ICL
INNOVATIVE
COMPUTING LABORATORY

2015/16 REPORT

THE UNIVERSITY OF TENNESSEE
KNOXVILLE

IN MEMORIAM

Nathan Garner

1973-2015

ICL 2015/16 REPORT

EDITED BY **Sam Crawford** DESIGNED BY **David Rogers**

ICL 2015/16 REPORT

FROM THE DIRECTOR

Jack Dongarra
DIRECTOR, ICL

With any luck, this past summer's announcement of President Obama's National Strategic Computing Initiative (NSCI) should usher in a national environment for scientific research that will help the Innovative Computing Laboratory continue to thrive. Aspiring to "… create systems that can apply exaflops of computing power to exabytes of data," the NSCI proposes to establish a coordinated, long term, multi-agency strategy for improving the nation's economic competitiveness and research prowess by raising its high performance computing and data analysis capabilities to unprecedented heights. I remember very well the last time—more than 15 years ago—such an ambitious federal initiative was launched, because it was my long time friend and collaborator, the late Ken Kennedy, who led the President's Information Technology Advisory Committee (PITAC) that produced the Information Technology Research: Investing in Our Future report. If the NSCI generates, over time, the same kind of national research environment that Ken's PITAC report did, then the future prospects for ICL will indeed be bright.

Of course my optimism about our ability to compete in such an opportunity space is due largely to the quality of the team we have and to the record of accomplishment that this team, including all of our researchers, staff, and students, has achieved. In fact it has become something of a tradition, over the last few years, to include in the annual report a snapshot of this remarkable history in the form of a timeline that lines up the tenure of some exceptional ICL alumni with the lifespan of major ICL projects. The names of the projects that they have helped to lead over the years (e.g., PVM, LAPACK, ScaLAPACK, BLAS, ATLAS, Netlib, TOP500, PAPI, NetSolve, Open-MPI, HPC Challenge, HPCG benchmark, PaRSEC, MAGMA, and many others) are familiar to HPC users around the world. In terms of our reputation for producing both world-class research papers, and world-class software infrastructure for scientific computing, ICL has very few peers around the world.

I am certainly confident that we made another set of important additions to ICL's stellar publication record in 2015. If you look at the two full pages of publications listed later in this report, you will find that papers from all of our core research areas—numerical libraries, distributed computing, performance optimization, and performance monitoring and benchmarking—were published in various major journals, or at well known, high quality conferences. The paper on “Exascale Computing and Big Data” that was published with Dan Reed in the Communications of the ACM had an especially broad impact because it explicates, in a clear way, the problem that the NSCI puts squarely on the table, namely, the problem of how to integrate traditional HPC with new forms of high-end data analytics. But just as notable, from an ICL point of view, is the fact that we published papers at SC15 and the IEEE Big Data conference that show how this integration problem can be successfully attacked through targeted innovations in algorithms and numerical libraries.

ICL also had a very strong year in terms of contending for funding for creative new projects. This is particularly apparent in the case of new NSF awards, where the competition is fierce and the peer review process is arduous. As with this year’s publications, we received NSF funding for new projects in each of our major research areas: In numerical libraries, the SparseKaffe project (joint venture with Sanjay Ranka at the Univ. of Florida and Tim Davis at Texas A&M) aims to achieve order of magnitude performance gains (with energy efficiency) for sparse direct methods on multicore platforms with accelerators, at all scales. In distributed computing, the project on Task-based Environment for Scientific Simulation at Extreme Scale (TESSE) is a collaboration with Robert Harrison at SUNY Stony Brook and Ed Valeev at Virginia Tech that aims to create a general-purpose, production-quality software framework for achieving both programmer productivity and portable performance on massively-parallel and hybrid platforms. In performance optimization, the Data-driven Autotuning for Runtime Execution (DARE) project breaks new ground in empirical autotuning techniques, in which ICL has pioneered, applying them to distributed computing platforms, and thereby providing unprecedented application-level performance tuning capabilities to end users. Finally, in performance monitoring and measurement, the PAPI-EX project (joint venture with Vince Weaver from the Univ. of Maine) will extend ICL’s well known and widely used PAPI package in ways that enable the monitoring of performance-critical resources that are shared by multicore and hybrid processors (e.g., on-chip communication networks, memory hierarchy, I/O interfaces, and power management logic), and that support PAPI integration with new task-based runtime systems.

Thus, the foundation has been laid for great things in all the domains in which ICL has excelled over the past quarter century. If, as I hope, the NSCI begins to receive the kind of support that many of us believe it should, the coming year will certainly see ICL’s research efforts continue to flourish and our engagement with the broader scientific community continue to grow. During these exciting times, I am grateful to our sponsors for their ongoing support. And as always, my special thanks and congratulations go to ICL researchers, staff, and students for their skill, dedication, and tireless efforts to make the ICL one of the best centers in the world for scientific computing technologies.

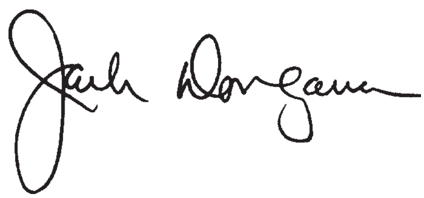

INTRODUCTION

The Innovative Computing Laboratory (ICL) is a large Computer Science research and development group situated in the heart of the University of Tennessee's Knoxville campus. ICL's mission is to establish and maintain the University of Tennessee as a world leader in advanced scientific and high performance computing through research, education, and collaboration.

KEY RESEARCH

- 1989: Level 3 BLAS, PVM
- 1992: BLACS, LAPACK
- 1993: TOP500
- 1994: MPI
- 1995: ScaLAPACK
- 1996: DOD PET
- 1997: ATLAS, NetSolve, RIB
- 1999: HARNESS, PAPI
- 2000: HPL

ICL MILESTONES

ICL ESTABLISHED IN AYRES HALL

PVM R&D100

ATLAS & NETSOLVE R&D100

MOVE TO CLAXTON BUILDING

1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999 2000 2001

PAPI R&D100

SELECT ALUMNI

1989

- 1994-2004 **Phil Mucci**
- 1994 **Frederic Desprez** — ENS-LYON, FRANCE
- 1996-1997 **Yves Robert** — ENS-LYON, FRANCE
- 1989-2001 **Susan Blackford** — CSP, INC.
- 1991-1992 **Adam Beguelin** — SENSR.NET
- 1994-2009 **Keith Seymour**
- 1992-2005 **Victor Eijkhout** — UNIVERSITY OF TEXAS, AUSTIN
- 1992-1995 **Jaeyoung Choi** — SOONGSIL UNIVERSITY, KOREA
- 1996-2001 **Dorian Arnold** — UNIVERSITY OF NEW MEXICO
- 1993-2001 **Antoine Petitet** — ESI GROUP, FRANCE
- 1989-1991 **Ed Anderson** — EPA
- 1993-1994 **Bernard Tourancheau** — UNIVERSITÉ JOSEPH FOURIER DE GRENOBLE
- 1996-2006 **Graham Fagg** — MICROSOFT
- 1991-1994 **Roldan Pozo** — NIST
- 1995-2001 **Erich Strohmaier** — LAWRENCE BERKELEY NATIONAL LABORATORY
- 1989-1990 **Zhaojun Bai** — UNIVERSITY OF CALIFORNIA, DAVIS
- 1992-1994 **Richard Barrett** — SANDIA NATIONAL LABORATORIES
- 1995-1998 **Henri Casanova** — UNIVERSITY OF HAWAII, MANOA
- 1999-2004 **Ken Roche** — UNIVERSITY OF WASHINGTON
- 1995-1996 **Greg Henry** — INTEL
- 1999-2003 **Sathish Vadhiyar** — INDIAN INSTITUTE OF SCIENCE, INDIA
- 1990-1996 **Bob Manchek** — STRATUS TECHNOLOGIES
- 1999-2010 **David Cronk** — LOCKHEED-MARTIN
- 1990-1991 **Robert van de Geijn** — UNIVERSITY OF TEXAS, AUSTIN
- 1995-1996 **Andy Cleary** — AMAZON
- 1995-1996 **Sven Hammarling** — NUMERICAL ALGORITHMS GROUP
- 1996-2005 **Kevin London** — MICROSOFT
- 1996-1999 **Martin Swany** — INDIANA UNIVERSITY
- 1991-2001 **Clint Whaley** — LOUISIANA STATE UNIVERSITY
- 1990-1996 **Reed Wade** — WETA DIGITAL
- 1997 **Françoise Tisseur** — MANCHESTER UNIVERSITY, ENGLAND
- 1993-2012 **Shirley Moore** — UNIVERSITY OF TEXAS AT EL-PASO
- 1989-2007 **Keith Moore**

ICL's founder, Dr. Jack Dongarra, established the lab in 1989 when he received a dual appointment as a Distinguished Professor at UTK and as a Distinguished Scientist at Oak Ridge National Laboratory. Since then, ICL has grown into an internationally recognized research laboratory, specializing in Numerical Linear Algebra, Distributed Computing, and Performance Evaluation and Benchmarking. The lab now employs forty researchers, students, and staff, and has earned many accolades, including four R&D100 awards.

