THE UNIVERSITY OF TENNESSEE KNOXVILLE

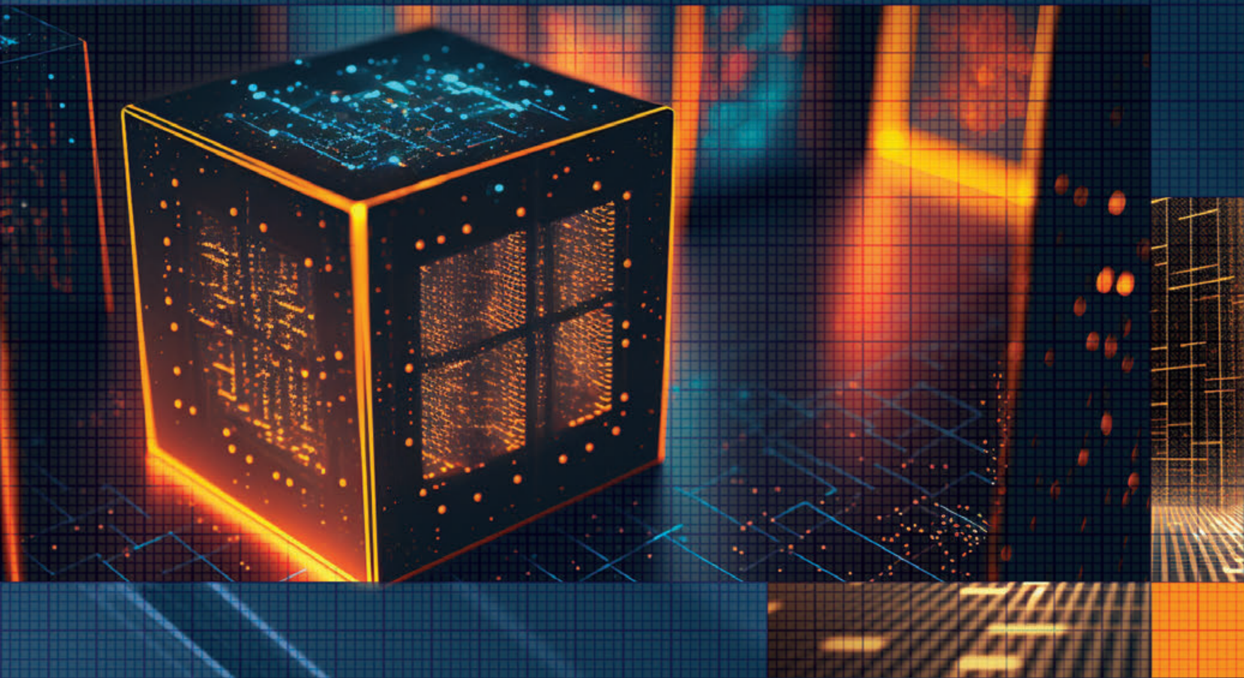

# INNOVATIVE COMPUTING LABORATORY

# ICL 2022 Year in Review

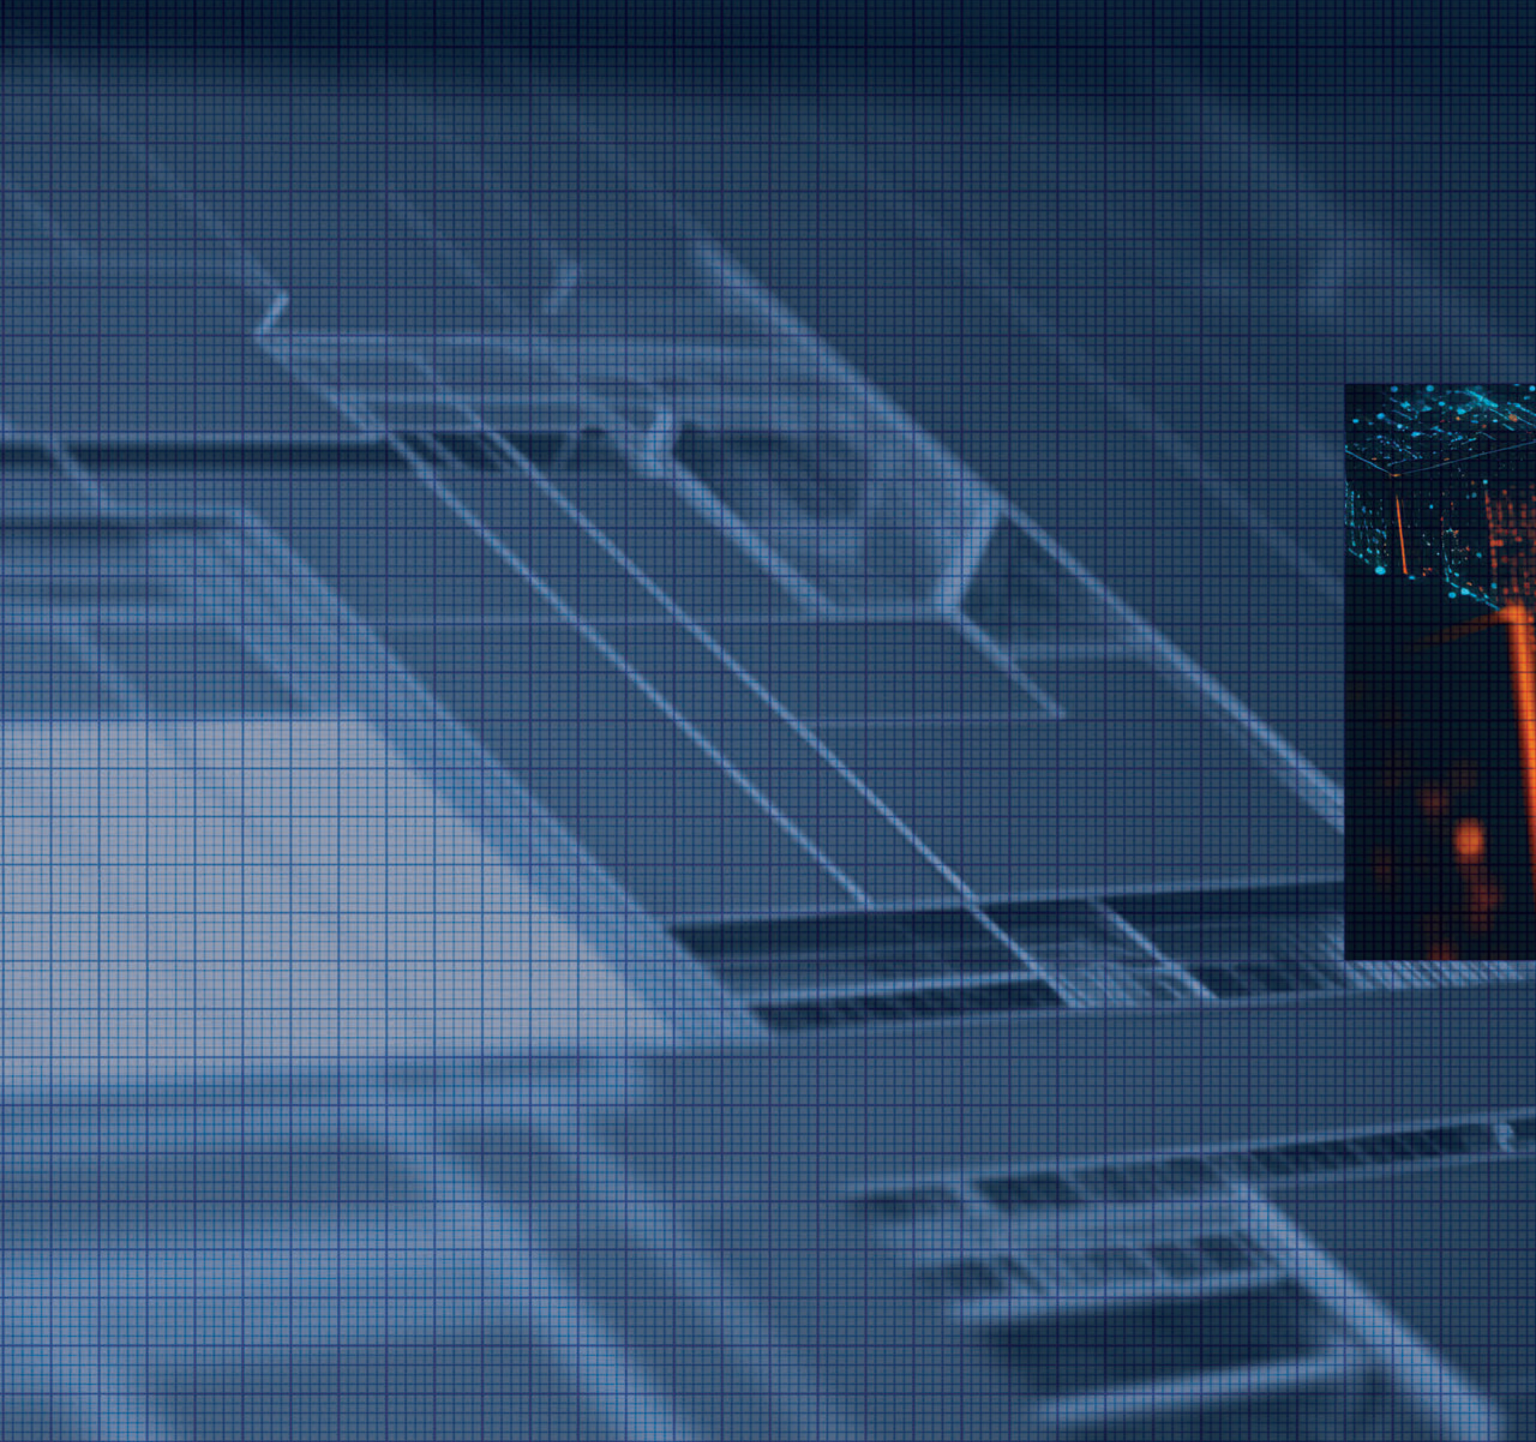

# INNOVATIVE COMPUTING LABORATORY 2022 Year in Review


ICL 2022 Year in Review

# FROM THE DIRECTOR

AS WE REFLECT on this year with heartfelt gratitude, we recognize 2022 as a year that will go down in ICL history as one packed with remarkable successes, highlights, events, and innovation.

