist at Oak Ridge National Laboratory. Over three decades, ICL has established itself as an internationally recognized research laboratory in advanced scientific and high-performance computing. The following timeline highlights just some of the significant projects and initiatives from ICL that have significantly influenced the landscape of high-performance and scientific computing.

1989

Parallel Virtual Machine (PVM) was a parallel networking tool that enabled a user to leverage a network of heterogeneous Unix and Windows machines as a single distributed parallel processor. Widely used in the 1990s and early 2000s, PVM played a significant role in the development of distributed computing techniques and applications and was awarded with a R&D 100 Award in 1994.

1990

The **Level-3 Basic Linear Algebra Subprograms (BLAS)** specification extended earlier BLAS versions, introducing more complex and efficient routines for matrix-matrix operations, which significantly enhanced computational performance for high-level scientific and engineering applications.

1992

The **Basic Linear Algebra Communication Subprograms (BLACS)** project was created to make linear algebra applications easier to program and more portable. BLACS provided a standard communication layer critical for the development and performance optimization of distributed matrix computations.

Still developed today, the **Linear Algebra Package (LAPACK)** is a foundational software library for numerical linear algebra. It provides routines for solving systems of linear equations, eigenvalue problems, and singular value decomposition, building on the legacy of BLAS. LAPACK's significance lies in its efficiency on shared-memory processors, greatly influencing computational methods in various scientific and engineering applications.

1993

The **TOP500** list ranks the world's 500 most powerful supercomputers. It serves as a benchmark for tracking the development of supercomputing technology and its evolution over time, reflecting shifts in computing power, architecture, and geopolitical dominance in high-performance computing. The list, updated biannually, has become a key indicator of technological and scientific progress in the field of computing.

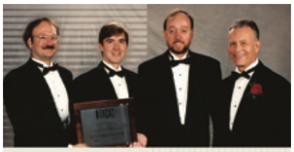

1994
PVM becomes ICL's first R&D 100 Award winner. Often referred to as the "Oscars of Invention," an R&D 100 Award signifies a significant contribution to science and technology and is seen as a mark of excellence in innovation.

1994

Version 1.0 of a standardized and portable message-passing system, the **Message Passing Interface (MPI)**, was released. MPI has become the *de facto* standard for communication among processes in a distributed memory system. The development of MPI is guided by the **MPI Forum**, a group comprised of representatives from academia, industry, and government institutions. ICL was one of the founding members and still participates in MPI Forum meetings.

1995

The **Scalable LAPACK (ScaLAPACK)** library extended the LAPACK library for high-performance distributed-memory architectures, providing routines for solving linear equations, eigenvalue problems, and singular value problems. ScaLAPACK's design reflects the shift in computational science towards parallel processing. Together, LAPACK and ScaLAPACK form foundational technologies for high-performance computing.

1996

NetSolve was a client-server system that enabled users to solve complex scientific problems using remote resources. NetSolve envisioned scientists, engineers, research professionals, and students working with the powerful and flexible toolset provided by their familiar desktop computing environment and yet being able to easily draw on the vast, shared resources of the Grid for exceptional computational resource needs.

1997

Automatically Tuned Linear Algebra Software (ATLAS) is a software library for linear algebra that automatically optimized itself for specific hardware configurations to maximize performance, a pioneering approach that greatly influenced subsequent software development in high-performance computing.

1999

ICL's flagship performance analysis tool, the **Performance Application Programming Interface (PAPI)**, is a standardized, easy-to-use interface that provides a consistent methodology for collecting performance data from a variety of hardware counters on most major processor platforms. PAPI received an R&D 100 Award in 2002 and is still developed and widely used today. The PAPI team released version 7.1.0 in December 2023.

1999
ATLAS and NetSolve each received R&D 100 Awards.

2000
ICL moved from Ayres Hall to the newly-built Claxton building on UTK campus.

2000

High-Performance Linpack (HPL) is a benchmark for distributed-memory computers that solves a (random) dense linear system in double-precision (64-bit) arithmetic. HPL is often one of the first programs to run on large HPC machines, producing a result that can be submitted to the TOP500 list of the world's fastest supercomputers.

2002
PAPI was recognized with an R&D 100 Award.

2002

Fault Tolerant MPI (FT-MPI) extended the standard MPI by adding fault tolerance capabilities. This innovation allows programs to continue functioning despite node failures, enhancing parallel applications' robustness and resilience. FT-MPI was later merged with LA-MPI, LAM/MPI, and PACX-MPI to form the Open MPI project.

2003

The **HPC Challenge (HPCC)** benchmark suite provided a comprehensive set of tests to evaluate the performance of HPC systems. HPCC extended beyond the capabilities of the LINPACK benchmark, offering a broader assessment of memory bandwidth, computational speed, and network communication to better understand the varied performance aspects of HPC systems across a range of complex computational tasks.

2004

In 2004, four institutions merged efforts in the **Open Source Message Passing Interface (Open MPI)**: FT-MPI from UTK/ICL, LA-MPI from LANL, and LAM/MPI from Indiana University, with contributions from PACX-MPI at the University of Stuttgart. Open MPI is an open-source implementation of the Message Passing Interface (MPI) standard that is fundamentally centered around component concepts.

2008

ICL initiated two projects in response to the evolving landscape of HPC hardware architectures. **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. **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)** was a global series of workshops spanning three years that brought together representatives of the international HPC community focused on developing a cohesive software infrastructure for exascale computing. At the end of the series, the *IESP Roadmap* was published, a strategic plan outlining the necessary steps to successfully create such an infrastructure.

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. PaRSEC's innovative approach optimizes dataflow and dependencies in parallel applications, significantly improving performance and scalability in modern HPC environments, particularly for dynamic and irregular computations.

2012

User Level Failure Mitigation (ULFM) is an extension of the MPI standard, designed to enhance fault tolerance in parallel computing environments. ULFM provides mechanisms for detecting and recovering from failures at the user level, allowing applications to continue operation despite node or process failures, thereby improving resilience and reliability in large-scale high-performance computing systems.

2013
Jack Dongarra was awarded with the ACM/IEEE Ken Kennedy Award "for his leadership in designing and promoting standards for mathematical software used to solve numerical problems common to high-performance computing."

2013

Building upon the work of the IESP, the **Big Data and Extreme-scale Computing (BDEC)** workshop series was initiated to tackle the emerging challenges and opportunities at the nexus of large-scale data analysis and exascale computing. A follow-up series, **BDEC2**, later extended the effort through 2020 and culminated in the publication of the "*Pathways to Convergence*" strategic roadmap, a comprehensive guide outlining the future direction for integrating these emerging technological fields.

2013

The **High Performance Conjugate Gradients (HPCG)** benchmark was designed as a complement to the HPL benchmark that drives the TOP500. HPCG focuses on iterative methods and data access patterns more representative of real-world applications, providing a more comprehensive view of a system's computational capabilities, particularly in memory-bound operations. HPCG awards are now presented in conjunction with TOP500 awards biannually.

2016
ICL landed seven awards through the Exascale Computing Project (ECP) and was the lead institution on four projects focused on enabling technologies for exascale computing.

2016

Distributed Tasking for Exascale (DTE) extended the capabilities of PaRSEC for exascale architectures.

Exascale Performance Application Programming Interface (Exa-PAPI) built on PAPI and extended it with performance counter monitoring capabilities for advanced ECP hardware and software technologies.

The **Production-ready, Exascale-enabled Krylov Solvers for Exascale Computing (PEEKS)** effort aimed to advance the capabilities of the ECP software stack by making new scalable algorithms accessible within the **Ginkgo** software ecosystem.

Software for Linear Algebra Targeting Exascale (SLATE) was designed to provide fundamental dense linear algebra capabilities to the ECP software ecosystem, building a new modern version of the venerable ScaLAPACK library.

2016
At SC16, Clint Whaley and Jack Dongarra's 1998 paper on ATLAS was recognized with a Test of Time Award, which "recognizes an outstanding paper that has deeply influenced the HPC discipline."

2017

The **Batched BLAS (BBLAS)** effort was initiated to create an API for efficiently handling multiple independent BLAS operations simultaneously (in batch), and to serve as a working forum for establishing this strategy as the next official BLAS standard. This approach is particularly useful for optimizing performance when dealing with a large number of small-sized matrix or vector operations, which is common in many scientific and engineering applications.

2018

The **ECP Fast Fourier Transform (ECP-FFT)** project was tasked with designing and implementing a fast and robust 2-D and 3-D fast Fourier transform (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 resulting library, known as **heFFTe**, delivers highly efficient FFTs for exascale computing.

2019

Now known as **HPL-MxP**, a new benchmark initiative was introduced in 2019 to highlight the emerging convergence of high-performance computing and artificial intelligence (AI) workloads. Machine learning methods that fuel advances in AI can deliver results that utilize much lower precision compared to the requirements in traditional HPC fields.

2019

The **Ecosystem for Programming and Executing eXtreme Applications (EPEXA)** team is creating a production-quality, open-source software ecosystem to allow rapid composition of advanced scientific applications for modern distributed memory heterogeneous platforms. A key outcome of EPEXA is the **Template Task Graph (TTG)** programming model, which facilitates an intuitive expression of parallelism in algorithms handling irregular and uneven data sets that are poorly served by current models on modern HPC systems.

2020

The **Surrogate Benchmark Initiative** (**Surrogates** for short) aims to provide benchmarks and tools for assessing deep neural network "surrogate" models, which can imitate part or all of a simulation and produce the same outcomes while requiring fewer resources, enabling faster simulations and predictions while maintaining a high level of accuracy. Tools developed under Surrogates will evaluate these models to measure progress and inform the co-design of new HPC systems to support their use.

2021
Jack Dongarra received the 2021 ACM A.M. Turing Award "for pioneering contributions to numerical algorithms and libraries that enabled high performance computational software to keep pace with exponential hardware improvements for over four decades."