2002	2003	2006	2008	2009	2010	2011	2012	2013	2014	2015
FT-MPI GCO	HPC Challenge LAPACK for Clusters	FT-LA Open MPI	MAGMA PLASMA	Blackjack IESP MuMMI	DPLASMA	PULSAR	PaRSEC ULFM	BDEC BEAST HPCG	ARGO RaPyDLI	DARE PAPIEX SparseKaffe TESSE

25YEARS
OF ENABLING TECHNOLOGIES FOR SCIENTIFIC COMPUTING

2002 2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014 2015

2003-2004 **Edgar Gabriel**
UNIVERSITY OF HOUSTON

2006-2008 **Karl Fuerlinger**
LUDWIG-MAXIMILIANS-UNIVERSITY MUNICH

2010-2012 **Mathieu Faverge**
UNIVERSITY OF BORDEAUX

2003-2007 **Thara Angskun**
SURANAREE UNIVERSITY OF TECHNOLOGY

2002-2009 **Haihang You**
CHINESE ACADEMY OF SCIENCES

2003-2006 **Julien Langou**
UNIVERSITY OF COLORADO, DENVER

2003-2006 **Felix Wolf**
JÜLICH SUPERCOMPUTING CENTRE

2009 **Emmanuel Agullo**
INRIA, FRANCE

2001-2006 **Zizhong (Jeffrey) Chen**
UNIVERSITY OF CALIFORNIA, RIVERSIDE

2008-2011 **Hatem Ltaief**
KAUST, SAUDI ARABIA

2004-2007 **Alfredo Buttari**
CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE

2001-2006 **Nathan Garner**

2010-2012 **Vince Weaver**
UNIVERSITY OF MAINE

2001-2014 **Dan Terpstra**
LIVING WATERS FOR THE WORLD

2015 HIGHLIGHTS

A MILESTONE

OVER 25 YEARS OF
ENABLING TECHNOLOGIES
FOR SCIENTIFIC COMPUTING

The 2014–2015 academic year marked the 25th anniversary of ICL's inception. ICL has been conducting innovative research in high performance computing and cyberinfrastructure since 1989, while making significant contributions to the HPC community along the way. What started as two graduate students and two Post-docs has now grown to a fully functional research center with a staff of 40 researchers, students, and administrators.

This quarter-century milestone was celebrated with a "25 Years of Innovative Computing" workshop on March 31 – April 2, 2015. Held at the University of Tennessee Conference Center, the workshop included 45 talks presented by ICL alumni over the span of 2 days. Around 90 total participants came from all over the globe—including attendees from as far away as Japan and Saudi Arabia—to share their current interests and activities and to celebrate their time at ICL.

ACCOLADES

Intel Parallel Computing Center

The Intel Parallel Computing Center (IPCC) program, renewed with ICL in 2015, is comprised of universities, institutions, and labs that are leaders in their field, focusing on modernizing applications to increase parallelism and scalability through optimizations that leverage cores, caches, threads, and vector capabilities of microprocessors and coprocessors.

The objective of the Innovative Computing Laboratory's IPCC is the development and optimization of numerical linear algebra libraries and technologies for applications, while tackling current challenges in heterogeneous Intel® Xeon Phi™ coprocessor-based High Performance Computing. In collaboration with Intel's MKL team, the IPCC at ICL will modernize the popular LAPACK and ScaLAPACK libraries to run efficiently on current and future manycore architectures, and will disseminate the developments through the open source MAGMA MIC library.

Jack Dongarra Receives HPCwire Readers' Choice Award for Outstanding Leadership in HPC

At SC15 in Austin, TX, ICL's Jack Dongarra received the HPCwire Readers' Choice Award for Outstanding Leadership in HPC. These awards, which are nominated and voted on by the worldwide HPC community, recognize the best and brightest developments in HPC in the last 12 months.

Best Paper Awards

HPEC'15

On September 15th, ICL's Azzam Haidar and his co-authors were presented with a best paper award for their entry in the 2015 IEEE High Performance Extreme Computing Conference (HPEC'15). The paper, "MAGMA Embedded: Towards a Dense Linear Algebra Library for Energy Efficient Extreme Computing," proposes the design and implementation of embedded system aware algorithms that target the challenges of lower power consumption and small form factors found in embedded computing devices like smartphones and watches.

Haidar, A., S. Tomov, P. Luszczek, and J. Dongarra, "**MAGMA Embedded: Towards a Dense Linear Algebra Library for Energy Efficient Extreme Computing**," *2015 IEEE High Performance Extreme Computing Conference (HPEC '15)*, Waltham, MA, IEEE, September 2015.

SpringSim'15

Khairul Kabir and his co-authors earned a best paper award for their entry in the Spring Simulation Multi-Conference 2015 (SpringSim'15) on April 12th. The paper, "Performance Analysis and Design of a Hessenberg Reduction using Stabilized Blocked Elementary Transformations for New Architectures," revisits solving nonsymmetric eigenvalue problems by using stabilized elementary transformations in the context of new architectures like multi-core CPUs and Intel Xeon Phi coprocessors.

Kabir, K., A. Haidar, S. Tomov, and J. Dongarra, "**Performance Analysis and Design of a Hessenberg Reduction using Stabilized Blocked Elementary Transformations for New Architectures**," *The Spring Simulation Multi-Conference 2015 (SpringSim'15)*, Alexandria, VA, April 2015.

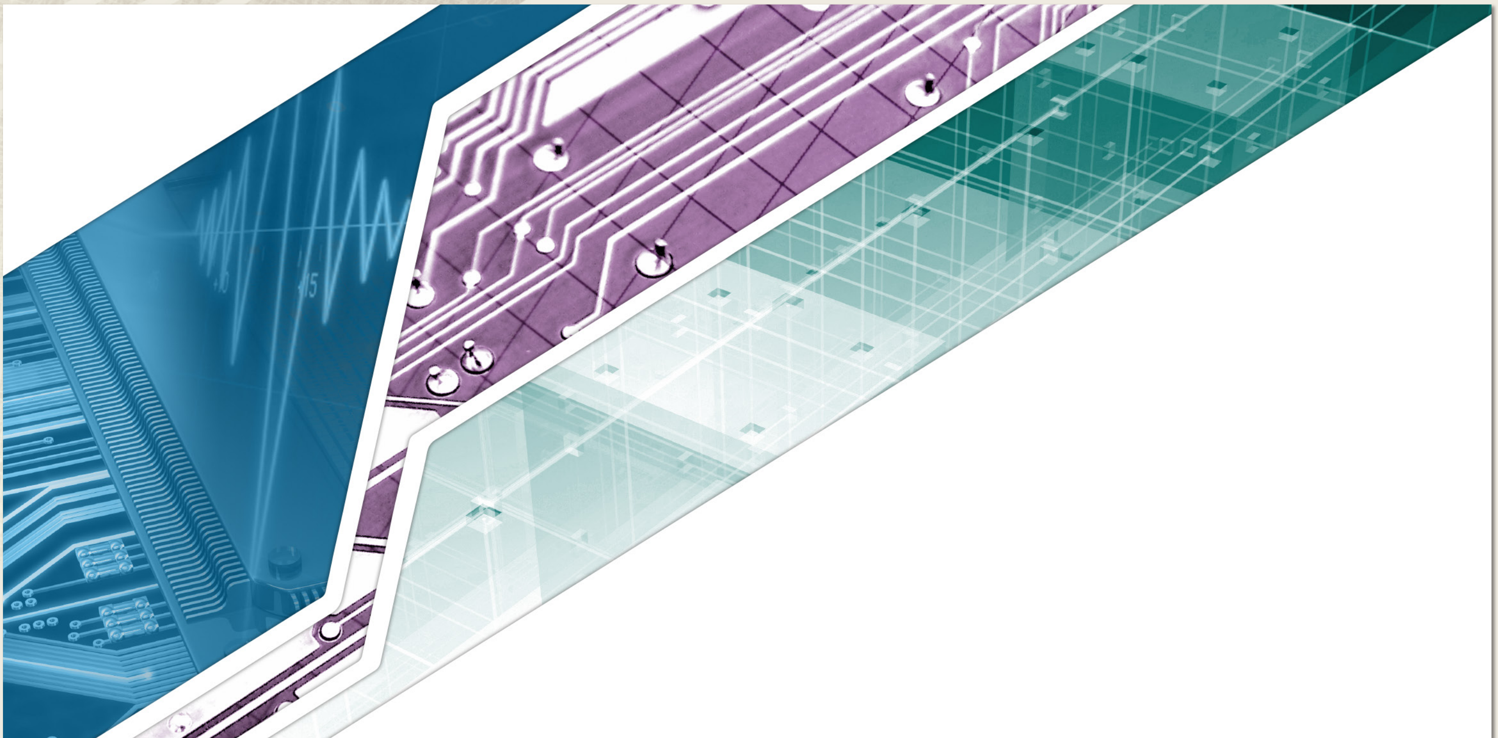

RESEARCH

Increased efforts to keep pace with the evolution in HPC hardware and software present unique challenges that only a handful of enabling technology researchers are capable of addressing successfully. Our cutting-edge research efforts of the past have provided the foundation for addressing these challenges and serve as catalysts for success in our ever growing research portfolio. Our vision, our expertise, our determination, and our track record continue to position ICL as a leader in academic research.

What originally began over 25 years ago as in-depth investigations of the numerical libraries that encode the use of linear algebra in software, has grown into an extensive research portfolio. We have evolved and expanded our research agenda to accommodate the aforementioned evolution in HPC, which includes a focus on algorithms and libraries for multicore and hybrid computing. As we have gained a solid understanding of the challenges presented in these domains, we have further expanded our scope to include work in performance evaluation and benchmarking for high-end computers, as well as work in high performance parallel and distributed computing, with efforts focused on message passing and fault tolerance.

Demonstrating the range and diversity of our work, we will be engaged in more than 25 significant research projects during 2015-2016 across our main areas of focus. On the following pages, we provide brief summaries of some of our efforts in these areas. For more detailed information about our work, visit our website – http://icl.utk.edu/.