The year started like the last one ended, under the effect of a global pandemic. COVID-19 continues to impact our daily life but there is a widespread acceptance of this being the new normal, and ICL returned to in-person operations in March 2022. Soon after, the major highlight of 2022 was announced: Jack Dongarra received the 2021 ACM A.M. Turing Award, the highest distinction in Computer Science, often referred to as the Nobel Prize in Computer Science. The Turing Award is a well-deserved individual recognition of Jack's accomplishments, but also a recognition of the high-performance computing (HPC) community as a whole. At the end of May, the HPC community saw another highlight: Oak Ridge National Lab's *Frontier* supercomputer officially became the first exascale system to operate after it successfully ran the HPL benchmark at an execution rate of more than one exaFLOP. This is a milestone in the history of supercomputing. A decade ago, experts were still uncertain whether an exaFLOP system was achievable. Still, they dared to imagine what a system delivering such performance could look like. In hindsight, it is surprising how closely the predictions about the design and characteristics of a potential exascale system match the *Frontier* supercomputer.

In June, Jack retired from his role as ICL director—a position he held for 33 years—and we hosted a party celebrating his career and legacy. The high number of current and former ICLers who joined the celebration reveals that the ICL family is a living community. We are happy that Jack remains an active part of the ICL family as professor emeritus, after his retirement as Electrical Engineering and Computer Science (EECS) distinguished professor and ICL director.

In August, I was granted the MathWorks Professorship in the EECS Department of the University of Tennessee (UTK) and assumed functions as ICL's new director. I am grateful to be back at UTK after being a postdoctoral associate at ICL, and having built a highly successful and collaborative research group at the Karlsruhe Institute of Technology (KIT) in Germany. I am honored to have the opportunity to contribute my experience to this institution, and I am looking forward to learning from the ICL team. Together, we will develop and implement progressive strategies.

Also in August, we had the largest retreat in ICL history: Together with Prof. Michela Taufer's Global Computing Lab (GCLab) and researchers from my group at KIT, the retreat hosted a group of 55 people presenting research activities and discussing challenges and opportunities in the national and international HPC community.

In November, all eyes were on Jack Dongarra's Turing Lecture that replaced this year's keynote at SC22. I would dare to say that, even for a few minutes, his participation turned the event into a rock concert.

Jack used the Turing Lecture as an opportunity to revisit ICL's success story and to look back at how the HPC community has developed over the years. The most impactful project in the nation's HPC history is the Department of Energy's (DOE) Exascale Computing Project (ECP), now entering its 7th year. This project invested approximately half of its budget into scientific software and applications, a game changer in the

community. We now see the fruits of this plan as a portfolio of software libraries and applications developed by ECP to enable scientific discoveries in *Frontier*. Our team is very proud of our heavy involvement in ECP and the development of software products to meet DOE's mission in exascale computing. These efforts include PAPI, PaRSEC, OpenMPI, SLATE, heFFTe, MAGMA, and Ginkgo. We are also part of the ECP's Center for Exascale Discretizations (CEED) and the Extreme-scale Software Development Kit (xSDK). These products are now heavily used by scientific applications and are great examples of the high-quality work achieved by ICL members. In 2023, ECP will reach the end of its course and face a thorough evaluation of its success. This will be a prime opportunity to, once again, highlight the technical innovation and high-quality software developed by our team.

Being the developers of a long list of high-quality software stacks puts us in pole position for future research activities and collaborations. The HPC landscape is changing, artificial intelligence and machine learning applications are driving the hardware design, and cloud computing is rapidly changing how and where we use computing resources. These trends come with new opportunities, and we have the skillset and infrastructure to be part of it. The legacy of outstanding algorithm research and high-quality software development makes ICL a global leader. This legacy is also a commitment: We have to work hard to fulfill the expectations the community has for us. While I am honored and excited to accept the challenge of leading ICL, I am also fully aware that our enterprise will require intentional team effort, and I am proud to say that I believe we have some of the best minds with us. I want to thank the ICL family and in particular its current members and outstanding staff for the help I have received since I joined as director. I am also grateful for the continued support and trust of our partners and collaborators in industry, academia, and government facilities, colleagues in the Tickle College of Engineering, and the University of Tennessee. Here's to a new year of exciting challenges and continued success.

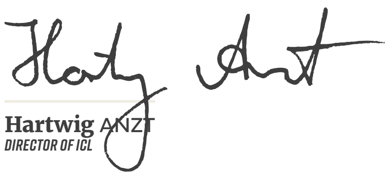

**Hartwig** ANZT
*DIRECTOR OF ICL*

# INTRODUCTION

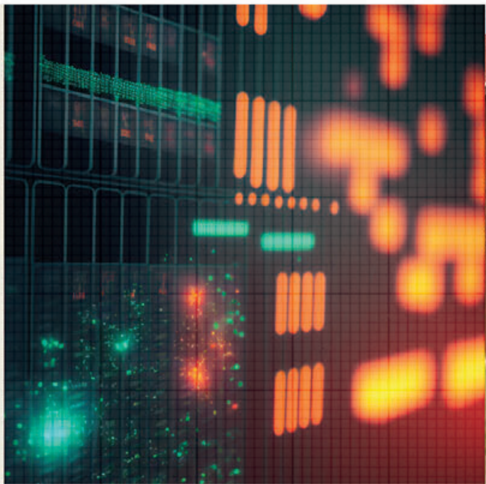

Claxton Building, Home of ICL

## ENABLING SCIENTIFIC HIGH-PERFORMANCE COMPUTING TODAY AND TOMORROW

THE Innovative Computing Laboratory (ICL) is a computer science research and development group 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 keen attention to message passing and fault tolerance.

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

"I'M PROUD AND EXCITED TO HAVE BEEN CHOSEN TO LEAD ICL AND TO BE WORKING WITH ITS BRIGHTEST MINDS TO CONTRIBUTE TO THE FUTURE OF HPC SCIENCE."

# INTRODUCING ICL'S NEW DIRECTOR

In July 2022, the Tickle College of Engineering named Hartwig Anzt as the new director of the Innovative Computing Laboratory, as well as the MathWorks Associate Professor in Scientific Computing. The professorship is the result of a $1 million donation from Mathworks, Inc. to create the MathWorks Endowed Professorship in Scientific Computing, based in the EECS department.

Anzt is a familiar face at ICL, having served as a postdoctoral researcher, research scientist, and consultant since joining the lab in June 2013. ICL's founder Jack Dongarra officially retired as director at the end of June.

Anzt previously served as a Helmholtz-Young-Investigator Group leader at the Steinbuch Centre for Computing at the Karlsruhe Institute of Technology. He earned his undergraduate degree at the Karlsruhe Institute of Technology (KIT) in 2006, graduate degrees from the University of Ottawa in 2008 and KIT in 2009, and his doctorate in 2012 from KIT, with Dongarra serving as one of his advisors.

He is the author of the MAGMA-sparse open-source software package, managing lead and developer of the Ginkgo numerical linear algebra library, and part of the Department of Energy's Exascale Computing Project delivering production-ready numerical linear algebra libraries. He is a member of the Institute of Electrical and Electronics Engineers Society, the Society of Industrial and Applied Mathematics, and the Association for Computing Machinery's Special Interest Group on High Performance Computing in the Regional Group of Middle Tennessee.

# HISTORY

JACK DONGARRA established ICL in 1989 when he received a dual appointment as a Distinguished Professor at UTK and Distinguished Scientist at Oak Ridge National Laboratory. Over thirty years later, ICL has grown into an internationally recognized research laboratory specializing in numerical linear algebra, distributed computing, and performance evaluation and benchmarking.

## 1989

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

The ***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.

## 1992

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

Still developed today, the ***Linear Algebra Package (LAPACK)*** is a standard software library for numerical linear algebra and is used in HPC systems worldwide.

## 1993

The ***TOP500*** was launched to improve and renew the Mannheim supercomputer statistics, which—at the time—had been in use for seven years. Updated twice a year since June 1993, TOP500 provides a reliable basis for tracking and detecting trends and milestones in high-performance computing.

1994
PVM becomes ICL's first R&D100 Award winner

## 1994

Version 1.0 of a standardized and portable message-passing system, called the ***Message Passing Interface (MPI)***, was introduced. MPI has since become the *de facto* standard for communication in parallel distributed computing systems.

## 1995

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

## 1997

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

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

## 1999

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

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

1999
ATLAS and NetSolve each received R&D 100 Awards

## 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 recognized with an R&D100 Award

## 2003

The ***HPC Challenge*** benchmark suite was developed for the Defense Advanced Research Projects Agency (DARPA) and measured a range of memory access patterns in order to provide a new in-depth analysis of system performance. An awards competition based on the benchmark was part of the annual SC conference from 2005 through 2016.

## 2006

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

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.

## 2008

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

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

## 2009

The ***International Exascale Software Project (IESP)*** brought together representatives of the global HPC community to plan and create a new software infrastructure for the extreme-scale systems that represent the future of computational science. At the conclusion of the worskhop series in 2012, *The International Exascale Software Roadmap* was published.

## 2010

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

## 2011

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

## 2012

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

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

2013
Jack Dongarra was awarded the ACM/IEEE Ken Kennedy Award at SC13

## 2013

The ***High Performance Conjugate Gradients (HPCG)*** benchmark was designed to measure performance that is representative of modern HPC capability by simulating patterns commonly found in real science and engineering applications. HPCG benchmark results have since been incorporated into TOP500 rankings.

The ***Big Data and Extreme-scale Computing (BDEC)*** workshop series was initiated to map out and account for the ways in which the major issues associated with "big data" intersect with national (and international) plans being laid out for achieving exascale computing. A follow-up worskhop series, ***BDEC2***, concluded in 2020.

## 2015

ICL's PAPI performance tool first released in 1999 continues to be widely used and actively developed today. In 2015, ***PAPI-Ex*** extended PAPI with measurement tools for changing hardware and software paradigms.

## 2016

ICL landed seven awards through the US Department of Energy's (DOE's) Exascale Computing Project (ECP) during the fall of 2016 and was the lead institution on four projects:

The ***Distributed Tasking for Exascale (DTE)*** project is focused on extending the capabilities of the PaRSEC runtime system on exascale architectures.

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

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

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

## 2017

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

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

## 2018

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

## 2019

Now known as ***HPL-MxP***, a new benchmark 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

***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.

The ***Scalable Run Time for Highly Parallel, Heterogeneous Systems (ScaRT)*** project aims to increase the scientific throughput of existing and future cyberinfrastructure platforms by reducing communication overheads and better matching the functionality of communication libraries to modern communication adapters.

## 2020

The ***Surrogate Benchmark Initiative (SBI)*** aims to provide benchmarks and tools for assessing deep neural network "surrogate" models. A surrogate model can imitate part or all of a simulation and produce the same outcomes while requiring less resources. Tools developed under SBI will evaluate these surrogate models to measure progress and inform the co-design of new HPC systems to support their use.