2023

As high-performance computing expertise becomes more crucial in scientific research, ICL has become increasingly involved in more interdisciplinary projects in the field of Computational Science.

ICL leads the HPC/AI cross-cutting focus area within the new **Consortium for Nuclear Forensics (CNF).** Sponsored by the National Nuclear Security Administration (NNSA), the CNF comprises 16 Universities and 7 National Laboratories that contribute to important research fields within nuclear forensics.

2023 HIGHLIGHTS

ECP RETROSPECTIVE

ICL members in attendance at the 2023 ECP Annual Meeting in Houston, Texas

THE Exascale Computing Project (ECP) marked a monumental shift in the computational science landscape. The program was initiated in 2016 by the U.S. Department of Energy (DOE) and the National Nuclear Security Administration (NNSA) with the ambitious goal of developing the technologies required to support a computing ecosystem capable of breaking through the exascale barrier, a benchmark of computing performance exceeding a quintillion (10^{18}) calculations per second. This leap represents not just a quantitative enhancement in computing speed, but a qualitative transformation in computational capabilities, enabling unprecedented scientific research and innovation.

The ECP's goals included not only building the hardware required to reach these computational speeds but also developing the software, applications, and tools necessary to run efficiently on exascale systems. This comprehensive strategy was designed to tackle some of the most complex challenges in various domains, including energy sustainability, environmental science, materials engineering, and more. By pushing computational frontiers, ECP aspired to unlock new realms of scientific inquiry, simulating intricate processes and phenomena with unmatched precision and scale.

ICL played a prominent role in the ECP, having contributed multiple projects in the program's Software Technologies and Application Development thrusts. These projects spanned ICL's core research areas of Numerical Linear Algebra, Distributed Computing, and Performance Analysis, demonstrating the foundational importance of ICL's expertise in high-performance computing.

The ECP proved triumphant when the *Frontier* supercomputer at Oak Ridge National Laboratory officially broke the exascale barrier. This achievement marked a new era in computational science, opening up new possibilities for scientific discovery and technological innovation.

Hear a conversation with Jack Dongarra about the ECP on episode 105 of the *Let's Talk Exascale* podcast

tiny.utk.edu/ecp-podcast

CEED

ICL LEAD
Stanimire Tomov

FIND OUT MORE AT
ceed.exascaleproject.org/

"THE CENTER FOR Efficient Exascale Discretizations (CEED) is a focused team effort within the U.S. Department of Energy Exascale Computing Project that aims to develop the next-generation discretization software and algorithms to enable a wide range of finite element applications to run efficiently on future hardware." – *EXASCALEPROJECT.ORG*

The Lawrence Livermore National Laboratory (LLNL)-led **Center for Efficient Exascale Discretizations** (CEED) co-design effort developed 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 NNSA applications to run efficiently on 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 is instrumental in identifying, developing, and optimizing tensor contractions that are essential building blocks for these kinds of DOE/NNSA applications. The ICL team also plays an integral role in co-designing application programming interfaces (APIs) with the LLNL scientists, external partners, and vendors. It delivers a high-performance tensor contractions package through the Matrix Algebra on GPU and Multicore Architectures (MAGMA) library.

CLOVER

ICL LEADS
Hartwig Anzt
Mark Gates
Stanimire Tomov

"SCIENTIFIC APPLICATIONS NEED to apply efficient and scalable implementations of numerical operations, such as matrix-vector products and Fourier transforms, in order to simulate their phenomena of interest. Software libraries are powerful means of sharing verified, optimized numerical algorithms and their implementations. The CLOVER project is delivering scalable, portable numerical algorithms to facilitate efficient simulations." – *EXASCALEPROJECT.ORG*

The ECP **CLOVER** project focused on developing three advanced, vendor-agnostic GPU-accelerated numerical libraries. This project addressed the need for efficient computational tools that are compatible across various GPU platforms, catering to the diverse and evolving landscape of high-performance computing.

Each of the three libraries in CLOVER has a distinct focus: **SLATE** is dedicated to dense linear algebra computations, offering optimized performance for these tasks on GPU systems. **Ginkgo** specializes in sparse linear algebra, providing efficient solutions for computations involving sparse matrices. Lastly, **heFFTe** focuses on multidimensional Fast Fourier Transforms, a critical component in many scientific calculations, ensuring accelerated performance on GPUs.

By spearheading the development of these three libraries, CLOVER positioned itself as a cornerstone in the realm of GPU-accelerated distributed computing. This initiative represents a significant stride in the ongoing journey to harness the full potential of GPU technologies, underscoring the importance of adaptable, cross-platform solutions in high-performance computing.

2023 HIGHLIGHTS

ECP RETROSPECTIVE

DTE

ICL LEAD
George Bosilca

FIND OUT MORE AT
icl.utk.edu/dte

"ONE DIFFICULTY ASSOCIATED with programming exascale systems is expressing the tasks comprising a scientific simulation and then mapping them to the heterogenous computational resources on that system, while achieving high performance. PaRSEC supports the development of domain-specific languages and tools to simplify and improve the productivity of scientists when using a task-based system and provides a low-level runtime." – *EXASCALEPROJECT.ORG*

The **Distributed Tasking for Exascale** (DTE) project has extended the capabilities of ICL's Parallel Runtime and Execution Controller (PaRSEC)—a generic framework for architecture-aware scheduling and management of microtasks on distributed, many-core, heterogeneous architectures. The PaRSEC environment 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. It 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 project has been a vital extension and continuation of this effort. It has ensured that PaRSEC meets the critical needs of ECP application communities regarding scalability, interoperability, and productivity.

Exa-PAPI

ICL LEADS
Heike Jagode
Anthony Danalis

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

"UNDERSTANDING THE PERFORMANCE characteristics of exascale applications is necessary in order to identify and address the barriers to achieving performance goals. This becomes more difficult as the architectures become more complex. The Performance Application Programming Interface (PAPI) provides both library and application developers with generic and portable access to low-level performance counters found across the exascale machine." – *EXASCALEPROJECT.ORG*

The **Exa-PAPI** project has expanded PAPI's capabilities in performance counter monitoring, incorporating support for power management across cutting-edge hardware and software technologies. This includes performance and power monitoring for AMD GPUs through integration with AMD ROCm and ROCm-SMI, Intel Ponte Vecchio GPUs via Intel's oneAPI Level Zero, and NVIDIA GPUs through the Perfworks API. Additionally, PAPI is compatible with interconnects, the latest CPUs, and ARM chips. These enhancements have been implemented while preserving the standard PAPI interface and methodology for utilizing low-level performance counters in CPUs, GPUs, on/off-chip memory, interconnects, and the I/O system, encompassing energy and power management. To strengthen PAPI's sustainability, ECP has facilitated its integration into the Spack package manager and the Extreme-scale Scientific Software Stack (E4S), ensuring software robustness through continuous integration and continuous deployment. In addition to hardware counter-based data, PAPI now supports the registration and monitoring of Software-Defined Events. This exposes the internal behavior of runtime systems and libraries like PaRSEC, SLATE, MAGMA, and Ginkgo to applications utilizing those libraries, broadening the scope of performance events to include software-based information. Additionally, PAPI has been expanded with the Counter Analysis Toolkit, aiding in native performance counter disambiguation through micro-benchmarks. These micro-benchmarks probe various essential aspects of modern chips, contributing to the classification of raw performance events.

In summary, the Exa-PAPI project has enabled PAPI to add comprehensive counter analysis capabilities, advanced performance and power monitoring support for exascale hardware components, and broadened the scope of performance events to encompass not only hardware-related metrics but also software-based information.

OMPI-X

ICL LEADS
George Bosilca
Joseph Schuchart

FIND OUT MORE AT
icl.utk.edu/ompi-x

"THE MESSAGE PASSING Interface (MPI) is a community standard for inter-process communication and is used by the majority of DOE's parallel scientific applications running on pre-exascale systems. The MPI standard can be implemented on all the large systems. The OMPI-X project ensures that the MPI standard and its specific implementation in Open MPI meet the needs of the ECP community in terms of performance, scalability, and capabilities." – EXASCALEPROJECT.ORG

The **Open MPI for Exascale** (OMPI-X) project focused on preparing the 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. Open MPI is a widely used, high-quality, open-source implementation of the MPI standard. However, despite their history and popularity, neither Open MPI nor the MPI standard itself were ready for the changes in hardware and software that has accompanied exascale computing.

To mitigate this concern, OMPI-X addressed 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.

xSDK

ICL LEAD
Piotr Luszczek

FIND OUT MORE AT
xsdk.info/ecp/

"THE LARGE NUMBER of software technologies being delivered to the application developers poses challenges, especially if the application needs to use more than one technology at the same time, such as using a linear solver from the PETSc/TAO mathematics library in conjunction with a time integrator from the SUNDIALS library. The xSDK project is an effort to create a value-added aggregation of mathematics and scientific libraries, to increase the combined usability." – EXASCALEPROJECT.ORG

The **Extreme-Scale Scientific Software Development Kit** (xSDK) was a collaborative effort between Argonne National Laboratory, ICL, Karlsruhe Institute of Technology, Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, Sandia National Laboratories, and the University of California, Berkeley. The project aimed to enable seamless integration and combined use of diverse, independently developed software packages for ECP applications. This includes a wide range of high-quality software libraries, solver packages, and two applications that address the strategic requirements of DOE's Office of Science.

xSDK streamlined processes for system administrators and application developers by standardizing naming conventions, runtime behavior, and installation procedures through its community policies. Each xSDK package provides a Spack installation script that can be invoked independently or through the installation of xSDK's Spack package. In addition, xSDK ships with a set of curated examples showing potential package integrations into application exemplars. ICL's MAGMA, PLASMA, SLATE, Ginkgo, and heFFTe libraries were included in the latest release of xSDK versioned as 1.0.