NUMERICAL LINEAR ALGEBRA

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

PERFORMANCE EVALUATION AND BENCHMARKING

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

DISTRIBUTED COMPUTING

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

ADAPT

Next-Gen Open MPI

The ADAPT project proposes to enhance, harden, and modernize the Open MPI library in the context of the ongoing revolution in processor architecture and system design. On the large systems expected before the end of this decade, the degree of parallelism (intra and inter node) will presumably increase by several orders of magnitude (based on the Exascale Roadmap predictions). To efficiently handle such systems, MPI implementations will have to adopt more asynchronous and thread-friendly behaviors to extract the best performance from more complex architectures.

The project team seeks to create a viable foundation for a new generation of Open MPI components, which enables a rapid exploration of new physical capabilities, provides greatly improved performance portability, and works toward full interoperability between classes of components. ADAPT explores process placement, distributed topologies, file accesses, point-to-point and collective communications, and different approaches to fault tolerance.

ARGO

A new architecture for Exascale computing

The ARGO project is developing a new Exascale Operating System and Runtime (OS/R) designed to support extreme-scale scientific computation. Disruptive new computing technologies, such as 3D memory, ultra-low-power cores, and embedded network controllers, are changing the scientific computing landscape. As it becomes clear that incremental approaches to operating systems and runtimes (OS/R) cannot grow into an Exascale solution, we propose a novel radical approach.

ARGO is designed on a new, agile modular architecture that supports both global optimization and local control. It aims to efficiently leverage new chip and interconnect technologies while addressing the new modalities, programming environments, and workflows expected at Exascale. It is designed from the ground up to run future HPC applications at extreme scales.

BEAST

Bench-testing Environment for Automated Software Tuning

FIND OUT MORE AT
http://icl.utk.edu/beast/

The objective of the Bench-testing Environment for Automated Software Tuning (BEAST) project is to embrace the nature of accelerators, such as Graphics Processing Units (GPUs) from NVIDIA and AMD, and Xeon Phi coprocessors from Intel, which offer an order of magnitude more computing power and an order of magnitude more memory bandwidth than standard processors, and create an unprecedented opportunity for breakthroughs in science and technology.

Accelerators are a different kind of beast when it comes to performance tuning, and their architectural features usually pose unique programming challenges. For example, accelerators have massive numbers of simple cores, with static pipelines, and no branch prediction, and a multitude of constraints that can obliterate performance in numerous situations. BEAST allows a user to write high performance kernels in a tunable manner, sweep through a large search space, collect massive amounts of performance data, and plow through that data with machine learning techniques. These features enable a user to optimize code for extreme performance, without descending into the dark abyss of assembly programming.

CoDAASH

Co-design Approach for Advances in Software and Hardware

The Co-design Approach for Advances in Software and Hardware (CoDAASH) project focuses on understanding the relationship between algorithms and hardware platforms and how to jointly optimize the software and hardware in order to achieve efficient implementations for applications in materials science, chemistry, and physics. CoDAASH is a joint effort between the University of Tennessee, Knoxville, Iowa State University, University of Texas, El Paso, and the University of California, San Diego, and is funded by the United States Air Force Office of Scientific Research (AFOSR).

ICL's contribution focuses on expressing certain computational chemistry algorithms in the form of a data flow graph (DAG) 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 enables computation at extreme scale in the era of many-core and highly heterogeneous platforms, by utilizing the hardware components (e.g., CPU or GPU) that perform best for the type of computational task under consideration. The dataflow-based form of these algorithms makes them compatible with next generation task scheduling systems, such as PaRSEC.

DARE

Data-driven Autotuning for Runtime Execution

The objective of the Data-driven Autotuning for Runtime Execution (DARE) project is to provide application-level performance tuning capabilities to the end user. DARE's development motivation stems from the never-ending hurdles of performance tuning of the PLASMA and MAGMA linear algebra libraries. These hurdles motivated the development of a software architecture that combines three components: hardware analysis, kernel modeling, and workload simulation.

With DARE, the hardware analysis block builds a detailed model of the hardware, its computational resources (CPU cores, GPU accelerators, Xeon Phi coprocessors), and its memory system (host memories, device memories, multiple levels of cache). The kernel modeling block builds accurate performance models for the computational kernels involved in the workload, depending on granularity, place of execution, induced memory traffic, etc., and the workload simulation block rapidly simulates a large number of runs in order to find the best execution conditions, while relying on the information provided by the other two blocks. The ultimate objective of DARE is to arrange the blocks in a continuous refinement loop that can serve as a framework for optimizing applications beyond the field of dense linear algebra.

DPLASMA

Distributed Parallel Linear Algebra Software for Multicore Architectures

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

The Distributed Parallel Linear Algebra Software for Multicore 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 multicore processors, and if available, accelerators like GPUs or Intel Xeon Phi. DPLASMA achieves this objective by deploying the Parallel Linear Algebra Software for Multicore Architectures (PLASMA) algorithms on distributed memory systems by leveraging 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 transparently operate on local data, or introduce implicit communications to resolve dependencies, removing the burden of initial data re-shuffle, and providing to the user a novel approach to address load balance.

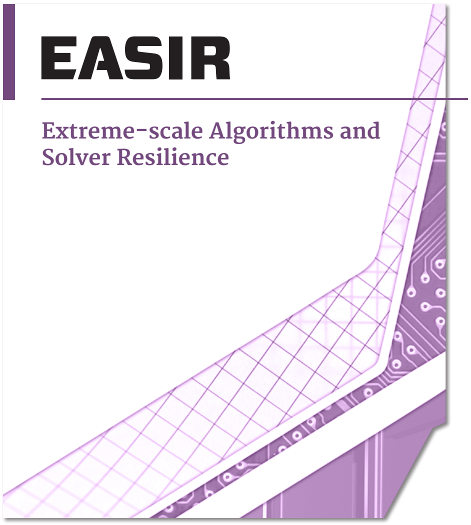

The mission of the Extreme-scale Algorithms and Solver Resilience (EASIR) project is to close the performance gap between the peak capabilities of HPC hardware and the performance realized by high performance computing applications. To carry out this mission, the EASIR project team develops architecture-aware algorithms and libraries, and the supporting runtime capabilities, to achieve scalable performance and resilience on heterogeneous architectures.

The project team includes personnel from ORNL, Sandia National Laboratories, the University of Illinois, the University of California Berkeley, and the University of Tennessee (ICL). ICL's efforts focus on providing components and services in a vertically integrated software stack, from low-level runtime process and thread scheduling to multicore aware library interfaces, multicore dense linear algebra, scalable iterative methods, and advanced parallel algorithms that break traditional parallelism bottlenecks.

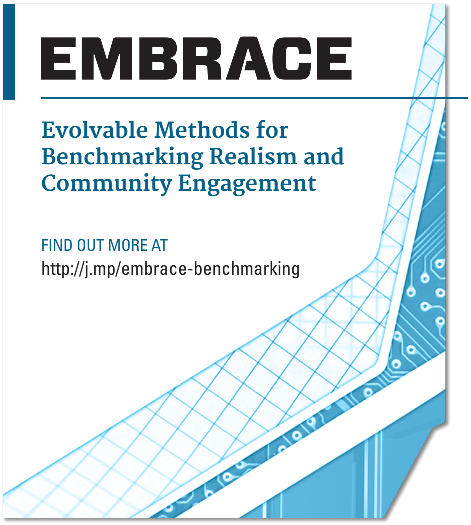

The primary goal of the Evolvable Methods for Benchmarking Realism and Community Engagement (EMBRACE) project is to find the most rational way to design and develop forward-looking benchmarks followed by correctly interpreting their results. This concept has become difficult over time because the performance engineering side of the scientific computing community is broadening, its applications are diverse, and the hardware platforms are quickly evolving. The aim of the EMBRACE project then is to facilitate a technical and reasoned debate about benchmarking, and to keep this debate alive over time.

EMBRACE is a collaboration between ICL/UTK and Georgia Tech. ICL's contribution includes expertise in benchmark design and implementation, and community engagement. In November 2015, an EMBRACE Birds-of-a-Feather session was held at the SC15 conference in Austin, TX. The reception was very positive and the need for reimagining the benchmarking process, creation, and interpretation was clearly voiced by the attendees. A workshop for SC16 is in the planning stages and was met with enthusiastic response as a forum for discussion about the community's need for performance benchmarking.

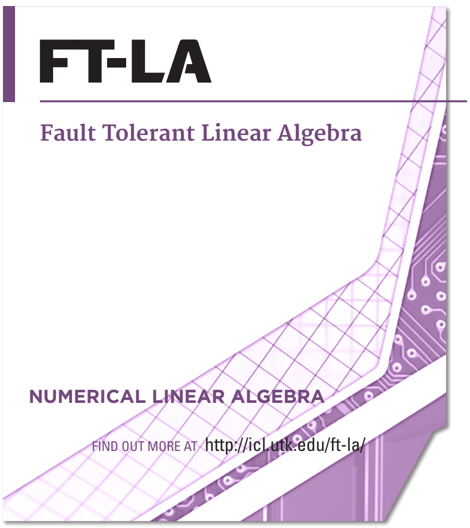

The Fault Tolerant Linear Algebra (FT-LA) research effort is aimed at understanding and developing Algorithm Based Fault Tolerance (ABFT) into major dense linear algebra kernels. With parallel machines currently reaching up to 300,000 cores, fault-tolerance has never been so paramount. The scientific community has to tackle process failures from two directions: first, efficient middleware needs to be designed to detect failures, and second, the numerical applications have to be flexible enough to permit the recovery of the lost data structures.