2022
Jack Dongarra was awarded the ACM A.M. Turing Award and Hartwig Anzt became the new director of ICL

# HIGHLIGHTS

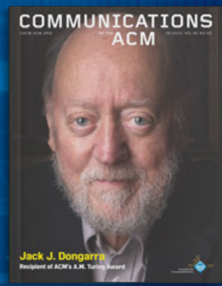


*COMMUNICATIONS OF THE ACM*
SPECIAL ISSUE:
tiny.utk.edu/acm-june22

## JACK DONGARRA HONORED WITH ACM A.M. TURING AWARD

ICL received a major recognition in 2022 when it was announced in March that ICL founder Jack Dongarra would be the recipient of the 2021 ACM A.M. Turing Award. The official citation for the award reads: "For his pioneering contributions to numerical algorithms and libraries that enabled high performance computational software to keep pace with exponential hardware improvements for over four decades." This achievement is the capstone of an illustrious career for Dongarra, as well as a reflection of the success of ICL under his tenure as director.

The ACM A.M. Turing Award is a prestigious award presented annually to individuals who have made significant contributions to the field of computer science. It is named after Alan Turing, a British mathematician and computer scientist who is considered to be the father of modern computer science. The award is often referred to as the "Nobel Prize in Computing" and is considered one of the highest honors that can be bestowed upon a computer scientist or engineer. The award is presented by the Association for Computing Machinery (ACM), a professional society for computer scientists, and the award winners are selected by a panel of experts in the field.

Dongarra was formally presented with the ACM A.M. Turing Award at the annual ACM Awards Banquet, which was held this year on Saturday, June 11 at the Palace Hotel in San Francisco.

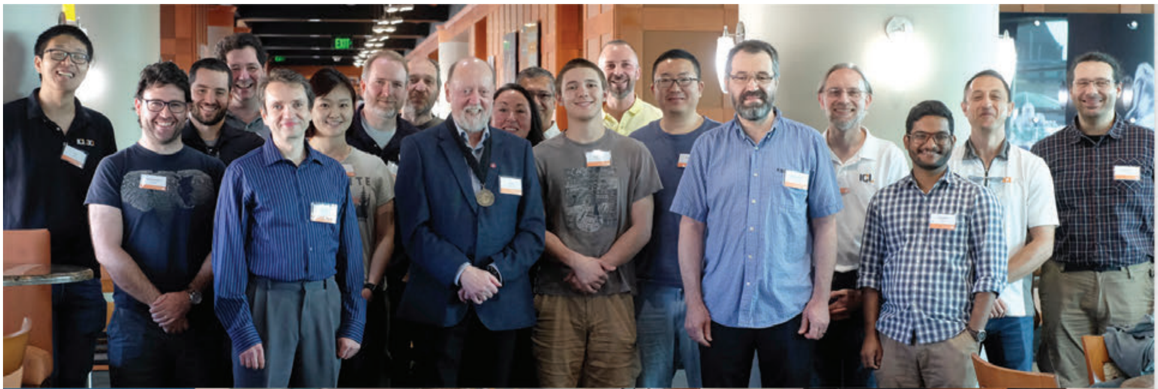

ICL group in attendance


(left to right)
**Richard Briggs, Donde Plowman, Jack Dongarra, Matthew Mench, and John Ragan**

# CHANCELLOR'S RECEPTION

UTK Chancellor Donde Plowman hosted a reception on May 10, 2022, to commemorate Jack Dongarra's Turing Award achievement. Attendees gathered at Neyland Stadium West Club for refreshments and mingling before the ceremony. In attendance, along with the Chancellor, were State Senator Richard Briggs and State Representative John Ragan. Following opening remarks from the Tickle College of Engineering Dean Matthew Mench and ICL's Piotr Lusczcek, Representative Ragan read to the assembly the contents of Tennessee House Joint Resolution 1381, a resolution to formally congratulate Dongarra from the state. Following, Chancellor Plowman presented to Dongarra the Chancellor's Medal. Dongarra then delivered some remarks highlighting the arc of his career from humble beginnings growing up in Chicago and spoke about the important role of the community that helped shape his success along the way, including his "many colleagues, post-docs, students and staff of ICL who really helped and influenced me over the years."

# HIGHLIGHTS


RECAP AND MORE PHOTOS AT
tiny.utk.edu/icl33

## CELEBRATING 33 YEARS OF ICL

ICL33 focused on Jack Dongarra's retirement, and was a celebration of all the small kindnesses, contributions, and collaborations along the path of his career. Current and former members traveled from nearby and abroad to spend an evening sharing stories, memories, and reminiscences of times at ICL. ICL's administrative and technical services groups worked jointly to provide a modest production for Jack that included an auspicious speech and toast from incoming MathWorks Professor and ICL director Hartwig Anzt, a tribute video with contributions from current and former ICLers from across the globe, an online chat between Jack and Terry Moore, and a presentation of a few special gifts to Jack and Sue Dongarra.

The evening saw plenty of laughs, especially during the showing of a music video tribute by ICL's "Fantastic Four", toasts and shared stories of Jack's inspiration and encouragement, and of course a few tears at heartfelt moments. While Jack's major accomplishments were mentioned and honored, it was the small stories that painted the true picture of his influence and impact. Of course, family featured prominently as well. Sue Dongarra was recognized for her contributions to ICL, and many photos were taken with the extended Dongarra family. The evening ended with photo sessions with all in attendance as well as individuals and groups represented at the event.

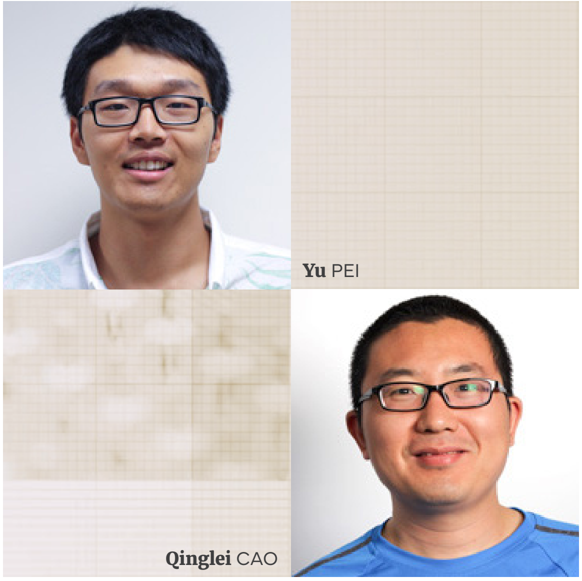


## 2022 ICL Graduates

In 2022, two of ICL's Graduate Research Assistants earned their PhDs from UTK under the guidance and mentorship of Jack Dongarra and ICL's research scientists. Congratulations to Yu Pei and Qinglei Cao. Both worked in ICL's distributed computing group under the guidance of George Bosilca.

## Jack Dongarra on *Sciencetown* Podcast

On a recent episode of the podcast *Sciencetown*, Jack Dongarra chatted with KAUST colleagues David Keyes, Hatem Ltaief, and Bilel Hadri about supercomputing developments and connections made throughout his career and at KAUST. Ltaief and Hadri, ICL alumni, demonstrate the global reach of ICL's vibrant network of friends and collaborators.

**https://soundcloud.com/sciencetown/20-jack-dongarra-on-supercomputing-connections-with-kaust**


## Natalie Beams on ECP Podcast

ICL's Natalie Beams was a featured guest on ECP 's *Let's Talk Exascale* Podcast in May 2022. The host of the podcast is ICL alumnus Scott Gibson. In the episode, the focus of the discussion is on the integration of Ginkgo as a numerical backend for the MFEM software ecosystem.

**https://www.exascaleproject.org/enabling-cross-project-research-to-strengthen-math-libraries-for-scientific-simulations/**

# RESEARCH

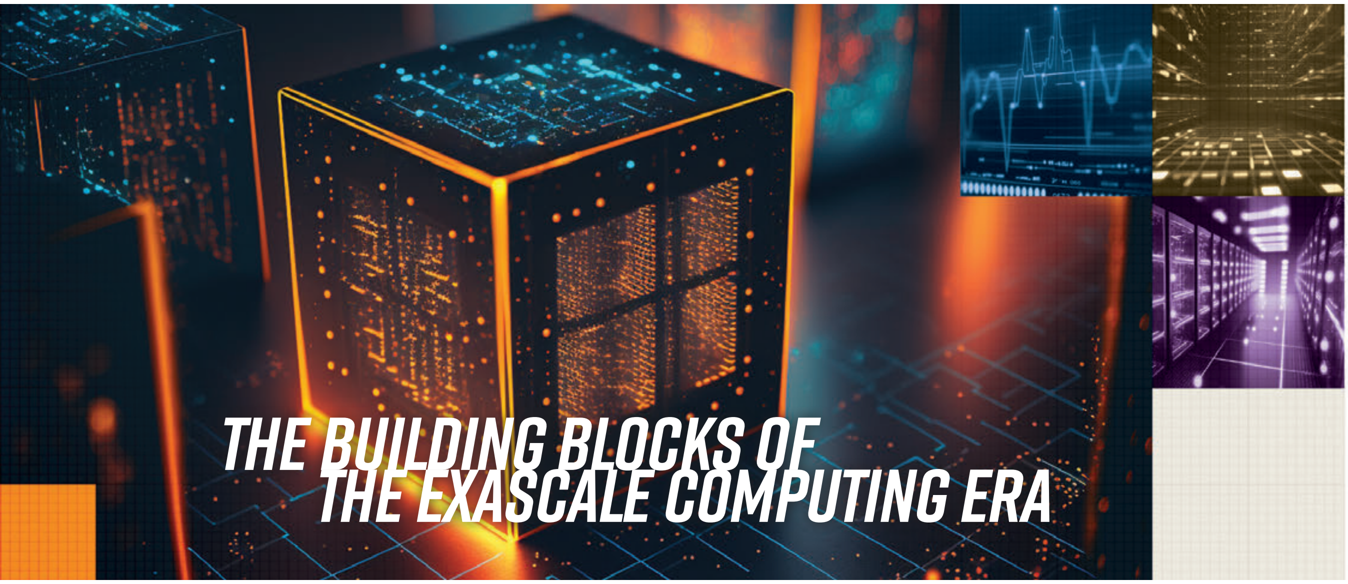

## THE BUILDING BLOCKS OF THE EXASCALE COMPUTING ERA

A MAJOR MILESTONE in high-performance computing was achieved in 2022 when ORNL's *Frontier* supercomputer officially became the first machine to achieve 1 exaFLOP performance on the Linpack benchmark. In some ways, the arrival of the exascale era can be seen as a culmination of years of research and development by the high-performance computing community. It is therefore fitting that this milestone coincided this year with the recognition of the work of Jack Dongarra and ICL in the form of the ACM A.M. Turing Award. In the journey to exascale computing, ICL has played a major role.