2023 HIGHLIGHTS

CONFERENCE HIGHLIGHTS



SC23

The International Conference for High Performance Computing, Networking, Storage, and Analysis returned to Denver in 2023, November 12th–17th. This year's conference General Chair was ICL alumnus Dorian Arnold. ICL was once again well represented at the conference. ICL's Associate Director, Deborah Penchoff, served as SC23's Plenary Prodcutions Chair, and ICL members led and participated in workshops, tutorials, and Birds of a Feather sessions, demonstrating their leadership in the field.

The annual ICL Alumni Dinner was held at Henry's Tavern in Denver on November 15th, and gathered around 40 past and present ICL members. It was a vibrant evening fostering connections and re-connections among attendees, embodying the collaborative spirit fundamental to ICL's legacy.

AWARDS

BEST PAPER FINALIST
ICS 2023 International Conference on Supercomputing
Orlando, FL, June 2023

Lindquist, N., P. Luszczek, and J. Dongarra
Using Additive Modifications in LU Factorization Instead of Pivoting

BEST PAPER FINALIST
37th IEEE International Parallel & Distributed Processing Symposium (IPDPS)
St. Petersburg, FL, May 2023

Sid-Lakhdar, W., S. Cayrols, D. Bielich, A. Abdelfattah, P. Luszczek, M. Gates, S. Tomov, H. Johansen, D. Williams-Young, T. Davis, J. Dongarra, and H. Anzt
PAQR: Pivoting Avoiding QR factorization

BEST POSTER
29th International European Conference on Parallel and Distributed Computing (Euro-Par)
Limassol, Cyprus, August 2023

Cojean, T., F. Göbel, and H. Anzt
BDDC Preconditioning in the Project MICROCARD

BEST STUDENT POSTER
IEEE International Conference on Cluster Computing 2023
Sante Fe, NM, October 2023

Mishler, D., J. Ciesko, S. Olivier and G. Bosilca
Performance Insights into Device-initiated RMA using Kokkos Remote Spaces

HOSTED WORKSHOPS & COLLABORATIONS

ICL/ORNL COLLABORATION MEETING

On May 1, ICL hosted a group of researchers from ORNL for a one-day workshop with ICL members. The event was organized by ICL's George Bosilca as an effort to foster collaboration between the two East Tennessee area HPC research groups.

LARGE SCALE SYSTEMS WORKSHOP

ICL hosted the 16th Scheduling for Large Scale Systems Workshop from May 22 through May 24. The workshop featured thematic half-day sessions focused on topics related to scheduling and algorithms for large-scale systems. George Bosilca and Yves Robert organized the event, which attracted 25 researchers from around the globe, including ICL alumni Mathieu Faverge and Amina Guermouche.

AI-ENVIRONMENT KICK-OFF MEETING

In July, ICL hosted the formal Kick-off meeting for the AI for Environmental Remediation project, which aims to spearhead the design of novel AI-driven solutions to reduce environmental contamination involving REEs and actinides through HPC-exascale enabling capabilities. Joining ICL's George Bosilca, Deborah Penchoff, and Stan Tomov were Niki Labbe from the UT Institute of Agriculture, Charles Sims from UTK's Howard H. Baker Jr. School of Public Policy and Public Affairs, Vasileios Anagnostopoulos from University of Central Florida, and Charles Peterson from University of California, Los Angeles.

SPARSE BLAS WORKSHOP

ICL hosted a Sparse BLAS Workshop November 7-9, 2023. The workshop brought together scientists from major research institutions and industry partners, including AMD, ARM, Intel, MathWorks, and NVIDIA, and focused on developing a common understanding and interface for basic operations on sparse linear algebra objects.

2023 HIGHLIGHTS

CONGRATULATIONS


Neil LINDQUIST

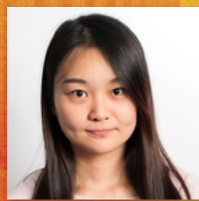
Jiali LI

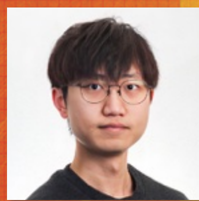
Patrick HON MAN LAU

2023 Graduates

ICL provides an excellent environment for students to gain knowledge and experience, working alongside its expert research staff. In 2023, ICL celebrated the graduation of three students. Neil Lindquist completed his PhD in Computer Science, advised by Jack Dongarra and Piotr Luszczek, focusing on *Reducing Communication in the Solution of Linear Systems*. Jiali Li, under the guidance of Dongarra and George Bosilca, earned her PhD with her dissertaion titled *Optimizing Collective Communication for Scalable Scientific Computing and Deep Learning*. Patrick Hon Man Lau earned his Master's in Computer Science, advised by Stanimire Tomov and Hartwig Anzt.

Jack Dongarra Elected to the National Academy of Sciences

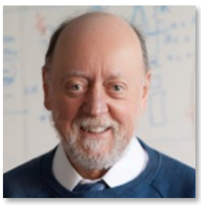

In May 2023, ICL founder Jack Dongarra was elected to the National Academy of Sciences. This prestigious recognition is a testament to his contributions to the field of computer science and high-performance computing. Dongarra has been at the forefront of developing numerical algorithms, software libraries, and performance analysis tools that are widely used in scientific and engineering applications.

TCE Commitment to Inclusive Community Award

In April 2023, ICL Associate Director Deborah Penchoff was awarded with the Tickle College of Engineering Commitment to Inclusive Community Award. The award recognizes an individual within the college who has demonstrated a commitment to diversity and inclusion in consideration of affordability, ability, identity, and access. The Commitment to Inclusive Community Award recipient demonstrates an exceptional understanding of diversity and inclusiveness beyond the call of duty.

SOFTWARE RELEASES

Ginkgo 1.7

Ginkgo is a high-performance linear algebra library for GPU systems, with a focus on solution of sparse linear systems. Version 1.7 was released in November 2023.

AVAILABLE AT
ginkgo-project.github.io/

heFFTe 2.4.0

Version 2.4.0 of the heFFTe software library was released in October 2023. heFFTe is focused on providing highly efficient and scalable FFT computations on distributed GPU architectures, designed specifically to meet the demands of exascale computing.

AVAILABLE AT
github.com/icl-utk-edu/heffte/

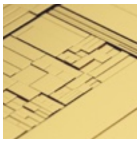

LAPACK 3.12.0

First released in 1992, the LAPACK software library for numerical linear algebra is a foundational HPC technology used in systems worldwide. Version 3.12.0 was released in November 2023, featuring new features, performance improvements, and bug fixes.

AVAILABLE AT
www.netlib.org/lapack/

MAGMA 2.7.2

MAGMA is a collection of next-generation dense linear algebra libraries that is designed for heterogeneous architectures, such as multiple GPUs and multi- or many-core CPUs. MAGMA had two releases in 2023, with the latest version, 2.7.2, in August.

AVAILABLE AT
icl.utk.edu/magma

Open MPI 5.0.1

The Open MPI Project is an open-source Message Passing Interface implementation that is developed and maintained by a consortium of academic, research, and industry partners, including ICL's Distributed Computing team. The latest major version, v5, was released in 2023.

AVAILABLE AT
www.open-mpi.org

PAPI 7.1.0

ICL's flagship performance analysis tool, PAPI, provides a consistent interface and methodology for collecting performance counter information from various hardware and software components. The latest major version, 7.1.0, was released in December 2023.

AVAILABLE AT
icl.utk.edu/papi

PLASMA 23.8.2

PLASMA is a dense linear algebra package at the forefront of multicore computing, designed to deliver the highest possible performance from a system of multiple sockets of multi-core CPUs. Version 23.8.2 was released in August 2023.

AVAILABLE AT
github.com/icl-utk-edu/plasma

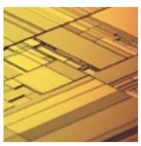

SLATE 2023.11.05

SLATE is a distributed, GPU-accelerated, dense linear algebra library targeting current and upcoming extreme-scale high-performance computing systems.

AVAILABLE AT
github.com/icl-utk-edu/slate/

2023 HIGHLIGHTS

VISITING RESEARCHERS

THE SPIRIT OF collaboration is fundamental to ICL's approach to research. The following researchers visited ICL in 2023, each bringing their unique expertise and perspectives to the group.

Andreas ADELMANN
PAUL SCHERRER INSTITUTE
VILLIGEN, SWITZERLAND

Terry COJEAN
KARLSRUHE INSTITUTE OF TECHNOLOGY
KARLSRUHE, GERMANY

Lisa GAEDKE-MERZHÄUSER
UNIVERSITÀ DELLA SVIZZERA ITALIANA
LUGANO, SWITZERLAND

Thomas GRÜTZMACHER
KARLSRUHE INSTITUTE OF TECHNOLOGY
KARLSRUHE, GERMANY

Sascha HUNOLD
VIENNA UNIVERSITY OF TECHNOLOGY
VIENNA, AUSTRIA

Albert KAHIRA
JÜLICH SUPERCOMPUTING CENTRE
JÜLICH, GERMANY

Sonali MAYANI
PAUL SCHERRER INSTITUTE
VILLIGEN, SWITZERLAND

Veronica MONTANARO
ETH ZÜRICH
ZÜRICH, SWITZERLAND

Brieuc NICOLAS
ENSEIRB-MATMECA
TALENCE, FRANCE

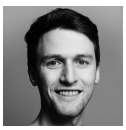

Tobias RIBIZEL
KARLSRUHE INSTITUTE OF TECHNOLOGY
KARLSRUHE, GERMANY

Yves ROBERT
ENS LYON
LYON, FRANCE

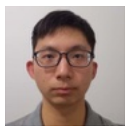

Yu-Hsiang MIKE TSAI
KARLSRUHE INSTITUTE OF TECHNOLOGY
KARLSRUHE, GERMANY

Robert UNDERWOOD
ARGONNE NATIONAL LABORATORY

You WU
ETH ZÜRICH
ZÜRICH, SWITZERLAND

ICL SEMINAR SERIES

FRIDAY LUNCH TALKS are a long-standing tradition at ICL that offer further opportunities to foster collaborations and bring new ideas to the group. ICL's 2023 Seminar Series featured a distinguished group of invited speakers from UTK and ORNL to other groups around the globe.