At ICL, we have successfully developed Fault Tolerant MPI middleware and, more recently, an FT-LA library that will efficiently handle several process failures. The project now supports soft (bit-flips) and hard (process failure) failures on a significant range of dense linear algorithms, from one sided factorizations (Cholesky, LU and QR) to two-sided factorizations (Hessenberg, tri-diagonalization, and bi-diagonalization). Future work in this area involves the development of scalable fault-tolerant singular value decompositions and eigendecomposition, following the ABFT principles.

The High Performance Conjugate Gradients (HPCG) benchmark is designed to measure performance that is representative of modern scientific applications. It does so by exercising the computational and communication patterns that are commonly found in real science and engineering codes, which are often based on sparse iterative solvers. HPCG exhibits the same irregular accesses to memory and fine-grain recursive computations that dominate large-scale scientific workloads used to simulate complex physical phenomena. As an emerging HPC metric of choice, HPCG implements the preconditioned conjugate gradient algorithm with a local symmetric Gauss-Seidel as the preconditioner. Additionally, the essential components of the geometric multigrid method are present in the code as a way to represent execution patterns of modern multigrid solvers.

HPCG 3.0 Reference Code was released on November 11, 2015 for the SC15 conference in Austin, TX. In addition to bug fixes, this release positions HPCG to even better represent modern PDE solvers, which reflect the behavior of explicit methods that involve unassembled matrices, and aids in running HPCG on production supercomputing installations. The reference version is accompanied by binary releases from Intel and NVIDIA that are carefully optimized for the vendors' respective hardware platforms. Since its inception in 2013, the community's reception of the benchmark has been overwhelmingly positive, and the constant feedback leads to the continuous improvement of the code and its scope. The current HPCG Performance List was also released at SC15 and now features over 60 supercomputing sites. HPCG is a

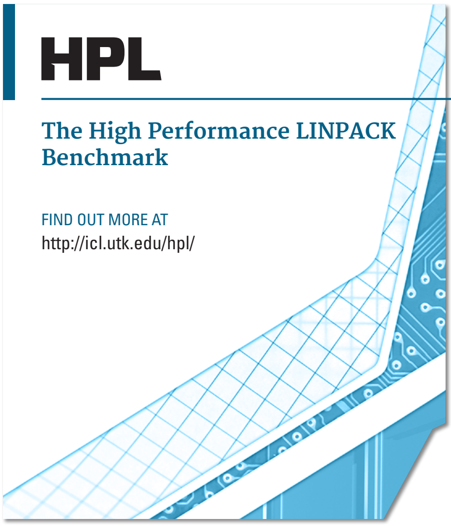

The High Performance LINPACK (HPL) benchmark is a software package that solves a (randomly generated) dense linear system in double precision (64-bit) arithmetic on distributed-memory computers. Written in a portable ANSI C and requiring an MPI implementation as well as either the BLAS or VSIPL library, HPL is often one of the first programs to run on large computer installations, producing a result that can be submitted to the biannual TOP500 list of the world's fastest supercomputers.

HPL 2.1, released in 2012, includes several major bug fixes and accuracy enhancements based on user feedback. The major focus of HPL 2.1 is to improve the accuracy of reported benchmark results, and ensure scalability of the code on large supercomputer installations with hundreds of thousands of computational cores. The last version also featured a detailed time-of-run accounting to help with assessing power requirements at the time of execution, a metric which has been reported with TOP500 results since 2007 and is also highlighted on the Green500 list. In 2011, the LINPACK benchmark app for iOS achieved performance of over 1 Gflop/s on an Apple iPad 2, with per-Watt performance easily beating supercomputing solutions, including the most power-efficient systems based on hardware accelerators. The App now achieves over 4 Gflop/s on the iPad Air.

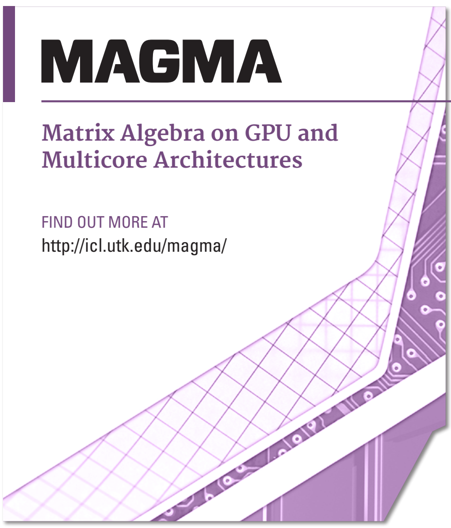

Matrix Algebra on GPU and Multicore Architectures (MAGMA) is a collection of next generation linear algebra (LA) libraries for heterogeneous architectures. The MAGMA package supports interfaces for current LA packages and standards, e.g., LAPACK and BLAS, to allow computational scientists to easily port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

New for 2015, MAGMA 1.7 features top performance and high accuracy LAPACK compliant routines for multicore CPUs enhanced with NVIDIA GPUs, and includes more than 400 routines, covering 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. In 2014, the MAGMA Sparse and MAGMA Batched packages were added with the MAGMA 1.6 release, providing support for sparse iterative and batched linear algebra on a set of small matrices in parallel, respectively. MAGMA provides multiple precision arithmetic support (S/D/C/Z, including mixed-precision). Most of the algorithms are hybrid, using both multi-core CPUs and GPUs, but starting with the 1.6 release, GPU-specific algorithms were added. MAGMA also supports AMD GPUs (clMAGMA 1.3) and Intel Xeon Phi coprocessors (MAGMA MIC 1.4).

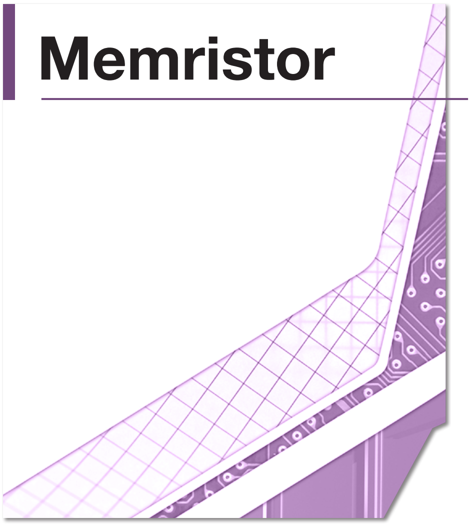

The Memristor project seeks to realize the potential of the recent advances in memristive technology and apply it for the purpose of computing rather than the more common application as a non-volatile storage medium. The project aims to create a specially crafted software layer that interacts with the memristive elements and provides computational services to applications, while providing a performance advantage and uncompromising numerical properties in a new power envelope.

Memristor is a collaboration between ICL/UTK, the University of Rochester, and Sandia National Laboratories. ICL contributes expertise in defining a software interface and providing reliable and efficient implementations that take into account the specific challenges associated with memristive technology which, in the case of computing, represents a new paradigm.

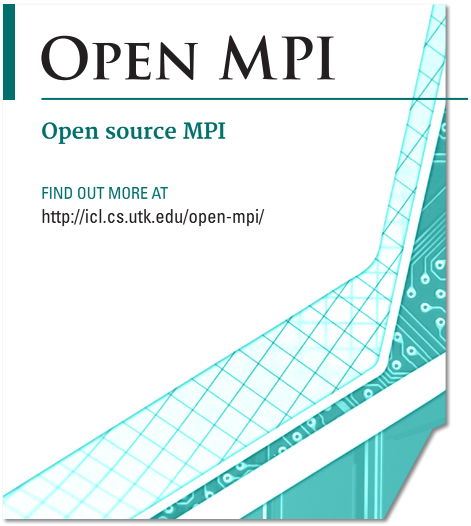

The Open MPI Project is an open source Message Passing Interface (MPI) implementation that is developed and maintained by a consortium of academic, research, and industry partners. MPI primarily addresses the message-passing parallel programming model, in which data is moved from the address space of one process to that of another process through cooperative operations on each process. Open MPI integrates technologies and resources from several other projects (HARNESS/FT-MPI, LA-MPI, LAM/MPI, and PACX-MPI) in order to build the best MPI library available. A completely new MPI-3.1 compliant implementation, Open MPI offers advantages for system and software vendors, application developers, and computer science researchers.

ICL's efforts in the context of Open MPI have significantly improved its scalability, performance on many-core environments, and architecture-aware capabilities, such as adaptive shared memory behaviors and dynamic collective selection, making it ready for the next generation Exascale challenges.

OpenSHMEM is a Partitioned Global Address Space (PGAS) library interface specification that aims to provide a standard Application Programming Interface (API) for SHMEM libraries to aid portability across multiple vendors—including SGI, Cray, IBM, HP, Mellanox, and Intel. OpenSHMEM supports one-sided communication and is a perfect fit for applications with irregular communication patterns with small/medium sized data transfers, since it is optimized for low-latency data transfers.

The OpenSHMEM Library API provides calls for data communication, group synchronization, data collection, data reduction, distributed locking of critical regions, and data and process accessibility to OpenSHMEM PEs. PEs put/get data to/from remotely accessible symmetric data objects on other PEs.

PAPI

The Performance API

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

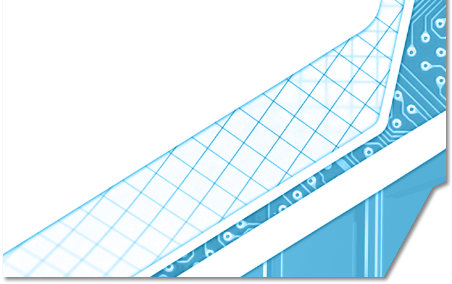

The Performance API (PAPI) provides simultaneous access to performance counters on CPUs, GPUs, and other components of interest (e.g., network and I/O systems). Provided as a linkable library or shared object, PAPI can be called directly in a user program, or used transparently through a variety of third-party tools, making it a de facto standard for hardware counter analysis. Industry liaisons with Bull, 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, usually with no changes at the user level. This provides a richer environment in which performance analysts can work.