However, this milestone also highlights the important work ahead related to the challenges of harnessing exascale capabilities to power scientific applications. To compute at a rate of 1 exaFLOP requires 1 billion floating point units performing 1 billion calculations per second each. This level of parallelism and speed requires pushing the envelope of the development and optimization of software components, as well as fault-tolerant strategies to detect and protect against failures.

Building upon a strong track record of over 30 years of work in the research areas of numerical linear algebra, distributed computing, and performance evaluation and benchmarking, ICL is well-positioned to continue to contribute to these efforts.

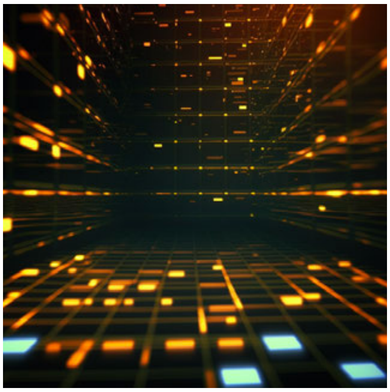

## NUMERICAL LINEAR ALGEBRA

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 for these libraries also be improved and modernized in order to achieve high 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.

## DISTRIBUTED COMPUTING

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.

## PERFORMANCE EVALUATION AND BENCHMARKING

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.

# RESEARCH

## 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 that will simplify the transition to the next generation of extreme-scale computer architecture through the leading-edge research it channels into its software deliverables.

FIND OUT MORE AT **icl.utk.edu/ballistic**

## Batched BLAS

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

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

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

## CEED

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

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

FIND OUT MORE AT **ceed.exascaleproject.org**

## CORES

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

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

FIND OUT MORE AT **icl.utk.edu/cores**

## DESC

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

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

FIND OUT MORE AT **hetcat-ccs.github.io**

## DPLASMA

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

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

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

## DTE

The Distributed Tasking for Exascale (DTE) project will extend 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 also provides a runtime component for dynamically executing tasks on heterogeneous distributed systems along with a productivity toolbox and development framework that supports multiple domain-specific languages (DSLs) and extensions and tools for debugging, trace collection, and analysis.

PaRSEC also enables fast prototyping DSLs to express the dependencies between tasks. 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 award is a vital extension and continuation of this effort. It will ensure that PaRSEC meets the critical needs of ECP application communities regarding scalability, interoperability, and productivity.

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

## EPEXA

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

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

FIND OUT MORE AT **icl.utk.edu/epexa**

# RESEARCH

## EVOLVE

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

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

## Exa-PAPI

The Exascale Performance Application Programming Interface (Exa-PAPI) project builds on the latest PAPI project and extends it with (1) performance counter monitoring capabilities for new and advanced ECP hardware, and software technologies; (2) fine-grained power management support; (3) functionality for performance counter analysis at "task granularity" for task-based runtime systems; and (4) "Software-defined Events" that originate from the ECP software stack and are currently treated as black boxes (i.e., communication libraries, math libraries, task-based runtime systems, etc.)

The objective is to enable monitoring of both types of performance events—hardware- and software-related events—in a uniform way, through one consistent PAPI interface. Third-party tools and application developers will have to handle only a single hook to PAPI in order to access all hardware performance counters in a system, including the new software-defined events.

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

## FFT-ECP / heFFTe

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

The main objective of the ECP FFT project is to design and develop a Highly Efficient FFTs for Exascale (heFFTe) library that provides fast and robust multidimensional FFTs for large-scale heterogeneous systems with multi-core processors and hardware accelerators. HeFFTe collects and leverages existing FFT capabilities while building a sustainable FFT library that minimizes data movements, optimizes MPI communications, overlaps computations with communications, and autotunes performance on various architectures and large scale-platforms. The current heFFTe v2.0 release achieves very good scalability on pre-exascale systems and performance close to 90% of the roofline peak.

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

## FIBER

The new benchmarking effort aims to benchmark a variety of implementations of parallel FFT implementations for the purposes of the Exascale Computing Project (ECP). Many scientific applications of strategic importance to DOE need multi-dimension Discrete Fourier transforms, including 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 to understand how the tradeoffs made by the developers affect the floating-point performance and mitigate communication overheads.

In its current incarnation, the FFT Benchmarking project released a technical report, ICL-UT-21-03, to provide an interim update on comprehensive testing and measuring various aspects of the FFT implementations for three-dimensional transforms on high performance systems. 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.

## Ginkgo

In the Ginkgo project, we develop high performance numerical linear algebra functionality reflecting the parallelism of modern HPC platforms. The focus is on GPU-accelerated systems, and Ginkgo can currently be used on AMD GPUs, Intel GPUs, and NVIDIA GPUs using backends written in the respective vendor languages. Ginkgo features a variety of iterative Krylov 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 features interfaces to several popular simulation frameworks, including MFEM, SUNDIALS, deal.ii, HyTeg, openCARP, and OpenFOAM. The Ginkgo library is open source and licensed under BSD 3-clause. In the development of Ginkgo, we aim for industry-level code quality standards, including Continuous Integration (CI) and Continuous Benchmarking (CB), comprehensive unit test coverage, and fulfilling the community policies of the extreme-scale Scientific Software Development Kit (xSDK) and the Extreme Scale Scientific Software Stack (E4S).

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

## HPCG

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

The HPCG 3.1 reference code was released in March 2019. In addition to bug fixes, 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 SC21 and features over 200 entries across the supercomputing landscape. TOP500.org has also tracked HPCG results since June 2017.

FIND OUT MORE AT **hpcg-benchmark.org/**

## HPL

The High Performance LINPACK (HPL) benchmark solves a dense linear system in double precision (64-bit arithmetic on distributed-memory computers. HPL is written in a portable ANSI C and requires an MPI implementation and either BLAS or 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.

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

## HPL-MxP

The HPL-MxP benchmark seeks to highlight the emerging convergence of high-performance computing (HPC) and artificial intelligence workloads. While traditional HPC focused on simulation runs for modeling phenomena in physics, chemistry, biology, and so on, the mathematical models that drive these computations require, for the most part, 64-bit accuracy. On the other hand, the machine learning methods that fuel advances in AI achieve desired results at 32-bit and even lower floating-point precision formats. This lesser demand for accuracy fueled a resurgence of interest in new hardware platforms that deliver a mix of unprecedented performance levels and energy savings to achieve the classification and recognition fidelity afforded by higher-accuracy formats HPL-MxP strives to unite these two realms by delivering a blend of modern algorithms and contemporary hardware while simultaneously connecting to the solver formulation of the decades-old HPL framework of benchmarking the largest supercomputing installations in the world.

FIND OUT MORE AT **hpl-mxp.org**

# RESEARCH

## LAPACK/ScaLAPACK

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

The latest major release of LAPACK, version 3.11.0, was released in November 2022. This version added a number of improvements along with big fixes. Highlights include safe scaling algorithms for 2-norm and Givens rotation, multishift QZ algorithm with aggressive early deflation, Householder reconstruction, level-3 BLAS solver for triangular Sylvester equations, and faster algorithms for Least Squares. A new version of ScaLAPACK, version 2.2.0, was also released in 2022, and added ILP64 support.

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

## MAGMA

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

MAGMA features LAPACK-compliant routines for multi-core CPUs enhanced with NVIDIA or AMD GPUs. MAGMA 2.5.4 now includes more than 400 routines that cover one-sided dense matrix factorizations and solvers, two-sided factorizations, and eigen/singular-value problem solvers, as well as a subset of highly optimized BLAS for GPUs. A MagmaDNN package has been added and further enhanced to provide high-performance data analytics, including functionalities for machine learning applications that use MAGMA as their computational back end. The MAGMA Sparse and MAGMA Batched packages have been included since MAGMA 1.6.

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

## MATEDOR

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

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

## OMPI-X

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

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

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

# oneAPI

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 the Innovative Computing Lab to spearhead the development of numerical linear algebra functionality in the oneAPI ecosystem by deploying a DPC++ kernels in the Ginkgo linear algebra library. This allows to run Ginkgo and applications that use Ginkgo's functionality on any hardware supporting the oneAPI industry standard, including embedded Intel GPUs and Intel's discrete GPUs that are expected to power the Aurora supercomputer.

# Open MPI

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

A completely new MPI 4.0–compliant implementation, Open MPI offers advantages for system and software vendors, application developers, and computer science researchers. ICL's efforts in the context of Open MPI have significantly improved its scalability, performance on many-core environments, and architecture-aware capabilities—such as adaptive shared memory behaviors and dynamic collective selection—making it ready for next-generation exascale challenges.

FIND OUT MORE AT **open-mpi.org**

# PAPI

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

PAPI 7.0.0 was announced on November 15th, 2022. This release offers several new components, including "intel_gpu" with monitoring capabilities on Intel GPUs; "sysdetect" (along with a new user API) for detecting details of the available hardware on a given compute system; a significant revision of the "rocm" component for AMD GPUs; the extension of the "cuda" component to enable performance monitoring on NVIDIA's compute capabilities 7.0 and beyond. PAPI 7.0.0 ships with a standalone "libsde" library and a new C++ API for software developers to define software-defined events from within their applications.

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

# PaRSEC

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

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

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

# RESEARCH

## PLASMA

The Parallel Linear Algebra Software for Multi-core Architectures (PLASMA) implements a set of fundamental linear algebra routines using the latest updates to the OpenMP standard. 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 that now includes offloading functionality and enables porting to hardware accelerators. The latest PLASMA release, version 21.8.29 from August 2021, added a transpose option to xGETRS() and xGELS() functions, convenience scripts for C and Fortran examples, and a Python script for quickly launching tests in addition to fixing bugs for corner cases in LU and norm functions.

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

## REE-HPC

The Numerically-Exact Relativistic Many-Body Electronic Structure of Heavy Elements project is a collaborative effort focused on 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 heterogenous computing to assist with tuning the implementation of novel methods for execution on large-scale HPC platforms, including exascale machines.

This effort enables advances in global needs involving Rare Earth Elements (REEs) and actinides. REEs are essential in applied technologies that include communications, computing, medical capabilities, green energy, and defense. Actinides are amongst the least studied elements, with many being essential in nuclear efforts relevant to defense, energy, and medical treatments.