JAN 27
David ICOVE
UL Professor of Practice
UTK EECS


FEB 10
Slaven PELES
Senior Computational Scientist
ORNL


FEB 17
Jakob LUETTGAU
Research Assistant Professor
UTK Global Computing Lab


FEB 24
Andrea DELGADO
Distinguished Staff Fellow
ORNL


MAR 3
Corey AHL
PhD Student
UTK Nuclear Engineering


MAR 10
Rafael FERREIRA DA SILVA
Senior Research Scientist
ORNL

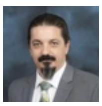
MAR 17
Vincent PAQUIT
Section Head, Secure & Digital Manufacturing
ORNL

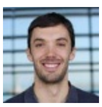
MAR 24
Mathieu FAVERGE
Assistant Professor
Bordeaux Institute of Technology


MAR 31
Sascha HUNOLD
Associate Professor
Vienna University of Technology


APR 14
Robert HARRISON
Professor / Director of Institute for Advanced Computational Science
Stony Brook University


APR 21
Reza ABEDI
Associate Professor
UT Space Institute


APR 28
Suzanne PARETE-KOON
HPC Engineer
ORNL


MAY 5
Nicholas PETERS
Section Head & Distinguished R&D Staff
ORNL


MAY 12
Sunita CHANDRASEKARAN
Associate Professor
University of Delaware


AUG 25
Brieuc NICOLAS
Computer Science Engineering Student
ENSEIRB-MATMECA


SEP 1
Devina SANJAYA
Assistant Professor
UTK MABE Department


SEP 8
Rinkle JUNEJA
R&D Associate
ORNL


SEP 22
Lisa GAEDKE-MERZHÄUSER
PhD Student
Università della Svizzera italiana


SEP 29
Steffen SCHOTTHOEFER
Householder Fellow
ORNL


OCT 6
Fabrizio BISETTI
Associate Professor
University of Texas at Austin


OCT 13
Thomas GILLIS
Postdoctoral Appointee
Argonne National Laboratory


OCT 20
Cory HAUCK
Applied Mathematician
ORNL


OCT 27
Yaohung (MIKE) TSAI
Research Scientist
Meta

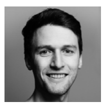
NOV 10
Tobias RIBIZEL
Research Software Engineer
Karlsruhe Institute of Technology


DEC 8
Yves ROBERT
Professor
ENS Lyon

RESEARCH

AI-Environment

ICL LEAD
Deborah Penchoff

FIND OUT MORE AT
icl.utk.edu/ai-env

The **AI For Environmental Remediation** Project is funded by the AI Tennessee Initiative, and it focuses on applying AI-driven tools to reduce environmental toxicity stemming from nuclear and radiochemical processes. The project contributes to the University's mission of leading research and development efforts to increase sustainability. The goal of the project is to design AI-driven solutions (1) to reduce environmental contamination from anthropogenic sources and (2) to develop environmentally friendly separations of Rare Earth Elements (REEs).

The team includes members from The University of Tennessee Knoxville, the UT Institute for Agriculture, the University of Central Florida, and the University of California, Los Angeles.

BALLISTIC

ICL LEAD
Piotr Luszczek

FIND OUT MORE AT
icl.utk.edu/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 pyramid 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.

Batched BLAS

ICL LEAD
Stanimire Tomov

FIND OUT MORE AT
icl.utk.edu/bblas

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.

Batched Sparse Linear Algebra

ICL LEAD
Piotr Luszczek

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 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.

BIOCHAR

ICL LEAD
Deborah Penchoff

The **Biochar Enhanced Ecosystem Services for Energy Crop Systems in the Southeast** project is sponsored by the Department of Energy Office of Energy Efficiency and Renewable Energy. The project focuses on decarbonization strategies in the supply chain to decrease the carbon intensity of sustainable aviation fuel (SAF) through the application of biochar, a carbon-rich soil enhancer made from organic waste. Biochar contributes to carbon sequestration and fertilization and allows for pathways in a circular system. ICL's contribution to the project focuses on applying HPC and AI-driven strategies to predict biochar properties to optimize its utilization in SAF production systems in the southeast.

Team members are from The University of Tennessee, Knoxville; The University of Tennessee Institute of Agriculture; Tennessee State University; Mississippi State University; Oak Ridge National Laboratory; and industry partners Kolmar-Americas and Genera Inc.

CETOP

ICL LEADS
Hartwig Anzt
Natalie Beams

FIND OUT MORE AT
icl.utk.edu/cetop

ICL will participate in the **Center for Edge of Tokamak Optimization** (CETOP), a new 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 will focus on improving GPU-capable preconditioning of solvers in the plasma physics codes and will make use of the Ginkgo library.

Consortium for Nuclear Forensics

ICL LEAD
Deborah Penchoff

FIND OUT MORE AT
icl.utk.edu/cnf

The **Consortium for Nuclear Forensics** (CNF) is the latest consortium sponsored by the US Department of Energy's National Nuclear Security Administration (NNSA). It focuses on nuclear forensics research and education involving rapid turnaround forensics, advanced analytical methods, ultrasensitive measurements, signature discovery, and prompt effects. These efforts are achieved through combined expertise in radiochemistry, geochemistry, analytical chemistry, nuclear material science, shock physics, quantum-enabling sensing, HPC, data science, and AI.

The CNF team includes members from the University of Florida; the University of Tennessee; the University of California, Berkeley; the University of Central Florida; the City University of New York; Clemson University; George Washington University; Iowa State University; the University of Michigan; the University of Nevada, Las Vegas; North Carolina State University; the University of Notre Dame; Oregon State University; Penn State; Lawrence Berkeley National Laboratory; Lawrence Livermore National Laboratory; Los Alamos National Laboratory; Oak Ridge National Laboratory; Pacific Northwest National Laboratory; Sandia National Laboratory; and Savannah River National Laboratory.

CORE-CM

ICL LEAD
Deborah Penchoff

The **Carbon Ore, Rare Earth and Critical Minerals** (CORE-CM) initiative for the U.S. Basins is funded by the Department of Energy Fossil Energy & Carbon Management Office of Resource Sustainability. Our project within this program focuses on applying HPC and AI to optimize processes to advance the manufacturing of valuable coal-derived products in southern Appalachia.

REEs are essential in applied technologies that include communications, computing, medical capabilities, green energy, and defense. Our project in CORE-CM seeks to develop HPC/AI-driven optimized strategies to extract REEs from coal ash ponds. This benefits society by providing potential alternatives to increase internal production of REEs while providing uses for coal ash. The team includes participants from The University of Tennessee, the Institute for Advanced Composites Manufacturing Innovation (IACMI), Oak Ridge National Laboratory, the University of Alabama, Southern Company, Roane State Community College, the Tennessee Geological Survey, and the Geological Survey of Alabama.

CORES

ICL LEAD
Stanimire Tomov

FIND OUT MORE AT
icl.utk.edu/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 enable us to solve, in real-time, convex optimization problems with significantly higher performance—and orders of magnitude greater size—compared to current state-of-the-art solvers.

DPLASMA

ICL LEADS
George Bosilca
Aurelien Bouteiller

FIND OUT MORE AT
icl.utk.edu/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.

RESEARCH

EPEXA

ICL LEADS
George Bosilca
Joseph Schuchart
Thomas Herault

FIND OUT MORE AT
icl.utk.edu/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.

Evolve

ICL LEAD
George Bosilca

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 improved efficiency 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 and has been fully integrated into Open MPI.

F-block Elements Separations

ICL LEAD
Deborah Penchoff

The project "**Membrane separations of f-block elements utilizing rhamnolipid-based chelants. A joint experimental and computational investigation**" is funded by the Department of Energy Small Business Technology Transfer (STTR), and it focuses on using machine learning to design better membranes to enable REE separations.

This effort is within technology areas relevant to the DOE Office of Science seeking to optimize selectivity and rate of extraction or rejection using AI/ML methodology membrane synthesis design, experimental design, and data analytics. The R&D is intended to lead to hybrid models that incorporate fundamental mechanisms and lead to feasible material structures that would be readily synthesizable for large-scale processes and complex systems.

The team includes The University of Tennessee and GlycoSurf, LLC.

Ginkgo

ICL LEAD
Hartwig Anzt

FIND OUT MORE AT
ginkgo-project.github.io/

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.

Greening the Southeast

ICL LEAD
Deborah Penchoff

The **Advancing Carbon-Neutral Crop Technologies to Develop Sustainable Consumer Goods (AL, GA, NC, TN)** ("Greening the Southeast" for short) is an NSF Engines Development Award in the Advanced Agriculture Topic and targets the BioTech, Advanced Materials, and Quantum Tech focus areas.

The NSF Engine Development Awards will help regional partners collaborate to advance technologies, address societal challenges, and create economic opportunities.

The team includes The University of Tennessee, HudsonAlpha Institute for Biotechnology, and other partners.

heFFTe

ICL LEAD
Stanimire Tomov

FIND OUT MORE AT
icl.utk.edu/fft

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 a **Highly Efficient FFTs for Exascale** (heFFTe) library that provided 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 v2.4 release achieves very good scalability on exascale systems and performance close to 90% of the roofline peak.

HPCG

ICL LEAD
Piotr Luszczek

FIND OUT MORE AT
hpcg-benchmark.org/

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 SC23 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.

HPFFT

ICL LEAD
Stanimire Tomov

The new **High Performance FFT** (HPFFT) benchmarking effort aims to comprehensively measure the performance of a variety of parallel FFT implementations for the ECP and its rich application ecosystem. Many scientific applications of strategic importance to DOE need multidimensional Discrete Fourier transforms, implemented by ICL's own heFFTe project. Over the course of the CLOVER project and, more specifically, its heFFTe component, the need for comparing FFT implementations became apparent in order to understand how tradeoffs made by developers affect floating-point performance and mitigate communication overheads.