In 2015, ICL began work on PAPI-EX, a project which will extend PAPI with measurement tools for changing hardware and software paradigms. New components of PAPI are under development that allow for multiple sources of performance data to be measured simultaneously via a common software interface. Specifically, a new component that controls power is now available with PAPI 5.4.2. In addition, PAPI-EX will incorporate a counter inspection toolkit designed to improve understanding of low-level hardware events. In response to new software paradigms—especially with the increasing popularity of dataflow-based programming models on distributed heterogeneous architectures—PAPI-EX will support integration with task-based runtime systems, enabling hardware performance counter measurements at true task granularity, as opposed to the thread/process granularity achievable today.

PaRSEC

Parallel Runtime Scheduling and Execution Controller

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

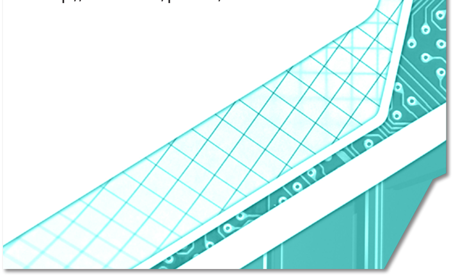

The Parallel Runtime Scheduling and Execution Controller (PaRSEC) is a generic framework for architecture-aware scheduling and management of micro-tasks on distributed many-core heterogeneous architectures. Applications we consider are expressed as a Direct Acyclic Graph (DAG) of tasks, with edges designating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

PaRSEC orchestrates the execution of an algorithm on a particular set of resources, assigns computation threads to the cores, overlaps communications and computations, and uses a dynamic, fully-distributed scheduler based on architectural features such as NUMA nodes and algorithmic features such as data reuse. PaRSEC includes a set of tools to generate the DAGs and integrate them in legacy codes, a runtime library to schedule the micro-tasks 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.

PLASMA

Parallel Linear Algebra Software for Multicore Architectures

FIND OUT MORE AT
http://icl.utk.edu/plasma/

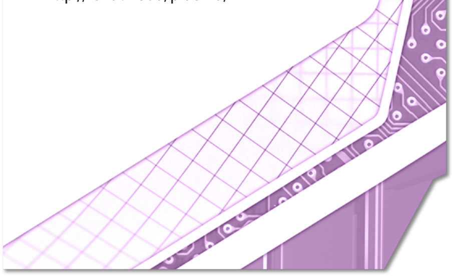

The Parallel Linear Algebra Software for Multicore Architectures (PLASMA) package is a dense linear algebra package at the forefront of multicore computing, designed to deliver the highest possible performance from a system with multiple sockets of multicore processors. PLASMA achieves this objective by combining state-of-the-art solutions in parallel algorithms, scheduling, and software engineering. Currently, PLASMA offers a collection of routines for solving linear systems of equations, least square problems, eigenvalue problems, and singular value problems.

PLASMA relies on runtime scheduling of parallel tasks, which is based on the idea of assigning work to cores based on the availability of data for processing at any given point in time. The concept, which is sometimes called data-driven scheduling, is closely related to the idea of expressing computation through a task graph, often referred to as the DAG (Directed Acyclic Graph), and the flexibility of exploring the DAG at runtime.

PULSAR

Parallel Ultra-Light Systolic Array Runtime

FIND OUT MORE AT
http://icl.utk.edu/pulsar/

The Parallel Ultra-Light Systolic Array Runtime (PULSAR), now in version 2.0, is a complete programming platform for large-scale distributed memory systems with multicore processors and hardware accelerators. PULSAR provides a simple abstraction layer over multithreading, message-passing, and multi-GPU, multi-stream programming. PULSAR offers a general-purpose programming model, suitable for a wide range of scientific and engineering applications.

This simple programming model allows the user to define the computation in the form of a Virtual Systolic Array (VSA), which is a set of Virtual Data Processors (VDPs), and is connected with data channels. This programming model is also accessible to the user through a very small and simple Application Programming Interface (API), and all the complexity of executing the workload on a large-scale system is hidden in the runtime implementation.

QUARK

QUeuing And Runtime for Kernels

FIND OUT MORE AT
http://icl.utk.edu/quark/

QUARK (QUeuing And Runtime for Kernels) provides a library that enables the dynamic, superscalar execution of tasks with data dependencies in a multi-core, multi-socket, shared-memory environment. QUARK infers data dependencies and precedence constraints between tasks based on the way the data is used, and then executes the tasks in a dynamic, asynchronous, superscalar fashion in order to achieve a high utilization of the available resources.

QUARK is designed to be easy to use, scales to large numbers of cores, and enables the efficient expression and implementation of complex algorithms. The QUARK runtime is codesigned with the PLASMA linear algebra library, and it contains optimizations inspired by the algorithms in PLASMA.

RaPyDLI

Rapid Python Deep Learning Infrastructure

The Rapid Python Deep Learning Infrastructure (RaPyDLI) project delivers productivity and performance to the Deep Learning community by combining high level Python, C/C++, and Java environments with carefully designed libraries supporting GPU accelerators and Xeon Phi coprocessors. Deep Learning (DL) has made major impacts in areas like speech recognition, drug discovery, and computer vision. This success relies on training large neural nets—currently, up to 10 billion connections trained on 10 million images—using either large scale commodity clusters or smaller HPC systems where accelerators perform with high efficiency. This approach is of prime importance as the hardware accelerators enable much more sophisticated neural networks by increasing the available computational power by more than an order of magnitude.

RaPyDLI is a collaboration between ICL/UTK, Indiana University, and Stanford University, with each institution contributing their long standing expertise in the field. Currently, ICL's focus for the RaPyDLI project is on efficient GPU kernel execution and optimization of scheduling strategies to reduce inefficiencies in the current code base in terms of performance and idle time.

LAPACK

The Linear Algebra PACKage

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

ScaLAPACK

Scalable LAPACK

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

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

LAPACK 3.6.0 was released in November 2015. LAPACK 3.6.0 includes BLAS3 routines for generalized SVD; new routines to compute a Subset of the Singular Value Decomposition, full or partial (subset) SVD of a bidiagonal matrix through an associated eigenvalue problem, full or partial (subset) SVD of a general matrix; new Complex Jacobi SVD routines; Recursive Cholesky routines; and some improvements to QR (NEP). Since 2011, LAPACK has included LAPACKE, a native C interface for LAPACK developed in collaboration with INTEL, which provides NAN check and automatic workspace allocation. ScaLAPACK 2.0.0, which includes the MRRR algorithm and new Nonsymmetric Eigenvalue Problem routines, was released in November 2011. Two additional ScaLAPACK versions (2.0.1 and 2.0.2) were released in 2012 for minor bug fixes.

SILAS

Sustained Innovation for Linear Algebra Software

The main goal of the NSF funded project on Sustained Innovation for Linear Algebra Software (SILAS) is to update two of the most widely used numerical libraries in the history of Computational Science and Engineering—LAPACK and ScaLAPACK—for the ongoing revolution in processor architecture and system design. Working with partners at the University of California, Berkeley and the University of Colorado, Denver, ICL is using SILAS to enhance and harden these essential libraries in order to prepare them for the kind of extreme scale systems and applications that are now coming online.

SILAS is organized around three complementary objectives: 1) Wherever possible, SILAS delivers seamless access to the most up-to-date algorithms, numerical implementations, and performance, by way of the familiar Sca/LAPACK programming interface; 2) Wherever necessary, SILAS makes advanced algorithms, numerical implementations and performance capabilities available through new interface extensions; and 3) SILAS provides a well engineered conduit through which new discoveries at the frontiers of research in these areas can be channeled as quickly as possible to all the application communities.

SparseKaffe

The Sparse direct methods via Run-time Scheduling and Execution of Kernels with Auto-tunable and Frequency-scaling Features for Energy-aware computing on heterogeneous architectures (SparseKafe) project will create fast and efficient sparse direct methods for platforms with multi-core processors with one or more accelerators (e.g., GPUs or Xeon Phi coprocessors). SparseKaffe spans the platform pyramid, from desktop machines to extreme scale systems consisting of multiple heterogeneous nodes connected through a high-speed network, with the goal of achieving orders of magnitude gains in computational performance, while also paying careful attention to the energy requirements.

The SparseKaffe project is a collaboration between ICL/UTK, the University of Florida, and Texas A&M University. ICL's work on the project will concentrate on dynamic runtime scheduling using the dataflow model, which will leverage, and be a natural extension of, ICL's work on runtimes as part of the MAGMA, PLASMA, and PaRSEC projects. The autotuning of the algorithm-specific computational kernels will apply the principles behind ICL's BEAST project.

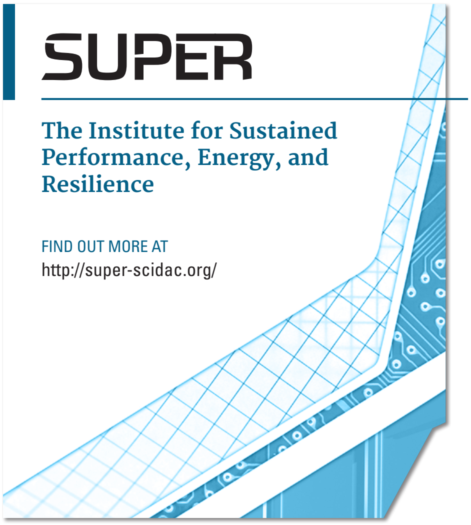

The Institute for Sustained Performance, Energy, and Resilience (SUPER), led by the University of Southern California, has organized a broad-based project involving several universities and DOE laboratories with expertise in compilers, system tools, performance engineering, energy management, and resilience to ensure that DOE's computational scientists can successfully exploit the emerging generation of high performance computing (HPC) systems.