## ScaRT

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

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

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

## SLATE

For decades, ICL has applied algorithmic and technological innovations to pioneering, implementing, and disseminating dense linear algebra software—including the LAPACK (Linear Algebra PACKage) and ScaLAPACK (Scalable Linear Algebra PACKage) libraries. The SLATE (Software for Linear Algebra Targeting Exascale) 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, some 25 years ago. In the past two decades, HPC has witnessed tectonic shifts in hardware and software technology, and many algorithmic innovations in scientific computing. SLATE was conceived to be a replacement for ScaLAPACK, boasting superior performance and scalability in HPC's modern, heterogeneous, distributed-memory environments.

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

## SBI

The DOE-funded Surrogate Benchmark Initiative (SBI) is a collaborative effort involving Indiana University, UTK/ICL, and Rutgers University to provide new benchmarks and tools for assessing deep neural network "surrogate" models. 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. SBI 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.

FIND OUT MORE AT **icl.utk.edu/research/surrogates**

## TOP500

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

In November 2022, the 60th TOP500 list was unveiled during SC22 held in Dallas, Texas. Yet again, U.S. took the crown with *Frontier*, with a mix of AMD CPUs and GPUs featuring Cray Slingshot 11 interconnect and built by HPE. The system's mix of nearly 9 thousand cores propelled *Frontier* into exascale era achieving 1.1 Eflop/s in the HPL benchmark, making it the fastest supercomputer in the world by over 2-fold factor over Japan's *Fugaku* machine, a CPU-only ARM system. *Frontier* now held the first spot for two consecutive lists.

FIND OUT MORE AT **top500.org**

## ULFM

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

Using the constructs defined by ULFM, applications and libraries drive the recovery of the 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 middlewares are building on top of ULFM to deliver scalable fault tolerance. Notable additions include the CoArray Fortran language and SAP databases. ULFM software is available in recent versions of MPICH and Open MPI.

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

## xSDK4ECP

The Extreme-Scale Scientific Software Development Kit for the Exascale Computing Project (xSDK4ECP) is a collaborative effort between Argonne National Laboratory, ICL, Karlsruhe Institute of Technology, Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, Sandia National Laboratory, and the University of California, Berkeley. The project aims 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.

To ensure the consistency of naming conventions, runtime behavior, and installation procedures, xSDK provides requirements and guidelines that are influential throughout the software development phase. xSDK lightens the burden on system administrators and application developers because each xSDK package provides a Spack installation script that can be invoked independently or through the installation of the xSDK's Spack package. In addition, xSDK now ships with a set of curated examples showing potential package integrations into application exemplars. ICL's MAGMA, PLASMA, SLATE, and heFFTe libraries are now included in the latest release, xSDK 0.8.

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

# PUBLICATIONS

Abdelfattah, A., P. Ghysels, W. Boukaram, S. Tomov, X. Sherry Li, and J. Dongarra
**Addressing Irregular Patterns of Matrix Computations on GPUs and Their Impact on Applications Powered by Sparse Direct Solvers**
*SC '22: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, Dallas, TX, IEEE Press, November 2022

Abdelfattah, A., S. Tomov, and J. Dongarra
**Batch QR Factorization on GPUs: Design, Optimization, and Tuning**
*Lecture Notes in Computer Science*, Vol. 13350, Cham, Springer International Publishing, June 2022

Abdulah, S., Q. Cao, Y. Pei, G. Bosilca, J. Dongarra, M. G. Genton, D. E. Keyes, H. Ltaief, and Y. Sun
**Accelerating Geostatistical Modeling and Prediction With Mixed-Precision Computations: A High-Productivity Approach With PaRSEC**
*IEEE Transactions on Parallel and Distributed Systems*, Vol. 33 Issue 4, pp. 964 - 976, April 2022

Agullo, E., M. Altenbernd, H. Anzt, L. Bautista-Gomez, T. Benacchio, L. Bonaventura, H-J. Bungartz, S. Chatterjee, F. M. Ciorba, N. DeBardeleben, et al.
**Resiliency in numerical algorithm design for extreme scale simulations**
*The International Journal of High Performance Computing Applications*, Vol. 36 Issue 2, pp. 251 - 285, March 2022

Aliaga, J. I., H. Anzt, T. Grützmacher, E. S. Quintana-Orti, and A. E. Tomás
**Compression and load balancing for efficient sparse matrix-vector product on multicore processors and graphics processing units**
*Concurrency and Computation: Practice and Experience*, Vol. 34 Issue 14, June 2022

Aliaga, J. I., H. Anzt, T. Grützmacher, E. S. Quintana-Ortí, and A. E. Tomás
**Compressed basis GMRES on high-performance graphics processing units**
*The International Journal of High Performance Computing Applications*, May 2022

Anzt, H., M. Casas, C. I. Malossi, E. S. Quintana-Ortí, F. Scheidegger, and S. Zhuang
**Approximate Computing for Scientific Applications**
*Approximate Computing Techniques*, 322: Springer International Publishing, pp. 415 - 465, January 2022

Anzt, H., T. Cojean, G. Flegar, F. Göbel, T. Grützmacher, P. Nayak, T. Ribizel, Y. Mike Tsai, and E. S. Quintana-Ortí
**Ginkgo: A Modern Linear Operator Algebra Framework for High Performance Computing,**
*ACM Transactions on Mathematical Software*, Vol. 48 Issue 12, pp. 1 - 33, March 2022

Ayala, A., S. Tomov, M. Stoyanov, A. Haidar, and J. Dongarra
**Performance Analysis of Parallel FFT on Large Multi-GPU Systems**
*2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, Lyon, France, August 2022

Bak, S., C. Bertoni, S. Boehm, R. Budiardja, B. M. Chapman, J. Doerfert, M. Eisenbach, H. Finkel, O. Hernandez, J. Huber, et al.
**OpenMP application experiences: Porting to accelerated nodes**
*Parallel Computing*, Vol. 109, March 2022

Benoit, A., Y. Du, T. Herault, L. Marchal, G. Pallez, L. Perotin, Y. Robert, H. Sun, and F. Vivien
**Checkpointing à la Young/Daly: An Overview**
*IC3-2022: Proceedings of the 2022 Fourteenth International Conference on Contemporary Computing*, Noida, India, ACM Press, pp. 701-710, August 2022

Bosilca, G., A. Bouteiller, T. Herault, V. Le Fèvre, Y. Robert, and J. Dongarra
**Comparing Distributed Termination Detection Algorithms for Modern HPC Platforms**
*International Journal of Networking and Computing*, Vol. 12 Issue 1, pp. 26 - 46, January 2022

Bouteiller, A. and Bosilca, G.
**Implicit Actions and Non-blocking Failure Recovery with MPI**
*Workshop on Fault Tolerance for HPC at eXtreme Scale*, Dallas, TX, November 2022

Cao, Q., G. Bosilca, N. Losada, W. Wu, D. Zhong, and J. Dongarra
**Evaluating Data Redistribution in PaRSEC**
*IEEE Transactions on Parallel and Distributed Systems*, Vol. 33 No. 8, pp. 1856-1872, August 2022

Cao, Q., R. Alomairy, Y. Pei, G. Bosilca, H. Ltaief, D. Keyes, and J. Dongarra
**A Framework to Exploit Data Sparsity in Tile Low-Rank Cholesky Factorization**
*2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2022

Cao, Q., S. Abdulah, R. Alomairy, Y. Pei, P. Nag, G. Bosilca, J. Dongarra, M. G. Genton, D. Keyes, H. Ltaief, and Y. Sun
**Reshaping Geostatistical Modeling and Prediction for Extreme-Scale Environmental Applications**
*SC '22: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, Dallas, TX, IEEE Press, November 2022

Cayrols, S., J. Li, G. Bosilca, S. Tomov, A. Ayala, and J. Dongarra
**Lossy all-to-all exchange for accelerating parallel 3-D FFTs on hybrid architectures with GPUs**
*2022 IEEE International Conference on Cluster Computing (CLUSTER)*, pp. 152-160, September 2022

Cojean, T., Y-H. Mike Tsai, and H. Anzt
**Ginkgo—A math library designed for platform portability**
*Parallel Computing*, Vol. 111, pp. 102902, February 2022

Dongarra, J.
**The Evolution of Mathematical Software**
*Communications of the ACM*, Vol. 65 No. 12, pp. 66-72, December 2022

Du, Y., G. Pallez, L. Marchal, and Y. Robert
**Optimal Checkpointing Strategies for Iterative Applications**
*IEEE Transactions on Parallel and Distributed Systems*, Volume 33 Issue 3, pp. 507-522, March 2022

Fortenberry, A., and S. Tomov
**Extending MAGMA Portability with OneAPI**
*The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC22), Ninth Workshop on Accelerator Programming Using Directives (WACCPD 2022)*, Dallas, TX, November 2022

Funk, Y., M. Götz, and H. Anzt
**Prediction of Optimal Solvers for Sparse Linear Systems Using Deep Learning**
*Proceedings of the 2022 SIAM Conference on Parallel Processing for Scientific Computing (PP)*, pp. 14-24, 2022

Kashi, A., P. Nayak, D. Kulkarni, A. Scheinberg, P. Lin, and H. Anzt
**Batched sparse iterative solvers on GPU for the collision operator for fusion plasma simulations**
*2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Lyon, France, IEEE, July 2022

Kovalchuk, S. V., V. V. Krzhizhanovskaya, M. Paszyński, D. Kranzlmüller, J. Dongarra, and P. M. A. Sloot
**Computational science for a better future**
*Journal of Computational Science*, Vol. 62, July 2022

Luszczek, P., and C. Brown
**Surrogate ML/AI Model Benchmarking for FAIR Principles' Conformance**
*2022 IEEE High Performance Extreme Computing Conference (HPEC)*, Waltham, MA, September 2022

Murray, R., J. Demmel, M. W. Mahoney, B. N. Erichson, M. Melnichenko, O. Asif Malik, L. Grigori, P. Luszczek, M. Dereziński, M. E. Lopes, et al.
**Randomized Numerical Linear Algebra: A Perspective on the Field with an Eye to Software**
*University of California, Berkeley EECS Technical Report*, No. UCB/EECS-2022-258, November 2022

Nance, D., S. Tomov, and K. Wong
**A Python Library for Matrix Algebra on GPU and Multicore Architectures**
*2022 IEEE 19th International Conference on Mobile Ad Hoc and Smart Systems (MASS)*, Denver, CO, IEEE, December 2022