In its current incarnation, the HPFFT project released a technical report, ICL-UT-21-03, to provide an interim update on comprehensive testing and measuring of various aspects of the FFT implementations for three-dimensional transforms on high-performance systems that may include hardware accelerators. The results from the report reveal significant differences in functionality, performance, and supported hardware and give the ECP application community a much better appreciation of the available software choices.

HPL

ICL LEADS
Piotr Luszczek
Jack Dongarra

FIND OUT MORE AT
icl.utk.edu/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 Vector Signal and Image Processing Library (VSIPL). HPL is often the first program to run on large HPC machines. Carefully optimized versions of HPL are available from major HPC hardware vendors.

The primary 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 is now hosted on GitHub and features more detailed reporting of the solution's scaled residual and the achieved performance number. Another addition is a software configuration tool based on GNU Autotools and the removal of deprecated MPI functions. The LINPACK app for iOS reached over 8 Gflop/s on the iPhone X. For the November 2021 TOP500 list, an optimized version of the HPL code achieved over 440 Pflop/s on the *Fugaku* supercomputer at RIKEN, Japan, and in June 2022, achieved 1.1 Eflop/s on *Frontier* at Oak Ridge National Laboratory and was further improved by almost 10% for June 2023.

HPL-MxP

ICL LEADS
Piotr Luszczek
Jack Dongarra

FIND OUT MORE AT
hpl-mxp.org/

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.

RESEARCH

LAPACK/ ScaLAPACK

ICL LEADS
Julie Langou
Piotr Luszczek
Jack Dongarra

FIND OUT MORE AT
netlib.org/lapack
netlib.org/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 maintenance of these two packages. LAPACK is single node and relies on a multi-core BLAS library for parallelism. ScaLAPACK is parallel and distributed and depends on the BLAS, LAPACK, MPI, and BLACS libraries.

The latest major release of LAPACK, version 3.12.0, was released in November 2023. This version added a number of improvements along with bug fixes. Highlights include the dynamic mode decomposition (DMD), truncated QR with column pivoting, reciprocal vector scaling, improvements in the 64-bit API, and a revised online documentation structure.

MAGMA

ICL LEAD
Stanimire Tomov

FIND OUT MORE AT
icl.utk.edu/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 multi-GPUs/coprocessors 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. MAGMA 2.7.2 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.

MAGMA-Edge

ICL LEAD
Stanimire Tomov

MAGMA-Edge is a fast deep learning (DL) library which aims to become the *de facto* standard DL framework for running AI applications on commodity edge devices. Edge devices such as the Raspberry Pi, NVIDIA Jetson Nano, or embedded processors from Intel and AMD are widely used in many common applications or appliances. These devices, when equipped with various sensors, provide efficient and economical Internet of Things (IoT) solutions for many common commercial needs and operations, most obviously on the food delivery bot seen around the UTK campus. Unlike traditional IoT devices that exchange data and instructions with a cloud server, the latest generation of edge devices is capable of performing data analytics and providing real-time AI solutions for a spectrum of designated tasks. MAGMA-Edge has been built on a set of fast linear algebra kernels for tensor operations extended from the MAGMA and MagmaDNN libraries.

MATEDOR

ICL LEAD
Stanimire Tomov

The **MAtrix, TEnsor, and Deep-learning Optimized Routines** (MATEDOR) project is performing the research required to define a standard interface for batched operations and provide a performance-portable software library demonstrating 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 defines modular, optimizable, and language-agnostic interfaces that can work seamlessly with a compiler. This modularity provides application, compiler, and runtime system developers with the option to use a single call to a routine from the new batch operation standard and allows the entire linear algebra community to collectively attack a wide range of small matrix or tensor problems.

Mixed-Precision Numerical Computing

ICL LEAD
Piotr Luszczek

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 them is this effort'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 floating-point representations such as limited precision, 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 in parallel to the mixed-precision aspects.

oneAPI

ICL LEAD
Hartwig Anzt

oneAPI is an industry initiative to create a single, unified, cross-architecture programming model for CPUs and accelerator architectures. Based on industry standards and its open development approach, the initiative will help streamline software development for high-performance computers, increase performance, and provide specifications for efficient and diverse architecture programming.

In the **oneAPI Center of Excellence**, Intel teams up with ICL to spearhead the development of numerical linear algebra functionality in the oneAPI ecosystem by deploying SYCL kernels in the Ginkgo linear algebra library. This allows running Ginkgo and applications that use Ginkgo's functionality on any hardware supporting the oneAPI industry standard, including Intel high-performance GPUs that are powering the *Aurora* supercomputer at Argonne National Laboratory.

Open MPI

ICL LEADS
George Bosilca
Joseph Schuchart

FIND OUT MORE AT
open-mpi.org

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.

PAPI

ICL LEADS
Heike Jagode
Anthony Danalis

FIND OUT MORE AT
icl.utk.edu/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 2023, PAPI released versions 7.0.1 and 7.1.0. These provide enhanced performance counter monitoring capabilities, extending support now to Intel Sapphire Rapids, AMD Zen4 CPUs, and ARM Neoverse V1 and V2 architectures. The ROCm component has been significantly refactored, enabling efficient monitoring of multiple GPUs in sampling mode and laying the foundation for future ROCm tools support. The ROCm SMI component was also expanded to include GPU-to-GPU XGMI event support, offering a wider range of monitoring capabilities. Furthermore, the CUDA component has also undergone significant refactoring and now provides multi-thread and multi-GPU support.

PaRSEC

ICL LEADS
George Bosilca
Thomas Herault

FIND OUT MORE AT
icl.utk.edu/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 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.

PESA

ICL LEAD
Stanimire Tomov

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 with MVAPICH and TAU using Emerging Communication Primitives** (PESA) project undertakes this challenge and pursues new performance engineering avenues – by exploiting a co-design approach using the MPI_T 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.

RESEARCH

PLASMA

ICL LEAD
Piotr Luszczek

FIND OUT MORE AT
icl.utk.edu/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 last decade, 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 23.8.2 from August 2023, added a tighter integration with vendor libraries, convenience scripts for easier configuration, and other general bug fixes and integration with external projects such as the ECP-sponsored Spack package manager.

RHE-HPC

ICL LEAD
Deborah Penchoff

The **Numerically-Exact Relativistic Many-Body Electronic Structure of Heavy Elements** (RHE-HPC for short) project is a collaborative effort focused on the development of a completely new and novel computational tool that enables, for the first time, fully predictive calculations on molecules containing f-block elements. This effort includes high-performance distributed and heterogeneous computing to assist with tuning the implementation of novel methods for execution on large-scale HPC platforms, including exascale machines.

RHE-HPC enables advances in global needs involving f-block elements, with a special focus on actinium-225, a highly promising radiotherapeutic to treat cancer. The team includes members from The University of Tennessee, Stony Brook University, and Virginia Tech.

ScaRT

ICL LEAD
George Bosilca

The **Scalable Run Time for Highly Parallel, Heterogeneous Systems** (ScaRT) project aimed 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 upon which they run, and by better matching the functionality of the underlying communication library to the capabilities of modern communication adapters.

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

SLATE

ICL LEAD
Mark Gates

FIND OUT MORE AT
icl.utk.edu/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 is the next iteration of our linear algebra libraries, using modern C++ technologies and targeting modern GPU-accelerated computer architectures to integrate into the ECP ecosystem.

For context, ScaLAPACK was first released in 1995, nearly 30 years ago. In the past three decades, HPC has witnessed tectonic shifts in hardware and software technology, as well as many algorithmic innovations in scientific computing. SLATE was conceived to be a replacement for ScaLAPACK, with new algorithms and boasting superior performance and scalability in HPC's modern, heterogeneous, distributed-memory environments.

SPADE

ICL LEADS
Heike Jagode
Anthony Danalis

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 15th, 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.

STEP

ICL LEAD
Heike Jagode

FIND OUT MORE AT
ascr-step.org/

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.

Surrogates

ICL LEAD
Piotr Luszczek

FIND OUT MORE AT
icl.utk.edu/surrogates

The DOE-funded **Surrogate Benchmark Initiative** (Surrogates) is a collaborative effort involving Indiana University, UTK/ICL, the University of Virginia, and Rutgers University to provide new benchmarks and tools for assessing "surrogate" models based on deep neural networks. When trained on data produced by ensemble runs of a given HPC simulation, a surrogate model can imitate—with high fidelity—part or all of that simulation and produce the same outcomes for a given set of inputs while requiring far less time and energy.

At present, however, there are no accepted benchmarks to evaluate these surrogate models, and there is no easy way to measure progress or inform the codesign of new HPC systems to support their use. Surrogates aims to address this fundamental problem by creating a community repository and a Findable, Accessible, Interoperable, and Reusable (FAIR) data ecosystem for HPC application surrogate benchmarks, including data, code, and all relevant collateral artifacts that the science and engineering community needs to use and reuse these data sets and surrogates.

Sustainable Energy Ecosystems

ICL LEAD
Deborah Penchoff

The **UTEC-UTK Collaboration for Sustainable Energy Ecosystems in Southern Cone** project focuses on accelerating workforce development, knowledge sharing, and climate action through international collaboration in the field of sustainable energy. The project is funded by the CAF Development Bank of Latin America, the U.S. Department of State (DOS), the Bureau of Western Hemisphere Affairs, and the U.S. Embassy in Chile, and involves higher education institutions in the United States and Argentina, Chile, Paraguay, and Uruguay. This effort aligns with the DOS goal to create a greener, more inclusive, and prosperous Hemisphere through enhanced regional education cooperation.

Participating institutions include The University of Tennessee, the UT Institute for Agriculture, and the Technological University of Uruguay.