SUPER is extending performance modeling and autotuning technology to heterogeneous and petascale computing systems, investigating application-level energy efficiency techniques, exploring resilience strategies for petascale applications, and developing strategies that collectively optimize performance, energy efficiency, and resilience. UTK work focuses on performance measurement, power and energy measurements, and resilience techniques for hard and soft errors.

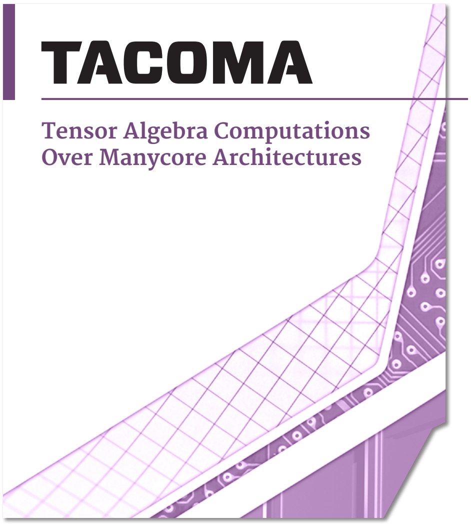

The goal of the Tensor Algebra Computations Over Manycore Architectures (TACOMA) project is to design a high-performance tensor contractions package for heterogeneous systems with multicore CPUs and GPU accelerators. TACOMA will target the acceleration of a number of important applications, including high-order finite element method (FEM) simulations, machine learning, big data analytics, multi-physics, and more.

The TACOMA project is a collaboration between UTK/ICL, Lawrence Livermore National Lab (LLNL), the University of Paris-Sud, and Inria to develop software infrastructure for tensor contractions to, in part, expedite the porting of packages like BLAST to next generation 150+ petaflop/s DOE machines (e.g., Sierra, LLNL's IBM Power9 machine). The TACOMA effort is multidisciplinary, incorporating linear algebra, languages, code generation and optimizations, domain science, and application-specific numerical algorithms.

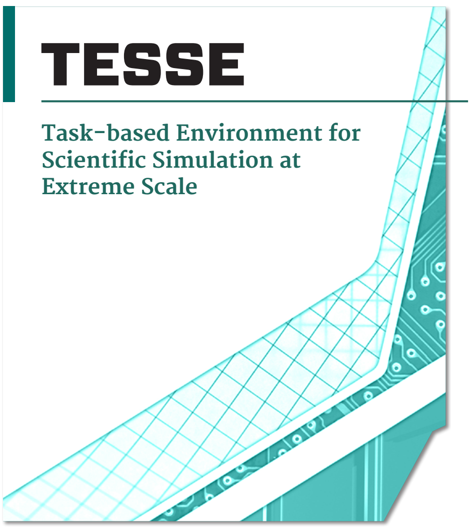

The goal of the Task-based Environment for Scientific Simulation at Extreme Scale (TESSE) project is to use an application-driven design to create a general-purpose production-quality software framework that attacks the twin challenges of programmer productivity and portable performance for advanced scientific applications on massively-parallel, hybrid, many-core systems of today and tomorrow.

The TESSE team is composed of researchers from Stony Brook, Virginia Tech, and ICL/UTK, who have designed a system that uses DAG-based data flow as the basis of the software. This capability, with the extensions being explored by the TESSE project, will provide significant potential advantages in ease of composition, performance, and ease of migration to future architectures for irregular parallel applications. The TESSE team's next major goal is the ubiquitous existence of a powerful DAG-based data flow tool that complements, and is completely interoperable with, mainstream standard parallel programming models such as OpenMP and MPI.

TOP500

Ranking the 500 fastest computers in the world

FIND OUT MORE AT
http://www.top500.org/

Since 1993, a ranking of the top 500 fastest computers in the world has been compiled biannually with published results released in June and November at major HPC gatherings in the US and Europe. Each machine on the TOP500 list is ranked based on its performance result from running the computationally intensive High Performance LINPACK (HPL) benchmark developed by ICL, which may be further optimized by the vendor, integrator, or the benchmarking engineer from the affiliated institution.

While new benchmarks, including HPCG, have been developed to measure performance of HPC systems, the TOP500 still relies on the HPL benchmark and remains the de-facto ranking used by commercial, industrial, government, and academic institutions. ICL continues to partner with NERSC/Lawrence Berkeley National Laboratory and Prometeus to produce the rankings.

In June 2013, China's Tianhe-2 supercomputer took the #1 spot on the TOP500 at an astounding 33.86 petaflop/s, and continues to maintain that position as of November 2015. The surprise appearance of Tianhe-2, two years ahead of the expected deployment, marks China's first return to the #1 position since November 2010, when Tianhe-1A was the top system. Tianhe-2 has 16,000 nodes, each with two Intel Xeon Ivy Bridge processors and three Xeon Phi coprocessors for a combined total of 3,120,000 computing cores.

UCX

Universal Communication X

The Universal Communication X (UCX) is a new collaborative effort between industry, laboratories, and academia to create an open-source, production grade communication framework for data centric and high-performance applications. UCX makes common network code available through a universal, portable, and performance driven communication API that exposes low level capabilities of high performance networks, henceforth providing critical services for implementing features needed by high level parallel programming constructs.

These services include lightweight remote memory access operations, lightweight synchronizations, active messages, atomic operations, etc. Open UCX has been successfully deployed to support communication middlewares, including OpenSHMEM, Open MPI, and PaRSEC.

ULFM

User Level Failure Mitigation

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

User Level Failure Mitigation (ULFM) is a set of new interfaces for the Message Passing Interface (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 whose completions were rendered impossible by failures.

Using the constructs defined by ULFM, applications and libraries drive the recovery of the MPI state. Consistency issues resulting from failures are addressed according to an application's needs, and the recovery actions are limited to the necessary MPI communication objects. Therefore, the recovery scheme is more efficient than a generic, automatic recovery technique, and can achieve the goals of enabling applications to resume communication after failure and maintaining extreme communication performance outside of recovery periods. A wide range of application types and middlewares are already building on top of ULFM to deliver scalable and user friendly fault tolerance.

2015 PUBLICATIONS

Anzt, H., E. Ponce, G. D. Peterson, and J. Dongarra, “**GPU-accelerated Co-design of Induced Dimension Reduction: Algorithmic Fusion and Kernel Overlap**”, *2nd International Workshop on Hardware-Software Co-Design for High Performance Computing*, Austin, TX, ACM, November 2015.

Anzt, H., J. Dongarra, and E. S. Quintana-Ortí, “**Adaptive Precision Solvers for Sparse Linear Systems**”, *3rd International Workshop on Energy Efficient Supercomputing (E2SC ‘15)*, Austin, TX, ACM, November 2015.

Anzt, H., W. Sawyer, S. Tomov, P. Luszczek, and J. Dongarra, “**Acceleration of GPU-based Krylov solvers via Data Transfer Reduction**”, *International Journal of High Performance Computing Applications*, 2015.

Anzt, H., E. Chow, and J. Dongarra, “**Iterative Sparse Triangular Solves for Preconditioning**”, *EuroPar 2015*, Vienna, Austria, Springer Berlin, August 2015.

Anzt, H., J. Dongarra, and E. S. Quintana-Ortí, “**Tuning Stationary Iterative Solvers for Fault Resilience**”, *6th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (ScalA15)*, Austin, TX, ACM, November 2015.

Anzt, H., E. Chow, D. Szyld, and J. Dongarra, “**Random-Order Alternating Schwarz for Sparse Triangular Solves**”, *2015 SIAM Conference on Applied Linear Algebra (SIAM LA)*, Atlanta, GA, SIAM, October 2015.

Anzt, H., S. Tomov, and J. Dongarra, “**Energy efficiency and performance frontiers for sparse computations on GPU supercomputers**”, *Sixth International Workshop on Programming Models and Applications for Multicores and Manycores (PMAM ‘15)*, San Francisco, CA, ACM, February 2015.

Anzt, H., S. Tomov, and J. Dongarra, “**Accelerating the LOBPCG method on GPUs using a blocked Sparse Matrix Vector Product**”, *Spring Simulation Multi-Conference 2015 (SpringSim’15)*, Alexandria, VA, SCS, April 2015.

Anzt, H., B. Haugen, J. Kurzak, P. Luszczek, and J. Dongarra, “**Experiences in autotuning matrix multiplication for energy minimization on GPUs**”, *Concurrency in Computation: Practice and Experience*, May 2015.

Aupy, G., and Y. Robert, “**Scheduling for fault-tolerance: an introduction**”, *Innovative Computing Laboratory Technical Report*, no. ICL-UT-15-02: University of Tennessee, January 2015.

Baboulin, M., J. Dongarra, A. Remy, S. Tomov, and I. Yamazaki, “**Dense Symmetric Indefinite Factorization on GPU acclerated architectures**”, *International Conference on Parallel Processing and Applied Mathematics (PPAM2015)*, Krakow, Poland, September 2015.

Bosilca, G., A. Bouteiller, T. Herault, Y. Robert, and J. Dongarra, “**Composing Resilience Techniques: ABFT, Periodic, and Incremental Checkpointing**”, *International Journal of Networking and Computing*, vol. 5, no. 1, pp. 2-15, January 2015.

Bouteiller, A., G. Bosilca, and J. Dongarra, “**Plan B: Interruption of Ongoing MPI Operations to Support Failure Recovery**”, *22nd European MPI Users’ Group Meeting*, Bordeaux, France, ACM, September 2015.