Penchoff, D. A., C. C. Peterson, E. M. Wrancher, G. Bosilca, R. J. Harrison, E. F. Valeev, and P. D. Benny
**Evaluations of molecular modeling and machine learning for predictive capabilities in binding of lanthanum and actinium with carboxylic acids**
*Journal of Radioanalytical and Nuclear Chemistry*, December 2022

Schuchart, J., P. Nookala, M. Mahdi Javanmard, T. Herault, E. F. Valeev, G. Bosilca, and R. J. Harrison
**Generalized Flow-Graph Programming Using Template Task-Graphs: Initial Implementation and Assessment**
*2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Lyon, France, IEEE, July 2022

Schuchart, J., P. Nookala, T. Herault, E. F. Valeev, and G. Bosilca
**Pushing the Boundaries of Small Tasks: Scalable Low-Overhead Data-Flow Programming in TTG**
*2022 IEEE International Conference on Cluster Computing (CLUSTER)*, Heidelberg, Germany, IEEE, September 2022

Sid-Lakhdar, W. M., M. Aznaveh, P. Luszczek, and J. Dongarra
**Deep Gaussian process with multitask and transfer learning for performance optimization**
*2022 IEEE High Performance Extreme Computing Conference (HPEC)*, Waltham, MA, September 2022


Tsai, Y-H. M., T. Cojean, and H. Anzt
**Providing performance portable numerics for Intel GPUs**
*Concurrency and Computation: Practice and Experience*, Vol. 17, October 2022

Tsai, Y. M., T. Cojean, and H. Anzt
**Porting Sparse Linear Algebra to Intel GPUs**
*Euro-Par 2021: Parallel Processing Workshops*, Vol. 13098, Lisbon, Portugal, Springer International Publishing, pp. 57 - 68, June 2022

Whitlock, M., N. Morales, G. Bosilca, A. Bouteiller, B. Nicolae, K. Teranishi, E. Giem, and V. Sarkar
**Integrating process, control-flow, and data resiliency layers using a hybrid Fenix/Kokkos approach**
*2022 IEEE International Conference on Cluster Computing (CLUSTER 2022)*, Heidelberg, Germany, September 2022


Zhong, D., Q. Cao, G. Bosilca, and J. Dongarra
**Using long vector extensions for MPI reductions**
*Parallel Computing*, Vol. 109, March 2022


# 2022 ICL TECHNICAL REPORT SERIES

Alomairy, R., M. Gates, S. Cayrols, D. Sukkari, K. Akbudak, A. YarKhan, P. Bagwell, and J. Dongarra
**Communication Avoiding LU with Tournament Pivoting in SLATE,**
SLATE Working Notes, No. 18, ICL-UT-22-01, January 2022

Ayala, A., S. Tomov, P. Luszczek, S. Cayrols, G. Ragghianti, and J. Dongarra
**FFT Benchmark Performance Experiments on Systems Targeting Exascale**
ICL Technical Report, No. ICL-UT-22-02, March 2022

Reed, D., D. Gannon, and J. Dongarra
**Reinventing High Performance Computing: Challenges and Opportunities**
ICL Technical Report, No. ICL-UT-22-03, March 2022

Cayrols, S., J. Li, G. Bosilca, S. Tomov, A. Ayala, and J. Dongarra
**Mixed precision and approximate 3D FFTs: Speed for accuracy trade-off with GPU-aware MPI and run-time data compression**
ICL Technical Report, No. ICL-UT-22-04, May 2022

Dongarra, J., and A. Geist
**Report on the Oak Ridge National Laboratory's *Frontier* System**
ICL Technical Report, No. ICL-UT-22-05, May 2022

Sid-Lakhdar, W. M., S. Cayrols, D. Bielich, A. Abdelfattah, P. Luszczek, M. Gates, S. Tomov, H. Johansen, D. Williams-Young, T. A. Davis, et al.
**PAQR: Pivoting Avoiding QR factorization**
ICL Technical Report, No. ICL-UT-22-06, June 2022

Ayala, A., S. Tomov, P. Luszczek, S. Cayrols, G. Ragghianti, and J. Dongarra
**Analysis of the Communication and Computation Cost of FFT Libraries towards Exascale**
ICL Technical Report, No. ICL-UT-22-07, July 2022

FOR A COMPLETE LIST OF PUBLICATIONS
AND TECH REPORTS
**icl.utk.edu/publications**

# CONFERENCES

## 2022 CONFERENCES ATTENDED

*JAN 25 - FEB 1* TRIESTE, ITALY
**The Master in High Performance Computing (MHPC)**

*FEB 2* KNOXVILLE, TN
**Virtual SPARKS: Artificial Intelligence**

*FEB 3-4* ROANOKE, VA
**EPEXA Meeting**

*FEB 23-26* VIRTUAL
**SIAM Conference on Parallel Processing for Scientific Computing (PP22)**

*MAR 20-24* SAN DIEGO, CA
**American Chemical Society Spring 2022 Meeting**

*MAR 21-24* VIRTUAL
**GTC 2022**

*APR 2-6* VIRTUAL
**Principles and Practice of Parallel Programming 2022 (PPoPP 2022)**

*APR 3-8* KONA, HI
**International Conference on Methods and Applications of Radioanalytical Chemistry (MARC)**

*MAY 2-6* VIRTUAL
**ECP Annual Meeting**

*MAY 29 - JUN 2* HAMBURG, GERMANY
**ISC High Performance 2022**

*MAY 30 - JUN 3* VIRTUAL
**36th IEEE International Parallel & Distributed Processing Symposium (IPDPS)**

*JUN 13-15* NATIONAL HARBOR, MD
**TechConnect World Innovation Special Symposium: Innovations in Rare Earths and Critical Minerals**

*JUN 19-23* LAKE TAHOE, CA
**Scalable Tools Workshop**

*JUN 19-24* ATHENS, GREECE
**International HPC Summer School 2022**

*JUN 21-23* LONDON, UK
**International Conference on Computational Science (ICCS)**

*JUL 10-14* BOSTON, MA
**PEARC22**

*AUG 3-5* DENVER, CO
**NSF Broadening Participation in Computing Plans Workshop**

*AUG 22-26* CHICAGO, IL
**American Chemical Society Fall 2022 Meeting**

*AUG 22-24* COLUMBUS, OH
**10th Annual MVAPICH User Group (MUG) Meeting**

*AUG 23-25* KINGSPORT, TN
**Smoky Mountains Computational Sciences & Engineering Conference**

*AUG 29 - SEP 1* MONTEREY, CA
**Monterey Data Conference 2022**

*SEP 6-9* HEIDELBERG, GERMANY
**IEEE Cluster 2022**

*SEP 6-9* LYON, FRANCE
**Workshop on Clusters, Clouds, and Data for Scientific Computing (CCDSC 2022)**

*SEP 11-14* GDANSK, POLAND
**14th International Conference on Parallel Processing and Applied Mathematics (PPAM 2022)**

*SEP 17-24* HEIDELBERG, GERMANY
**Heidelberg Laureate Forum**

*SEP 27-30* CHATTANOOGA, TN
**EuroMPI**

*SEP 28-30* URBANA, IL
**14th JLESC Workshop**

*SEP 28-30* CHATTANOOGA, TN
**MPI Forum**

*OCT 26-28* MÜNSTER, GERMANY
**MaRDI Workshop on Scientific Computing**

*OCT 31 - NOV 1* ATLANTA, GA
**Radiobioassay and Radiochemical Measurements Conference**

*NOV 12-18* DALLAS, TX
**SC22**

*NOV 15-18* LAUSANNE, SWITZERLAND
**CECAM: Challenges and Advances in Solving Eigenproblems for Electronic-Structure Theory**

*NOV 29 - DEC 2*
BUENOS AIRES, ARGENTINA
**2022 FEWSUS Annual International Symposium**

*DEC 5-8* VIRTUAL
**MPI Forum**

*DEC 5-9* KARLSRUHE, GERMANY
**EuroHPC Project Meeting**

*DEC 5-10* BANGKOK, THAILAND
**EU-ASEAN High-Performance Computing School 2022**

*DEC 18-21* BANGALORE, INDIA
**29th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC)**

# CONFERENCE HIGHLIGHTS


## ISC 2022 Special Session Honors Jack Dongarra

This year's International Supercomputing Conference (ISC 2022) was held from May 29th through June 2nd in Hamburg, Germany. ICL's Jack Dongarra and Piotr Luszczek were among 3,007 attendees and 137 exhibitors for the first in-person ISC in three years.

Dongarra and Luszczek were joined by incoming Mathworks Professor Hartwig Anzt to give a tutorial on Modern Mixed-Precision Methods wherein they spoke on recent algorithmic progress in exploiting multiple precisions for increased efficiency in performance, communication, and storage. Additionally, Dongarra was recognized in a special session for receiving the 2021 ACM A.M. Turing Award. Dongarra was also named as the ISC 2023 deputy chair.

WATCH THE SPECIAL SESSION AT **tiny.utk.edu/jack@isc22**

## Fastest Code Award at 2022 International HPC Summer School

The International HPC Summer School is an invitation-only event that aims to gather the best PhD students and young scientists in HPC from the US, Canada, Europe and Japan. ICL's Alan Alaya was first selected to go to the fully expense-paid event in 2019 in Kobe, Japan. This year, the HPC Summer School returned to in-person interaction in Greece. Ayala achieved the first place award in a competition where all participants are invited to optimize HPC code using programming tools like OpenMP, MPI, CUDA and OpenACC.

## Best Paper Award at PPAM 2022

In September, members of ICL attended and presented at PPAM 2022 in Gdansk, Poland. ICL director Hartwig Anzt and former director Jack Dongarra gave keynote talks titled *Lossy Compression and Mixed Precision Strategies for Memory-Bound Linear Algebra* and *Where We Are Today And A Look Into The Future*, respectively.

ICL members Natalie Beams and Hartwig Anzt (along with KIT's Yu-Hsiang "Mike" Tsai) took home a Best Paper Award with their work, "Mixed Precision Algebraic Multigrid on GPUs." The paper introduces Ginkgo's algebraic multigrid (AMG) solver and demonstrates its platform portability on NVIDIA, AMD, and Intel GPUs.

# CONFERENCE HIGHLIGHTS



## ICL Representation at SC22 Headlined by Turing Award Lecture

Jack Dongarra's Turing Award Lecture served as the keynote for the SC22 conference in Dallas, Texas, and was delivered to what was described as a rock concert-like atmosphere. Dongarra's lecture was a look back on his research activiites over the years in the context of the history of high-performance computing.