TOP500

ICL LEAD
Jack Dongarra

FIND OUT MORE AT
top500.org/

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 2023, the 62nd TOP500 list was unveiled during SC23, which was held in Denver, Colorado. Yet again, the U.S. took the crown with *Frontier*, with a mix of AMD CPUs and GPUs featuring a Cray Slingshot 11 interconnect and built by HP Enterprise. The system's mix of nearly 9 million cores propelled Frontier into the exascale era. It achieved nearly 1.2 Eflop/s in HPL, making it the fastest supercomputer in the world by more than a twofold factor over Japan's Fugaku machine, a CPU-only ARM system. Frontier has now held the first spot for four consecutive lists. The second place finally saw Argonne National Lab's *Aurora* system, but running only at about half of its capacity.

ULFM

ICL LEADS
George Bosilca
Aurelien Bouteiller

FIND OUT MORE AT
fault-tolerance.org/

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.

2023 CONFERENCES ATTENDED


JAN 16-21 HOUSTON, TX
ECP Annual Meeting


JAN 18-19 WASHINGTON, DC
Energy, Mobility, and Environment


FEB 1-8 VIRTUAL
MPI Forum


FEB 2-3 NEW YORK, NY
EPEXA Meeting


FEB 15-18 BATON ROUGE, LA
WAMTA 2023
Workshop on Asynchronous Many-Task Systems and Applications

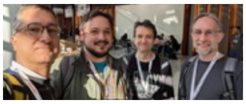
FEB 26 - MAR 3 AMSTERDAM
SIAM CSE23
Conference on Computational Science and Engineering


MAR 13-16 VIRTUAL
MPI Forum


MAR 21-23 BORDEAUX, FRANCE
JLESC 15


MAR 26-30 INDIANAPOLIS, IN
American Chemical Society Spring 2023


MAR 29-31 PARIS, FRANCE
Scheduling Variable Capacity Resources for Sustainability

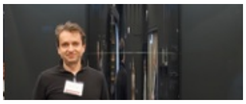
APR 10 YORKTOWN HEIGHTS, NY
IBM Quantum Lab visit


APR 13-15 KNOXVILLE, TN
American Nuclear Society 2023 Student Conference


APR 16-20 COPPER MOUNTAIN, CO
Copper Mountain Conference on Multigrid Methods


APR 20-21 ROANOKE, VA
EPEXA Project Meeting


APR 24-28 GLENEDEN BEACH, OR
Salishan HPC Conference


APR 26-27 GAINESVILLE, FL
DOE NNSA Consortium for Nuclear Forensics Kick-off Meeting


MAY 2-5 VIRTUAL
MPI Forum

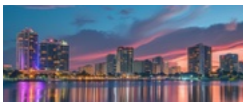
MAY 15-19 ST. PETERSBURG, FL
IPDPS 2023
IEEE International Parallel and Distributed Processing Symposium Workshops

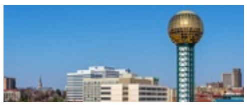
MAY 21-24 KNOXVILLE, TN
USPLF 2023
U.S. Precision Livestock Farming Conference


MAY 21-25 HAMBURG, GERMANY
ISC High Performance 2023

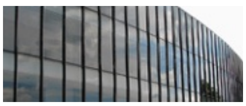
JUN 6-9 YORKTOWN HEIGHTS, NY
STEP Town Hall Meeting

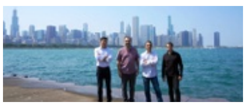
JUN 6-9 LEMONT, IL
PESO Community Workshop


JUN 17 ORLANDO, FL
CSC 2023
Correctness and Resiliency at Scale in Scientific Computing

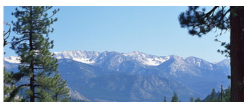
JUN 18-22 LAKE TAHOE, CA
Scalable Tools Workshop


JUN 21-23 ORLANDO, FL
ICS 2023
ACM International Conference on Supercomputing

JUN 26-28 DAVOS, SWITZERLAND

PASC 2023
The Platform for Advanced Scientific Computing

JUL 3-7 BERLIN, GERMANY

Berlin Summit for EVE
Earth Virtualization Engines

JUL 10-12 BUCHAREST, ROMANIA

ISPDC 2023
22nd IEEE International Symposium on Parallel and Distributed Computing

JUL 11-12 AMES, IA

STEP Midwest Town Hall

AUG 1-3 LIVERMORE, CA

ECP CEED Annual Meeting

AUG 9-10 WASHINGTON, DC

Systems and Applications Challenges for the Emerging Bazaar of Accelerators Workshop

AUG 13-17 SAN FRANCISCO, CA

American Chemical Society Fall 2023 Meeting

AUG 20-25 TOKYO, JAPAN

ICIAM 2023
10th International Congress on Industrial and Applied Mathematics

AUG 21-23 COLUMBUS, OH

11th Annual MVAPICH User Group (MUG) Conference

AUG 21-24 MONTEREY, CA

Monterey Data Conference

AUG 29-31 KNOXVILLE, TN

Smoky Mountains Computational Sciences & Engineering Conference

SEP 11-15 BRISTOL, UK

EuroMPI 2023 / MPI Forum

SEP 24-29 HEIDELBERG, GERMANY

Heidelberg Laureate Forum

SEP 25-29 VIRTUAL

IEEE High Performance Extreme Computing Virtual Conference (HPEC)

OCT 9-11 BUENOS AIRES, ARGENTINA

Enhancing Transdisciplinary Networks through Strategic Partnerships between the University of Tennessee and Argentina

OCT 16-17 DOHA, QATAR

ICSEC23
International Computational Science and Engineering Conference

OCT 18-19 CHICAGO, IL

Workshop on Sparse Tensor Computations

OCT 26-27 NEW YORK, NY

EPEXA Project Meeting

OCT 29 - NOV 3 WEST PALM BEACH, FL

Radiobioassay and Radiochemical Measurements Conference

OCT 31 - NOV 1 SANTA FE, NM

IEEE Cluster 2023

OCT 31 - NOV 1 ATLANTA, GA

ORNL Core Universities AI Workshop 2023

NOV 11-17 DENVER, CO

SC23

NOV 20-28 KARLSRUHE, GERMANY

WarmWorld Workshop Karlsruhe

DEC 4-7 VIRTUAL

MPI Forum

DEC 12-15 ALBUQUERQUE, NM

Kokkos User Group Meeting

2023 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 2023 papers, including journal articles, book chapters, and conference proceedings. Many of these are available for download from our website.

Abdelfattah, A., S. Tomov, P. Luszczek, H. Anzt, and J. Dongarra
GPU-based LU Factorization and Solve on Batches of Matrices with Band Structure
SC-W 2023: Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis, Denver, CO, ACM, November 2023

Barry, D., H. Jagode, A. Danalis, and J. Dongarra
Memory Traffic and Complete Application Profiling with PAPI Multi-Component Measurements
2023 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), St. Petersburg, Florida, IEEE, August 2023

Deshmukh, S., R. Yokota, G. Bosilca, and Q. Ma
O(N) distributed direct factorization of structured dense matrices using runtime systems
52nd International Conference on Parallel Processing (ICPP 2023), Salt Lake City, Utah, ACM, August 2023

Aggarwal, I., P. Nayak, A. Kashi, and H. Anzt
Preconditioners for Batched Iterative Linear Solvers on GPUs
SMC 2022: Accelerating Science and Engineering Discoveries Through Integrated Research Infrastructure for Experiment, Big Data, Modeling and Simulation, January 2023

Benoit, A., T. Herault, L. Perotin, Y. Robert, and F. Vivien
Revisiting I/O bandwidth-sharing strategies for HPC applications
INRIA Research Report, no. RR-9502: INRIA, March 2023

Grützmacher, T., H. Anzt, and E. S. Quintana-Ortí
Using Ginkgo's memory accessor for improving the accuracy of memory-bound low precision BLAS
Software: Practice and Experience, vol. 532, issue 1, January 2023

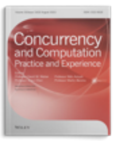

Aliaga, J. I., H. Anzt, E. S. Quintana-Orti, and A. E. Thomas
Sparse matrix-vector and matrix-multivector products for the truncated SVD on graphics processors
Concurrency and Computation: Practice and Experience, August 2023

Cao, Q., S. Abdulah, H. Ltaief, M. G. Genton, D. Keyes, and G. Bosilca
Reducing Data Motion and Energy Consumption of Geospatial Modeling Applications Using Automated Precision Conversion
2023 IEEE International Conference on Cluster Computing (CLUSTER), Santa Fe, NM, USA, IEEE, November 2023

Hoefler, T., et al.
Earth Virtualization Engines - A Technical Perspective
Computing in Science & Engineering, vol. 25, no. 03, May-June 2023

Barbut, Q., A. Benoit, T. Herault, Y. Robert, and F. Vivien
When to checkpoint at the end of a fixed-length reservation?
Fault Tolerance for HPC at eXtreme Scales (FTXS) Workshop, Denver, United States, August 2023

Deshmukh, S., R. Yokota, and G. Bosilca
Cache Optimization and Performance Modeling of Batched, Small, and Rectangular Matrix Multiplication on Intel, AMD, and Fujitsu Processors
ACM Transactions on Mathematical Software, vol. 49, issue 3, September 2023

Li, J., G. Bosilca, A. Bouteiller, and B. Nicolae
Elastic deep learning through resilient collective operations
SC-W 2023: Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis, Denver, CO, ACM, November 2023

BEST PAPER FINALIST

Lindquist, N., P. Luszczek, and J. Dongarra
Using Additive Modifications in LU Factorization Instead of Pivoting
37th ACM International Conference on Supercomputing (ICS23), Orlando, FL, ACM, June 2023

Ribizel, T., and H. Anzt
Parallel Symbolic Cholesky Factorization
SC-W 2023: Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis, Denver, CO, ACM, November 2023

Sukkari, D., M. Gates, M. Al Farhan, H. Anzt, and J. Dongarra
Task-Based Polar Decomposition Using SLATE on Massively Parallel Systems with Hardware Accelerators
SC-W '23: Proceedings of the SC '23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis, Denver, CO, ACM, November 2023