Bouteiller, A., T. Herault, G. Bosilca, P. Du, and J. Dongarra, “**Algorithm-based Fault Tolerance for Dense Matrix Factorizations, Multiple Failures, and Accuracy**”, *ACM Transactions on Parallel Computing*, vol. 1, issue 2, no. 10, pp. 10:1-10:28, January 2015.

Cao, C., G. Bosilca, T. Herault, and J. Dongarra, “**Design for a Soft Error Resilient Dynamic Task-based Runtime**”, *29th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Hyderabad, India, IEEE, May 2015.

Chow, E., H. Anzt, and J. Dongarra, “**Asynchronous Iterative Algorithm for Computing Incomplete Factorizations on GPUs**”, *International Supercomputing Conference (ISC 2015)*, Frankfurt, Germany, July 2015.

Danalis, A., H. Jagode, G. Bosilca, and J. Dongarra, “**PaRSEC in Practice: Optimizing a Legacy Chemistry Application through Distributed Task-Based Execution**”, *IEEE Cluster 2015*, Chicago, IL, IEEE, September 2015.

Donfack, S., J. Dongarra, M. Faverge, M. Gates, J. Kurzak, P. Luszczek, and I. Yamazaki, “**A Survey of Recent Developments in Parallel Implementations of Gaussian Elimination**”, *Concurrency and Computation: Practice and Experience*, vol. 27, issue 5, pp. 1292-1309, April 2015.

Dongarra, J., Editors: N. J. Higham, M. R. Dennis, P. Glendinning, P. A. Martin, F. Santosa, and J. Tanner, “**High-Performance Computing**”, *The Princeton Companion to Applied Mathematics*, Princeton, New Jersey, Princeton University Press, pp. 839-842, 2015.

Dongarra, J., M. A. Heroux, and P. Luszczek, “**High-Performance Conjugate-Gradient Benchmark: A New Metric for Ranking High-Performance Computing Systems**”, *The International Journal of High Performance Computing Applications*, 2015.

Dongarra, J., M. A. Heroux, and P. Luszczek, “**HPCG Benchmark: a New Metric for Ranking High Performance Computing Systems**”, *University of Tennessee Computer Science Technical Report*, no. ut-eecs-15-736: University of Tennessee, January 2015.

Dongarra, J., T. Herault, and Y. Robert, “**Fault Tolerance Techniques for High-performance Computing**”, *University of Tennessee Computer Science Technical Report* (also LAWN 289), no. UT-EECS-15-734: University of Tennessee, May 2015.

Faverge, M., J. Herrmann, J. Langou, B. Lowery, Y. Robert, and J. Dongarra, “**Mixing LU-QR Factorization Algorithms to Design High-Performance Dense Linear Algebra Solvers**”, *Journal of Parallel and Distributed Computing*, vol. 85, pp. 32-46, November 2015.

Gates, M., H. Anzt, J. Kurzak, and J. Dongarra, “**Accelerating Collaborative Filtering for Implicit Feedback Datasets using GPUs**”, *2015 IEEE International Conference on Big Data (IEEE BigData 2015)*, Santa Clara, CA, IEEE, November 2015.

Gates, M., S. Tomov, and A. Haidar, “**Comparing Hybrid CPU-GPU and Native GPU-only Acceleration for Linear Algebra**”, *2015 SIAM Conference on Applied Linear Algebra*, Atlanta, GA, SIAM, October 2015.

Haidar, A., A. Abdelfattah, S. Tomov, and J. Dongarra, “**Batched Matrix Computations on Hardware Accelerators Based on GPUs**”, *2015 SIAM Conference on Applied Linear Algebra (SIAM LA)*, Atlanta, GA, SIAM, October 2015.

Haidar, A., J. Kurzak, G. Pichon, and M. Faverge, “**A Data Flow Divide and Conquer Algorithm for Multicore Architecture**”, *29th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Hyderabad, India, IEEE, May 2015.

Haidar, A., P. Luszczek, S. Tomov, and J. Dongarra, “**Efficient Eigensolver Algorithms on Accelerator Based Architectures**”, *2015 SIAM Conference on Applied Linear Algebra (SIAM LA)*, Atlanta, GA, SIAM, October 2015.

Haidar, A., T. Dong, P. Luszczek, S. Tomov, and J. Dongarra, "**Optimization for Performance and Energy for Batched Matrix Computations on GPUs**", *8th Workshop on General Purpose Processing Using GPUs (GPGPU 8)*, San Francisco, CA, ACM, February 2015.

Haidar, A., T. Dong, P. Luszczek, S. Tomov, and J. Dongarra, "**Batched matrix computations on hardware accelerators based on GPUs**", *International Journal of High Performance Computing Applications*, February 2015.

Haidar, A., S. Tomov, P. Luszczek, and J. Dongarra, "**MAGMA Embedded: Towards a Dense Linear Algebra Library for Energy Efficient Extreme Computing**", *2015 IEEE High Performance Extreme Computing Conference (HPEC '15)*, (Best Paper Award), Waltham, MA, IEEE, September 2015.

Haidar, A., P. Luszczek, S. Tomov, and J. Dongarra, "**Towards Batched Linear Solvers on Accelerated Hardware Platforms**", *8th Workshop on General Purpose Processing Using GPUs (GPGPU 8)*, San Francisco, CA, ACM, February 2015.

Haidar, A., T. Dong, S. Tomov, P. Luszczek, and J. Dongarra, "**Framework for Batched and GPU-resident Factorization Algorithms to Block Householder Transformations**", *ISC High Performance*, Frankfurt, Germany, Springer, July 2015.

Haidar, A., J. Dongarra, K. Kabir, M. Gates, P. Luszczek, S. Tomov, and Y. Jia, "**HPC Programming on Intel Many-Integrated-Core Hardware with MAGMA Port to Xeon Phi**", *Scientific Programming*, vol. 23, issue 1, January 2015.

Haidar, A., A. YarKhan, C. Cao, P. Luszczek, S. Tomov, and J. Dongarra, "**Flexible Linear Algebra Development and Scheduling with Cholesky Factorization**", *17th IEEE International Conference on High Performance Computing and Communications*, Newark, NJ, August 2015.

Haidar, A., P. Luszczek, S. Tomov, and J. Dongarra, "**Batched Matrix Computations on Hardware Accelerators**", *EuroMPI/Asia 2015 Workshop*, Bordeaux, France, September 2015.

Haugen, B., S. Richmond, J. Kurzak, C. A. Steed, and J. Dongarra, "**Visualizing Execution Traces with Task Dependencies**", *2nd Workshop on Visual Performance Analysis (VPA '15)*, Austin, TX, ACM, November 2015.

Herault, T., A. Bouteiller, G. Bosilca, M. Gamell, K. Teranishi, M. Parashar, and J. Dongarra, "**Practical Scalable Consensus for Pseudo-Synchronous Distributed Systems: Formal Proof**", *Innovative Computing Laboratory Technical Report*, no. ICL-UT-15-01, April 2015.

Herault, T., A. Bouteiller, G. Bosilca, M. Gamell, K. Teranishi, M. Parashar, and J. Dongarra, "**Practical Scalable Consensus for Pseudo-Synchronous Distributed Systems**", *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC15)*, Austin, TX, ACM, November 2015.

Jagode, H., A. YarKhan, A. Danalis, and J. Dongarra, "**Power Management and Event Verification in PAPI**", *9th Parallel Tools Workshop*, Dresden, Germany, September 2015.

Jagode, H., A. Danalis, G. Bosilca, and J. Dongarra, "**Accelerating NWChem Coupled Cluster through dataflow-based Execution**", *11th International Conference on Parallel Processing and Applied Mathematics (PPAM 2015)*, Krakow, Poland, September 2015.

Kabir, K., A. Haidar, S. Tomov, and J. Dongarra, "**Performance Analysis and Optimization of Two-Sided Factorization Algorithms for Heterogeneous Platform**", *International Conference on Computational Science (ICCS 2015)*, Reykjavík, Iceland, June 2015.

Kabir, K., A. Haidar, S. Tomov, and J. Dongarra, "**Performance Analysis and Design of a Hessenberg Reduction using Stabilized Blocked Elementary Transformations for New Architectures**", *The Spring Simulation Multi-Conference 2015 (SpringSim'15)*, Best Paper Award, Alexandria, VA, April 2015.

Kabir, K., A. Haidar, S. Tomov, and J. Dongarra, "**On the Design, Development, and Analysis of Optimized Matrix-Vector Multiplication Routines for Coprocessors**", *ISC High Performance 2015*, Frankfurt, Germany, July 2015.

Kurzak, J., H. Anzt, M. Gates, and J. Dongarra, "**Implementation and Tuning of Batched Cholesky Factorization and Solve for NVIDIA GPUs**", *IEEE Transactions on Parallel and Distributed Systems*, no. 1045-9219, November 2015.

Mary, T., I. Yamazaki, J. Kurzak, P. Luszczek, S. Tomov, and J. Dongarra, "**Performance of Random Sampling for Computing Low-rank Approximations of a Dense Matrix on GPUs**", *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC15)*, Austin, TX, ACM, November 2015.

Reed, D., and J. Dongarra, "**Exascale Computing and Big Data**", *Communications of the ACM*, vol. 58, no. 7: ACM, pp. 56-68, July 2015.

Solcà, R., A. Kozhevnikov, A. Haidar, S. Tomov, T. C. Schulthess, and J. Dongarra, "**Efficient Implementation Of Quantum Materials Simulations On Distributed CPU-GPU Systems**", *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC15)*, Austin, TX, ACM, November 2015.