Beyond Jack's Turing Lecture, ICL was well represented by a contingent of 20 researchers, staff, and students who attended SC22 and participated in various capacities in the six-day program. ICL's Heike Jagode served as the SC22 Technical Program Vice Chair and Yves Robert served as the Posters Vice Chair.  Multiple ICL researchers presented in panels and sessions, and presented talks at the University of Tennessee booth which represented ICL as well as the Global Computing Lab (GCLab) and the University of Tennessee Chattanooga's SimCenter.

With SC22 once again featuring in-person attendance, members of ICL and other affiliated attendees were able to meet up for dinner in Dallas for the traditional ICL alumni dinner. The group of over 50 attendees included current ICL and GCLab members, and ICL alumni.

Yves Robert and Heike Jagode

TURING LECTURE RECAP AT
**tiny.utk.edu/turing-lecture**

Official SC22 photos by Jo Ramsey, SC Photography


## Gordon Bell Prize Finalist Work Highlights ICL's Mission of Enabling Technologies for HPC

The ACM's 2022 Gordon Bell Prize was announced at SC22. The award is the most prestigious recognition for achievement and innovation in the application of high-performance parallel computing to the field of science and technology. A team including members from ICL and King Abdullah University of Science and Technology (KAUST) was among the finalists.

Members of ICL's distributed computing group, led by George Bosilca, worked with the team at KAUST to integrate ICL's PaRSEC runtime system into the ExaGeoStat global weather prediction model. ExaGeoStat relies on PaRSEC for scheduling computational tasks on massively parallel systems, enabling statisticians to tackle computationally challenging scientific problems at extreme scales.

In the work which earned the Gordon Bell Prize Finalist distinction, the team first tested their method's effectiveness and scalability on notable HPC systems including KAUST's *Shaheen II* supercomputer, the High-Performance Computing Center Stuttgart's *Hawk* supercomputer, and the Oak Ridge National Laboratory's *Summit* supercomputer. This successful demonstration won them access to one of the world's most powerful supercomputers, the *Fugaku* system located at Japan's RIKEN Center for Computational Science.

Highlighting the success of the collaboration, Bosilca said, "In addition to its scientific and societal value, this collaboration highlights the importance of computer science research to develop solutions that address the increasingly varied needs of computational science, and the importance of co-design to realize successful integrations between science domains and runtimes. While this is true at any scale, it becomes critical for efficient exascale computing."

## Students@SC

Jack Dongarra donated his SC22 speaker's honorarium to the Students@SC program. This donation supported travel costs for the SC22 Lead Student Volunteers, who had the opportunity to meet with Dongarra to thank him, and to ask questions about his experiences and perspectives on HPC.

# PARTNERSHIPS

ICL has fostered relationships with academic institutions and research centers as well as HPC industry leaders in the United States and abroad.

## GOVERNMENT & ACADEMIC PARTNERSHIPS

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

## GCLab

In June 2018, Michela Taufer joined UTK's Department of Electrical Engineering and Computer Science as the Jack Dongarra Professor in High Performance Computing. Her research group, the Global Computing Lab (GCLab), focuses on various aspects of HPC and scientific computing—including computational chemistry and chemical engineering, pharmaceutical sciences, seismology, and mathematics.

FOR MORE INFO https://globalcomputing.group/

# INDUSTRY PARTNERSHIPS

AMD
Atos
CRAY THE SUPERCOMPUTER COMPANY
ENGILITY
EPAnalytics
IBM
IERUS
intel
MathWorks
nag
NVIDIA

# INTERNATIONAL COLLABORATIONS

***Barcelona Supercomputing Center***
BARCELONA, SPAIN

***Central Institute for Applied Mathematics***
JÜLICH, GERMANY

***Doshisha University***
KYOTO, JAPAN

***École Normale Supérieure de Lyon***
LYON, FRANCE

***École Polytechnique Fédérale de Lausanne***
LAUSANNE, SWITZERLAND

***ETH Zürich***
ZÜRICH, SWITZERLAND

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

***European Exascale Software Initiative***
EUROPEAN UNION

***Forschungszentrum Jülich***
JÜLICH, GERMANY

***High Performance Computing Center Stuttgart***
STUTTGART, GERMANY

***INRIA***
FRANCE

***Karlsruhe Institute of Technology***
KARLSRUHE, GERMANY

***Kasetsart University***
BANGKOK, THAILAND

***King Abdullah University of Science and Technology (KAUST)***
THUWAL, SAUDI ARABIA

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

***Moscow State University***
MOSCOW, RUSSIA

***National Institute of Advanced Industrial Science and Technology***
TSUKUBA, JAPAN

***Parallel and HPC Application Software Exchange***
TSUKUBA, JAPAN

***Prometeus GmbH***
MANNHEIM, GERMANY

***Regionales Rechenzentrum Erlangen (RRZE)***
ERLANGEN, GERMANY

***RIKEN***
WAKŌ, JAPAN

***Rutherford Appleton Laboratory***
OXFORD, ENGLAND

***Soongsil University***
SEOUL, SOUTH KOREA

***Technische Universität Wien***
VIENNA, AUSTRIA

***Technische Universität Dresden***
DRESDEN, GERMANY

***Tokyo Institute of Technology***
TOKYO, JAPAN

***Umeå University***
UMEÅ, SWEDEN

***Université Claude Bernard Lyon 1***
LYON, FRANCE

***University of Bordeaux***
BORDEAUX, FRANCE

***University of Cape Town***
CAPE TOWN, SOUTH AFRICA

***University of Manchester***
MANCHESTER, ENGLAND

***University of Paris-Sud***
PARIS, FRANCE

***University of Picardie Jules Verne***
AMIENS, FRANCE

***University of Tsukuba***
TSUKUBA, JAPAN

# PEOPLE

## STAFF AND STUDENTS

**Ahmad** ABDELFATTAH
*RESEARCH SCIENTIST II*

**Kadir** AKBUDAK
*RESEARCH SCIENTIST I*

**Robert** ANDERSON
*COMMUNICATIONS PROFESSIONAL II*

**Hartwig** ANZT
*DIRECTOR OF ICL /*
*MATHWORKS ASSOCIATE PROFESSOR, EECS*

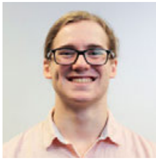

**Daniel** BARRY
*GRADUATE RESEARCH ASSISTANT*

**Natalie** BEAMS
*RESEARCH SCIENTIST I*

**Daniel** BIELICH
*POST DOCTORAL RESEARCH ASSOCIATE*

**George** BOSILCA
*RESEARCH ASSISTANT PROFESSOR*

**Aurelien** BOUTEILLER
*RESEARCH ASSISTANT PROFESSOR*

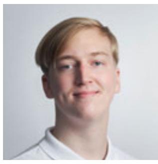

**Cade** BROWN
*UNDERGRADUATE RESEARCH ASSISTANT*

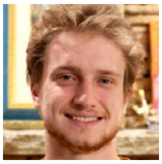

**Pierluigi**
CAMBIE-FABRIS
*BIWEEKLY STUDENT ASSISTANT*

**Sebastien** CAYROLS
*RESEARCH SCIENTIST I*

**Giuseppe** CONGIU
*RESEARCH SCIENTIST I*

**Anthony** DANALIS
*RESEARCH ASSISTANT PROFESSOR*

**Jack** DONGARRA
*RESEARCH PROFESSOR EMERITUS*

**Teresa** FINCHUM
*ADMINISTRATIVE SPECIALIST II*

**Mark** GATES
*RESEARCH ASSISTANT PROFESSOR*

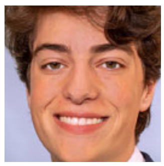

**Julian** HALLOY
*BIWEEKLY STUDENT ASSISTANT*

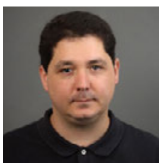

**Thomas** HERAULT
*RESEARCH ASSISTANT PROFESSOR*

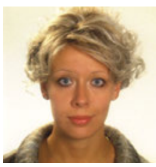

**Heike** JAGODE
*RESEARCH ASSISTANT PROFESSOR*

**Hannah** KNOCH
*INFORMATION SPECIALIST I*

**Julie** LANGOU
*RESEARCH LEADER*

**Tracy** LEE
*FINANCIAL SPECIALIST I*

**Jiali** LI
*GRADUATE RESEARCH ASSISTANT*

**Yicheng** LI
*GRADUATE RESEARCH ASSISTANT*

**Neil** LINDQUIST
*GRADUATE RESEARCH ASSISTANT*

**Piotr** LUSZCZEK
*RESEARCH ASSISTANT PROFESSOR*

**Ethan** MANASS
*UNDERGRADUATE*

**Noah** MARKS
*GRADUATE RESEARCH ASSISTANT*

**Maksim** MELNICHENKO
*GRADUATE RESEARCH ASSISTANT*

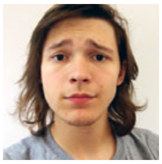

**Daniel** MISHLER
*GRADUATE RESEARCH ASSISTANT*

**Phuong** NGUYEN
*RESEARCH ASSOCIATE II*

**Anustuv** PAL
*RESEARCH SCIENTIST I*

**Deborah** PENCHOFF
*ASSOCIATE DIRECTOR*

**Gerald** RAGGHIANTI
*TECHNICAL SERVICES LEADER*

**Yves** ROBERT
*VISITING SCHOLAR*

**David** ROGERS
*COMMUNICATIONS SERVICES LEADER*

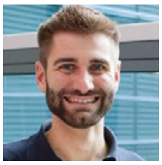

**Joseph** SCHUCHART
*POST DOCTORAL RESEARCH ASSOCIATE*

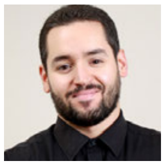

**Wissam** SID LAKHDAR
*RESEARCH SCIENTIST I*

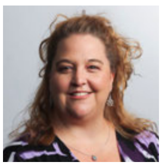

**Leighanne** SISK
*ADMINISTRATIVE SPECIALIST I*

**Joan** SNODERLY
*ASSISTANT DIRECTOR*

**Dalal** SUKKARI
*RESEARCH SCIENTIST I*

**Stanimire** TOMOV
*RESEARCH ASSISTANT PROFESSOR*

**Asim** YARKHAN
*RESEARCH SCIENTIST II*

# PEOPLE

## VISITORS

IN ADDITION to ICL's staff and students, ICL's offices are typically populated with visitors throughout the year. ICL often hosts visitors from groups who are collaborating on research projects, or from groups affiliated with ICL members. During the summer of 2022, ICL hosted a visiting group from KIT's Fixed-Point Methods for Numerics at Exascale research group. KIT's Fritz Goebel composed a recap of his group's visit.