Luszczek, P., W. M. Sid-Lakhdar, and J. Dongarra
Combining multitask and transfer learning with deep Gaussian processes for autotuning-based performance engineering
The International Journal of High Performance Computing Applications, March 2023

BEST PAPER FINALIST

Sid-Lakhdar, W., S. Cayrols, D. Bielich, A. Abdelfattah, P. Luszczek, M. Gates, S. Tomov, H. Johansen, D. Williams-Young, T. Davis, et al.
PAQR: Pivoting Avoiding QR factorization
2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS), St. Petersburg, FL, USA, IEEE, 2023

Thiyagalingam, J., G. von Laszewski, J. Yin, M. Emani, J. Papay, G. Barrett, P. Luszczek, A. Tsaris, C. Kirkpatrick, F. Wang, et al.
AI Benchmarking for Science: Efforts from the MLCommons Science Working Group
Lecture Notes in Computer Science, vol. 13387: Springer International Publishing, January 2023

Mishler, D., J. Ciesko, S. Olivier, and G. Bosilca
Performance Insights into Device-initiated RMA Using Kokkos Remote Spaces
2023 IEEE International Conference on Cluster Computing Workshops (CLUSTER Workshops), Santa Fe, NM, USA, IEEE, November 2023

Schuchart, J., S. Hunold, and G. Bosilca
Synchronizing MPI Processes in Space and Time
EUROMPI '23: 30th European MPI Users' Group Meeting, Bristol, United Kingdom, ACM, September 2023

Tsai, Y-H. Mike, N. Beams, and H. Anzt
Mixed Precision Algebraic Multigrid on GPUs
Parallel Processing and Applied Mathematics (PPAM 2022), vol. 13826, Cham, Springer International Publishing, April 2023

Mor, O., G. Bosilca, and M. Snir
Improving the Scaling of an Asynchronous Many-Task Runtime with a Lightweight Communication Engine
ICPP '23: Proceedings of the 52nd International Conference on Parallel Processing, Salt Lake City, Utah, ACM, September 2023

Schuchart, J., and G. Bosilca
MPI Continuations And How To Invoke Them
Sustained Simulation Performance 2021, Cham, Springer International Publishing, February 2023

Tsai, Y-H. Mike, N. Beams, and H. Anzt
Three-precision algebraic multigrid on GPUs
Future Generation Computer Systems, July 2023

Reed, D., D. Gannon, and J. Dongarra
HPC Forecast: Cloudy and Uncertain
Communications of the ACM, vol. 66, issue 2, January 2023

S. Slattery, K. A. Surjuse, C. Peterson, D. Penchoff, and E. Valeev
Economical Quasi-Newton Unitary Optimization of Electronic Orbitals
Physical Chemistry Chemical Physics, December 2023

Valeev, E., R. J. Harrison, A. Holmes, C. Peterson, and D. Penchoff
Direct determination of optimal real-space orbitals for correlated electronic structure of molecules
Journal of Chemical Theory and Computation, October 2023

ICL STAFF AND STUDENTS

DECEMBER 2023

Ahmad ABDELFATTAH
RESEARCH SCIENTIST II

Hartwig ANZT
DIRECTOR, 2023
MATHWORKS ASSOCIATE PROFESSOR, EECS

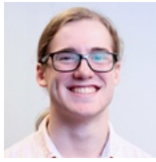

Daniel BARRY
GRADUATE RESEARCH ASSISTANT

Natalie BEAMS
RESEARCH SCIENTIST I

George BOSILCA
RESEARCH ASSOCIATE PROFESSOR

Aurelien BOUTEILLER
RESEARCH ASSISTANT PROFESSOR

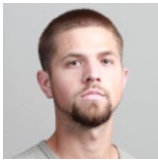

Treece BURGESS
RESEARCH ASSOCIATE II

Pierluigi CAMBIE-FABRIS
BIWEEKLY STUDENT ASSISTANT

Sebastien CAYROLS
RESEARCH SCIENTIST I

Giuseppe CONGIU
RESEARCH SCIENTIST I

Anthony DANALIS
RESEARCH ASSISTANT PROFESSOR

Tatiana DEHOFF
UT STUDENT ASSISTANT

Jack DONGARRA
RESEARCH PROFESSOR EMERITUS

Mark GATES
INTERIM DIRECTOR, 2024
RESEARCH ASSISTANT PROFESSOR

Thomas GRÜTZMACHER
RESEARCH ASSOCIATE II

Fritz GÖBEL
RESEARCH ASSOCIATE II

Julian HALLOY
BIWEEKLY STUDENT ASSISTANT

Thomas HERAULT
RESEARCH ASSISTANT PROFESSOR

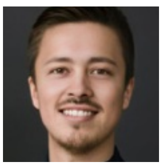

Mark HIPSHIRE
GRADUATE RESEARCH ASSISTANT

Heike JAGODE
RESEARCH ASSOCIATE PROFESSOR

Julie LANGOU
RESEARCH LEADER

Patrick LAU
GRADUATE RESEARCH ASSISTANT

Tracy LEE
BUSINESS MANAGER

Yicheng LI
GRADUATE RESEARCH ASSISTANT

Neil LINDQUIST
RESEARCH ASSOCIATE II

Piotr LUSZCZEK
RESEARCH ASSOCIATE PROFESSOR

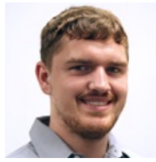

Noah MARKS
GRADUATE RESEARCH ASSISTANT

Maksim MELNICHENKO
GRADUATE RESEARCH ASSISTANT

Daniel MISHLER
GRADUATE RESEARCH ASSISTANT

Anustuv PAL
RESEARCH SCIENTIST I

Deborah PENCHOFF
ASSOCIATE DIRECTOR

Gerald RAGGHIANTI
TECHNICAL SERVICES LEADER

David ROGERS
COMMUNICATIONS SERVICES LEADER

Joseph SCHUCHART
RESEARCH ASSISTANT PROFESSOR

Wissam SID LAKHDAR
RESEARCH SCIENTIST I

Dalal SUKKARI
RESEARCH SCIENTIST I

Tokey TAHMID
GRADUATE RESEARCH ASSISTANT

Stanimire TOMOV
RESEARCH ASSOCIATE PROFESSOR

Tonya WIMBERLEY
FINANCIAL ASSOCIATE

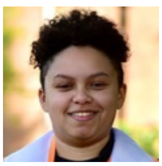

Eleigha WRANCHER
RESEARCH ASSISTANT

Asim YARKHAN
RESEARCH SCIENTIST II

ICL ALUMNI

THE SUCCESS OF ICL over the years was possible due to the efforts of its many talented staff and students.

Maksims ABALENKOVS
Carolyn AEBISCHER
Sudesh AGRAWAL
Bivek AGRAWAL
Emmanuel AGULLO
Kadir AKBUDAK
Mohammed AL FARHAN
Rabab AL-OMAIRY
Jennifer ALLGEYER
Wes ALVARO
Ed ANDERSON
Robert ANDERSON
Daniel ANDRZEJEWSKI
Thara ANGSKUN
Papa ARKHURST
Dorian ARNOLD
Rizwan ASHRAF
Cedric AUGONNET
Alan AYALA
Marc BABOULIN
Matthew BACHSTEIN
Paul BAGWELL
Zhaojun BAI
Ashwin BALAKRISHNAN
Richard BARRETT
Alex BASSI
John BATSON
David BATTLE
Micah BECK
Daniel BECKER
Dulceneia BECKER
Adam BEGUELIN
Annamaria BENZONI
Tom BERRY
Vincent BERTHOUX
Scott BETTS
Nikhil BHATIA
Daniel BIELICH
Noel BLACK
Laura BLACK
Susan BLACKFORD
Wesley BLAND
Kartheek BODANKI
David BOLT
Fernando BOND
Carolyn BOWERS
Vincent BRIDONNEAU
Barry BRITT
Cade BROWN
Randy BROWN
Murray BROWNE
Bonnie BROWNE
Cynthia BROWNE
Giuseppe BRUNO
Antonin BUKOVSKY
Greg BUNCH
Alfredo BUTTARI
Anthony CANINO
Domingo GIMENEZ CANOVAS
Chongxiao CAO
Qinglei CAO
Earl CARR
Henri CASANOVA
Cedric CASTAGNEDE
Tony CASTALDO
Ramkrishna CHAKRABARTY
Sharon CHAMBERS
Ali CHARARA
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
Sam CRAWFORD
David CRONK
Javier CUENCA
Manoel CUNHA
Yuanshun (SHAUN) DAI
Cricket DEANE
Remi DELMAS
Sameer DESHMUKH
Frederic DESPREZ
Jun DING
Jin DING
Ying DING
Martin DO
Simplice DONFACK
Leon DONG
Tingxing DONG
Nick DONGARRA
David DOOLIN
Joe DORRIS
Andrew DOWNEY
Mary DRAKE
Julio DRIGGS
Brian DRUM
Peng DU
David EBERIUS
Eduardo ECHAVARRIA
Victor EIJKHOUT
Brett ELLIS
Shawn ERICSON
Zachary EYLER-WALKER
Lisa EZZELL
Christoph FABIANEK
Graham FAGG
Hadeer FARAHAT
Mathieu FAVERGE
Diana FAYAD
Shengzhog FENG
Don FIKE
Salvatore FILIPPONE
Anna FINCHUM
Teresa FINCHUM
Mike FINGER
Jamie FINNEY
Markus FISCHER
Clement FLINT
Len FREEMAN
Xiaoquan FU
Erika FUENTES
Karl FUERLINGER
Megan FULLER
Edgar GABRIEL
Tracy GANGWER
Lynn GANGWER
Nathan GARNER
Kelley GARNER
Tina GARRISON
Adriana GARTIES
Peter GAULTNEY
Christoph GEILE
Jean PATRICK GELAS
Boris GELFEND
Damien GENET
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
Caleb HAN
Hanumantharayappa
Satomi HASEGAWA
Hidehiko 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
Matt JOHNSON
Aral JOHNSON
Sean JOLLY
Jan JONES
Kim 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
Hannah KNOCH
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
Weiran LI
Yinan LI
Xiang LI
Jiali LI
Chaoyang LIU
Kevin LONDON
Matt LONGLEY
Florent LOPEZ
Nuria LOSADA
Hatem LTAIEF
Yuechao LU
Daniel LUCIO
Richard LUCZAK
Xi LUO
Teng MA
Sticks MABAKANE
Ethan MANASS
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
Terry MOORE
Keith MOORE
Shirley MOORE
Robert MORGAN
Kishan MOTHERAMGARI
Steven MOULTON
Phil MUCCI
Daichi MUKUNOKI
Matthew NABITY
Shankar NARASIMHASWAMI
Rajib NATH
Fernando NAVARRO
John NELSON
Donnie NEWELL
Peter NEWTON
Phuong NGUYEN
Jonas NILSSON
Jakob OESTERGAARD
Caroline PAPADOPOULOS
Leelinda PARKER
Thibault PARPAITE
Thananon PATINYASAKDIKUL
Dilip PATLOLLA
Andy PEARSON
Yu PEI
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
Daniel SCHULTZ
Ken SCHWARTZ
Keith SEYMOUR
Farial SHAHNAZ
Hejer SHAIEK
Daniel SHARP
Brian SHEELY
Zhiao SHI
Sergei SHINKAREV
Majed SIDANI
Shilpa SINGHAL
Leighanne SISK
Matt SKINNER
Joan SNODERLY
Peter SOENDERGAARD
Raffaele SOLCA
Gwang SON
Fengguang SONG
Thomas SPENCER
Jeffrey STEILL
Erich STROHMAIER
Xiaobai SUN
Martin SWANY
Daisuke TAKAHASHI
Judi TALLEY
Ronald TAM
Chunyan TANG
Yuan TANG
Yusuke TANIMURA
Keita TERANISHI
Dan TERPSTRA
Joe THOMAS
John THURMAN
Francoise TISSEUR
Jude TOTH
Bernard TOURANCHEAU
Yaohung TSAI
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
Aaron WELCH
Scott WELLS
David WEST
R. CLINT WHALEY
Jody WHISNANT
James WHITE
Scotti WHITMIRE
Frank WINKLER
Susan WO
Felix WOLF
Stephen WOOD
Jiayi WU
You WU
Panruo WU
Wei 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
Dong ZHONG
Min ZHOU
Luke ZHOU