LEFT (left to right) **Mike Tsai, Thomas Grützmacher, Marcel Koch, Hartwig Anzt, Vasilis Georgiou, Fritz Goebel, Claudius Holeksa**
BELOW **The same group joined by Gerald Ragghianti**

*After having to cancel our planned visit in 2020 due to Covid, we, a group of 5 Ph.D. students (Claudius, Fritz, Mike, Thomas, and Vasilis) and one postdoc (Marcel) from the Karlsruhe Institute of Technology in Germany, finally got the chance to visit ICL for a few weeks this August.*

*Hartwig had just left Karlsruhe to take on his new position as ICL's director a few days before, and we were looking forward to a quick reunion in the States. After each taking a different trip to Knoxville we moved into an Airbnb on South Gay St, our base for the upcoming weeks. Our stay was off to a great start with a first day at ICL with a warm welcome and a first dinner downtown before we got to spend a few great days at the retreat in Gatlinburg filled with hiking and lots of interesting talks and discussions. For the remainder of our stay, various additional hikes, shenanigans like whitewater rafting and mountain biking, lunches and dinners with many friendly ICLers as well as early morning swim and gym sessions with Hartwig made for a great time.*

*We had a great time finally meeting a lot of people we so far only knew as little windows in ECP calls in person. Thank you very much to everyone at ICL for having us and making us feel like part of the family, some of us are already planning to come back for a while in 2023. A special thank you to Joan, Leighanne, and Teresa for organizing everything from on-site parking to pool and gym access for us!*

# SEMINAR SERIES

FRIDAY LUNCH TALKS are a long-standing tradition at ICL. They serve as a weekly opportunity for the group to connect. The speakers scheduled are a combination of ICL's research staff and students and guests. For ICL members, these talks are a chance to present their work to the group as a way to keep everyone informed and solicit feedback and ideas. Similarly, for visiting speakers, they offer an opportunity for new collaborations.

## 2022 ICL SEMINAR SERIES **Visiting Speakers**

FEB 25 **Laura Grigori**
INRIA PARIS

**Recent advances in randomization techniques for solving linear systems of equations**

MAR 4 **Ashlee Anderson**
UNIVERSITY OF TENNESSEE
TPTE DEPARTMENT

**Just what is critical race theory, and why is everyone talking about it?**

MAR 11 **Nigel Tan**
UNIVERSITY OF TENNESSEE
GLOBAL COMPUTING LAB

**Towards Access Pattern Aware Checkpointing For Kokkos Applications**

MAR 25 **Eddie Mitchell**
UNIVERSITY OF TENNESSEE
DEPARTMENT OF MATHEMATICS

**Simplicial Convolutional Neural Networks for Neural Spike Train Decoding**

APR 1 **Cleve Moler**
MATHWORKS

**Cleve's Corner Blog, April 1; Past and Present**

APR 8 **Sherry Li**
LAWRENCE BERKELEY NATIONAL LABORATORY

**Interplay of linear algebra, machine learning, and HPC**

APR 22 **Zhaojun Bai**
UC DAVIS

**Meeting the challenges of computing many eigenpairs**

APR 22 **David Keyes**
KAUST

**Meeting the challenges of computing many eigenpairs**

MAY 6 **James Demmel**
UNIVERSITY OF CALIFORNIA AT BERKELEY

**Communication-Avoiding Algorithms for Linear Algebra, ML and Beyond**

MAY 13 **Erin Carson**
CHARLES UNIVERSITY, CZECH REPUBLIC

**Challenges and Opportunities in Mixed Precision Numerical Linear Algebra**

MAY 20 **Swann Perarnau**
ARGONNE NATIONAL LABORATORY

**AML: Building Blocks for Advanced Memory Management on Heterogeneous Architectures**

AUG 19 **Swann Perarnau**
KARLSRUHE INSTITUTE OF TECHNOLOGY

**Using POSIT as an element-wise compression technique on GPUs**

AUG 26 **Catherine Schuman**
UNIVERSITY OF TENNESSEE
EECS DEPARTMENT

**Neuromorphic Computing from the Computer Science Perspective: Algorithms and Applications**

SEP 2 **Ulrike Yang**
CENTER FOR APPLIED SCIENTIFIC COMPUTING AT LLNL

**On the Design of Algebraic Multigrid Methods for Exascale Computers and their Use in Numerical Simulations**

SEP 9 **Ilse Ipsen**
NORTH CAROLINA STATE UNIVERSITY

**Why matrices can be better conditioned in lower precision**

SEP 23 **Jamie Finney**
ORNL

**Facility Software Deployment and User Environments**

SEP 30 **Edgar Solomonik**
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

**Scaling Numerical Algorithms and Software via Improved Performance Modeling**

OCT 7 **Aditya Kashi**
ORNL

**High-performance batched sparse iterative linear solvers and preconditioners for GPUs**

OCT 21 **Joshua Fu**
UNIVERSITY OF TENNESSEE / ORNL

**Chemistry Solver in the Energy Exascale Earth System Model**

NOV 11 **Hatem Ltaief**
KAUST

**Redesigning Matrix Computations for Computational Astronomy and Seismic Imaging Applications**

DEC 16 **Guojing Cong**
KAUST

**Distributed Training and Intelligent Simulation for Accelerated Discovery**

# PEOPLE

## *ALUMNI*

The success of ICL over the years was possible due to the efforts of many talented staff and students.

**Maksims** ABALENKOVS
**Carolyn** AEBISCHER
**Sudesh** AGRAWAL
**Bivek** AGRAWAL
**Emmanuel** AGULLO
**Mohammed** AL FARHAN
**Rabab** AL-OMAIRY
**Jennifer** ALLGEYER
**Wes** ALVARO
**Ed** 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
**Laura** BLACK
**Noel** BLACK
**Susan** BLACKFORD
**Wesley** BLAND
**Kartheek** BODANKI
**David** BOLT
**Fernando** BOND
**Carolyn** BOWERS
**Vincent** BRIDONNEAU
**Barry** BRITT
**Randy** BROWN
**Murray** BROWNE
**Cynthia** BROWNE
**Bonnie** 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
**Ying** DING
**Jin** 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
**Mike** FINGER
**Jamie** FINNEY
**Markus** FISCHER
**Clement** FLINT
**Len** FREEMAN
**Xiaoquan** FU
**Erika** FUENTES
**Karl** FUERLINGER
**Megan** FULLER
**Edgar** GABRIEL
**Lynn** GANGWER
**Tracy** GANGWER
**Kelley** GARNER
**Nathan** 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
**Hanumantharayappa**
**Hidehiko** HASEGAWA
**Satomi** HASEGAWA
**Chris** HASTINGS
**Blake** HAUGEN
**David** HENDERSON
**John** HENRY
**Greg** HENRY
**Julien** HERRMANN
**Holly** HICKS
**Alexandra** HICKS-HARDIMAN
**Sid** HILL
**Tomoyuki** HIROYASU
**George** HO
**Josh** HOFFMAN
**Reazul** HOQUE
**Jeff** HORNER
**Mitch** HORTON
**Yan** HUANG
**Harry** HUGHES
**Aurelie** HURAULT
**Chris** HURT
**Paul** JACOBS
**Emmanuel** JEANNOT
**Weizhong** JI
**Yulu** JIA
**Weicheng** JIANG
**Song** JIN
**Patrick** JOHANSSON
**Aral** JOHNSON
**Matt** JOHNSON
**Sean** JOLLY

**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
**Myung** HO KIM
**Youngbae** KIM
**Jenya** KIRSHTEIN
**Cindy** KNISLEY
**Michael** KOLATIS
**Anara** KOZHOKANOVA
**Moritz** KREUTZER
**Chandra** KRINTZ
**Tilman** KUESTNER
**Jakub** KURZAK
**Krerkchai** KUSOLCHU
**Coire** KYLE
**Amanda** LAAKE
**Xavier** LACOSTE
**Julien** LANGOU
**Jeff** LARKIN
**Brian** LAROSE
**Frank** LAUER
**Valentin** LE FEVRE
**DongWoo** LEE
**Pierre** LEMARINIER
**Todd** LETSCHE
**Sharon** LEWIS
**Yinan** LI
**Xiang** LI
**Weiran** LI
**Chaoyang** LIU
**Kevin** LONDON
**Matt** LONGLEY
**Florent** LOPEZ
**Florent** LOPEZ
**Nuria** LOSADA
**Hatem** LTAIEF
**Yuechao** LU
**Daniel** LUCIO
**Richard** LUCZAK
**Xi** LUO
**Teng** MA
**Sticks** MABAKANE
**Robert** MANCHEK
**Gabriel** MARIN
**Tushti** MARWAH
**Theo** MARY
**Ian** MASLIAH
**Donald** MCCASLAND
**Tyler** MCDANIEL
**Paul** MCMAHAN
**Eric** MEEK
**James** MEYERING
**Jeremy** MILLAR
**Michelle** MILLER
**Cindy** MITCHELL
**Stuart** MONTY
**Shirley** MOORE
**Keith** MOORE
**Terry** MOORE
**Erik** 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
**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
**Matt** SKINNER
**Peter** SOENDERGAARD
**Raffaele** SOLCA
**Gwang** SON
**Fengguang** SONG
**Thomas** SPENCER
**Jeffrey** STEILL
**Erich** STROHMAIER
**Xiaobai** SUN
**Martin** SWANY
**Daisuke** TAKAHASHI
**Judi** TALLEY
**Ronald** TAM
**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
**Mike** WALTZ
**Robert** 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
**Panruo** WU
**Wei** WU
**Jiayi** WU
**Qiu** XIA
**Tinghua** XU
**Ichitaro** YAMAZAKI
**Tao** YANG
**Erlin** YAO
**Kevin** YE
**Jin** YI
**Haihang** YOU
**Lamia** YOUSEFF
**Brian** ZACHARY
**Dmitry** ZAITSEV
**Omar** ZENATI
**Yuanlei** ZHANG
**Junlong** ZHAO
**Yong** ZHENG
**Dong** ZHONG
**Luke** ZHOU
**Min** ZHOU

# ACKNOWLEDGMENTS

Over 33 years, ICL's work has been made possible through the support of the federal agencies that are charged with allocating public research funding, and we thank the following agencies:

In addition to the support of the federal government, we have 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:

Administrative and technical support for ICL is provided by the Center for Information Technology Research (CITR).

INNOVATIVE
COMPUTING LABORATORY

FIND OUT MORE AT
https://icl.utk.edu/