PARTNERSHIPS

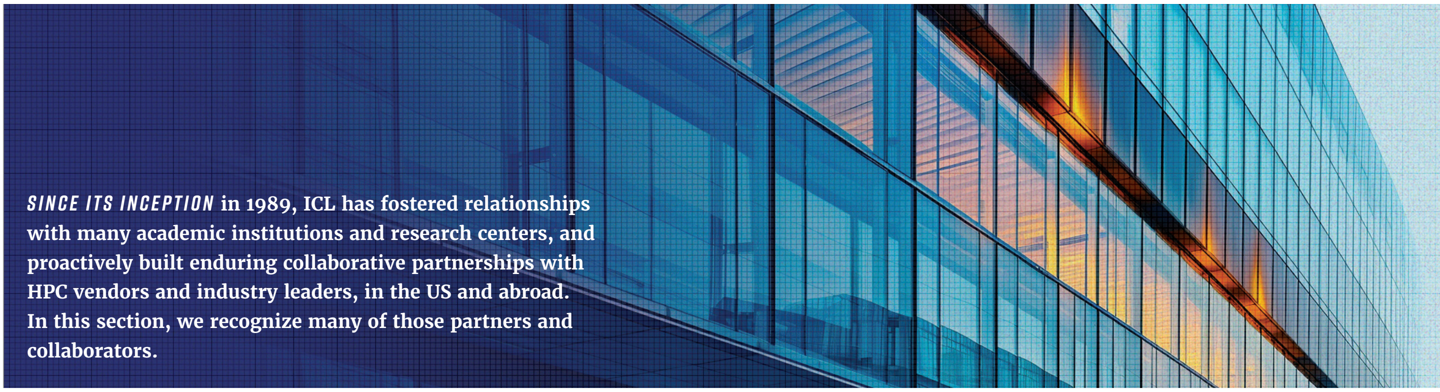

SINCE ITS INCEPTION **in 1989, ICL has fostered relationships with many academic institutions and research centers, and proactively built enduring collaborative partnerships with HPC vendors and industry leaders, in the US and abroad. In this section, we recognize many of those partners and collaborators.**

GOVERNMENT & ACADEMIC PARTNERSHIPS

- Argonne National Laboratory
- Mississippi State University
- SDSC San Diego Supercomputer Center
- University of Central Florida
- USC University of Southern California
- Emory University
- National Center for Supercomputing Applications, University of Illinois Urbana-Champaign
- Sandia National Laboratories
- The University of Chicago
- The University of Texas at El Paso
- Indiana University
- Northwestern University
- Stony Brook University
- University of Colorado Denver
- The University of Utah
- Iowa State University
- Oak Ridge National Laboratory
- Tennessee State University
- University of Florida
- University of Washington
- Berkeley Lab
- Pacific Northwest National Laboratory
- Texas A&M University
- University of Houston
- Virginia Tech
- Lawrence Livermore National Laboratory
- Princeton University
- Berkeley University of California
- Illinois University of Illinois at Urbana-Champaign
- Los Alamos National Laboratory
- Rice
- UCLA
- The University of Maine
- Michigan Technological University
- Rutgers The State University of New Jersey
- UC Riverside
- The University of New Mexico

INTERNATIONAL COLLABORATIONS

Central Institute for Applied Mathematics
JÜLICH, GERMANY

European Exascale Software Initiative
EUROPEAN UNION

Prometeus GmbH
MANNHEIM, GERMANY

Umeå University
UMEÅ, SWEDEN

Charles University
PRAGUE, CZECH REPUBLIC

Forschungszentrum Jülich
JÜLICH, GERMANY

RIKEN
WAKŌ, JAPAN

University of Bordeaux
BORDEAUX, FRANCE

Claude Bernard University Lyon 1
LYON, FRANCE

High Performance Computing Center Stuttgart
STUTTGART, GERMANY

Rutherford Appleton Laboratory
OXFORD, ENGLAND

University of Cape Town
CAPE TOWN, SOUTH AFRICA

Doshisha University
KYOTO, JAPAN

INRIA
FRANCE

Soongsil University
SEOUL, SOUTH KOREA

University of Manchester
MANCHESTER, ENGLAND

ENS de Lyon
LYON, FRANCE

Karlsruhe Institute of Technology
KARLSRUHE, GERMANY

Swiss Federal Institute of Technology Lausanne
SEOUL, SOUTH KOREA

University of Paris-Sud
PARIS, FRANCE

Erlangen Regional Data Center
ERLANGEN, GERMANY

King Abdullah University of Science and Technology
THUWAL, SAUDI ARABIA

Technical University of Dresden
DRESDEN, GERMANY

University of Picardie Jules Verne
AMIENS, FRANCE

ETH Zürich
ZÜRICH, SWITZERLAND

National Institute of Advanced Industrial Science and Technology
TSUKUBA, JAPAN

Technical University of Munich
MUNICH, GERMANY

University of Tsukuba
TSUKUBA, JAPAN

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

Parallel and HPC Application Software Exchange
TSUKUBA, JAPAN

Tokyo Institute of Technology
TOKYO, JAPAN

Vienna University of Technology
VIENNA, AUSTRIA

INDUSTRY PARTNERSHIPS

AMD
ENGILITY
IBM
MathWorks
Atos
EPAnalytics
IERUS
nag
CRAY
FUJITSU
intel
NVIDIA

ACKNOWLEDGMENTS

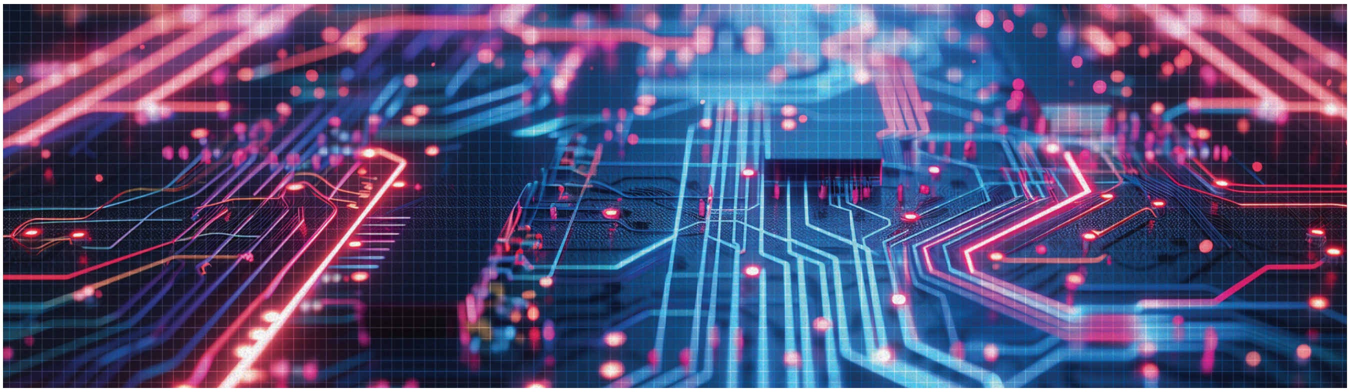

ICL'S WORK HAS been made possible through the support of the federal agencies that are charged with allocating public research funding, and we thank the following agencies:

AFOSR Air Force Office of Scientific Research

U.S. Department of Energy Office of Science

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. We gratefully acknowledge the following vendors for their generosity and support in recent years:

AMD

Atos

Engility

intel

MathWorks

NVIDIA

UT-Battelle

INNOVATIVE COMPUTING LABORATORY

FIND OUT MORE AT icl.utk.edu/