Song, F., and J. Dongarra, "**A Scalable Approach to Solving Dense Linear Algebra Problems on Hybrid CPU-GPU Systems**", *Concurrency and Computation: Practice and Experience*, vol. 27, issue 14, pp. 3702-3723, September 2015.

Strohmaier, E., H. Meuer, J. Dongarra, and H. D. Simon, "**The TOP500 List and Progress in High-Performance Computing**", *IEEE Computer*, vol. 48, issue 11, pp. 42-49, November 2015.

Wu, W., A. Bouteiller, G. Bosilca, M. Faverge, and J. Dongarra, "**Hierarchical DAG scheduling for Hybrid Distributed Systems**", *29th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Hyderabad, India, IEEE, May 2015.

Yamazaki, I., S. Tomov, and J. Dongarra, "**Mixed-Precision Cholesky QR Factorization and its Case Studies on Multicore CPU with Multiple GPUs**", *SIAM Journal on Scientific Computing*, vol. 37, no. 3, pp. C203-C330, May 2015.

Yamazaki, I., J. Kurzak, P. Luszczek, and J. Dongarra, "**Randomized Algorithms to Update Partial Singular Value Decomposition on a Hybrid CPU/GPU Cluster**", *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC15)*, Austin, TX, ACM, November 2015.

Yamazaki, I., S. Tomov, and J. Dongarra, "**Computing Low-rank Approximation of a Dense Matrix on Multicore CPUs with a GPU and its Application to Solving a Hierarchically Semiseparable Linear System of Equations**", *Scientific Programming*, 2015.

Yamazaki, I., J. Barlow, S. Tomov, J. Kurzak, and J. Dongarra, "**Mixed-precision orthogonalization process Performance on multicore CPUs with GPUs**", *2015 SIAM Conference on Applied Linear Algebra*, Atlanta, GA, SIAM, October 2015.

Yamazaki, I., J. Kurzak, P. Luszczek, and J. Dongarra, "**Randomized Algorithms to Update Partial Singular Value Decomposition on a Hybrid CPU/GPU Cluster**", *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC15)*, Austin, TX, ACM, November 2015.

YarKhan, A., A. Haidar, C. Cao, P. Luszczek, S. Tomov, and J. Dongarra, "**Cholesky Across Accelerators**", *17th IEEE International Conference on High Performance Computing and Communications (HPCC 2015)*, Elizabeth, NJ, IEEE, August 2015.

2015 EVENTS

JANUARY 27-29
Open MPI developers meeting
Dallas, TX

JANUARY 29-30
Big Data and Extreme-scale Computing (BDEC)
Barcelona, Spain

FEBRUARY 7-11
20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP 2015)
San Francisco, CA

FEBRUARY 17-18
NSF Software Infrastructure for Sustained Innovation (SI2) Principal Investigator (PI) Workshop
Arlington, VA

MARCH 2-5
Message Passing Interface (MPI) Forum
Portland, OR

MARCH 4
SUPER all hands meeting
San Diego, CA

MARCH 14-18
SIAM Conference on Computational Science and Engineering
Salt Lake City, UT

MARCH 17
GPU Technology Conference
San Jose, CA

MARCH 17-20
Supercomputing Frontiers
Singapore

APRIL 12-15
23rd High Performance Computing Symposium (HPC 2015)
Alexandria (Washington, DC), VA

MAY 25-29
30th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2015)
Hyderabad, India

JUNE 1-4
Message Passing Interface (MPI) Forum
Chicago, IL

JUNE 23-25
Open MPI Developers Meeting
San Jose, CA

JUNE 29-JULY 2
Sparse Days
Saint-Girons, France

JULY 12-16
ISC High Performance 2015
Frankfurt, Germany

AUGUST 4-6
OpenSHMEM 2015: Second Workshop on OpenSHMEM and Related Technologies
Annapolis, DE

AUGUST 24-28
21th International European Conference on Parallel and Distributed Computing (EuroPar 2015)
Vienna, Austria

AUGUST 24-26
17th IEEE International Conference on High Performance Computing and Communications (HPCC 2015)
Newark, NJ

AUGUST 31-SEP 2
Smoky Mountains Computational Sciences and Engineering Conference (SMC15)
Gatlinburg, TN

SEPTEMBER 2-3
9th Parallel Tools Workshop
Dresden, Germany

SEPTEMBER 6-9
11th International Conference on Parallel Processing and Applied Mathematics (PPAM15)
Krakow, Poland

SEPTEMBER 9-11
IEEE Cluster 2015
Chicago, IL

SEPTEMBER 9
ORNL Scientific Seminar Series
Oak Ridge, TN

SEPTEMBER 15-17
2015 IEEE High Performance Extreme Computing Conference (HPEC'15)
Waltham, MA

SEPTEMBER 21-23
EuroMPI 2015
Bordeaux, France

Every year, members of our research staff attend national and international conferences, workshops, and seminars. These meetings provide opportunities to present our research, share our knowledge, and exchange ideas with leading computational science researchers from around the world. Showing the work we have done and participating in the intellectual life of the scientific community in this way is an essential part of the research process. The following pages contain a list of events we have participated in over the past year.

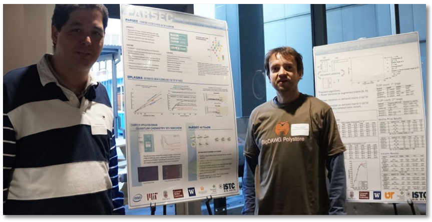

SEPTEMBER 24
Intel Big Data Retreat
Hillsboro, OR

SEPTEMBER 28-29
Russian Supercomputing Days
Moscow, Russia

OCTOBER 1
ARGO Members Meeting
Argonne, IL

OCTOBER 12-14
3rd International Workshops on Advances in Computational Mechanics (IWACOM-III)
Tokyo, Japan

OCTOBER 15-16
Workshop on Computational Science & Engineering Software Sustainability and Productivity Challenges (CSESSP)
Rockford, MD

OCTOBER 20-21
White House Workshop on the National Strategic Computing Initiative (NSCI)
Washington, DC

OCTOBER 26-30
SIAM Conference on Applied Linear Algebra (SIAM LA)
Atlanta, GA

NOVEMBER 15-20
27th International Conference for High Performance Computing, Networking, Storage and Analysis (SC15)
Austin, TX

NOVEMBER 20
2nd Workshop on Visual Performance Analysis (VPA)
Austin, TX

SUPERCOMPUTING '15

The International Conference for High Performance Computing Networking, Storage, and Analysis (SC), established in 1988, is a staple of ICL's November itinerary. SC is vital to the growth and evolution of high performance computing in the United States because it is the only US event that elicits substantial participation from all segments of the HPC community, including hundreds of users, developers, vendors, research institutions, and representatives of government funding agencies. Such a talent-rich gathering enables participants to discuss challenges, share innovations, and coordinate relationships and collaborations with some of the best minds in scientific and high performance computing.

SC15 was held in Austin, TX on November 15 – 20. As usual, ICL had a significant presence at SC, with faculty, research staff, and students giving talks, presenting papers, and leading "Birds-of-a-Feather" sessions. For the fourth consecutive year, ICL was active in the University of Tennessee's SC booth. The booth, which was organized and led by the National Institute for Computational Sciences (NICS), was visually designed with the help of ICL/CITR staff, manned with support from ICL researchers attending SC, and featured some of the lab's research projects.

COMPUTATIONAL RESOURCES

As the GPU hybrid computing paradigm continues to lead the evolution of computational hardware into Petascale computing, computing architectures are increasingly changing. However, the programming tools, applications, and algorithms that form the backbone of the ever growing need for greater performance are equally important. Such myriad hardware/software configurations present unique challenges that require testing and development of applications that are often unique to the platform on which they reside. For this reason, it is imperative that we have access to a wide range of computing resources in order to conduct our cutting-edge research.

INTEL PARALLEL COMPUTING CENTER

The Intel Parallel Computing Center (IPCC) program, renewed with ICL in 2015, is comprised of universities, institutions, and labs that are leaders in their field, focusing on modernizing applications to increase parallelism and scalability through optimizations that leverage cores, caches, threads, and vector capabilities of microprocessors and coprocessors.

The objective of the Innovative Computing Laboratory's IPCC is the development and optimization of numerical linear algebra libraries and technologies for applications, while tackling current challenges in heterogeneous Intel® Xeon Phi™ coprocessor-based High Performance Computing. In collaboration with Intel's MKL team, the IPCC at ICL will modernize the popular LAPACK and ScaLAPACK libraries to run efficiently on current and future manycore architectures, and will disseminate the developments through the open source MAGMA MIC library.

NVIDIA GPU CENTER OF EXCELLENCE

The Innovative Computing Laboratory joins a very small and select group of labs given a GPU Center of Excellence designation. ICL/UTK's GPU COE focuses on the development of numerical linear algebra libraries for CUDA-based hybrid architectures. ICL's work on the Matrix Algebra on GPU and Multicore Architectures (MAGMA) project further enables and expands our CUDA-based software library efforts, especially in the area of high-performance scientific computing.

The GPU COE designation also led to the establishment of a productive long-term collaboration between ICL and NVIDIA. As part of the collaboration and GPU COE designation, ICL has continuously received hardware, financial support, and other resources from NVIDIA.


Newton Cluster
UNIVERSITY OF TENNESSEE

CAMPUS RESOURCES

To meet these challenges, ICL has access to multiple state-of-the-art heterogeneous systems in house. In fact, our hardware industry partners, including AMD, Intel, and NVIDIA, provide us with bleeding edge hardware resources (often under NDA and prior to public release) which we use to upgrade and maintain the lab's infrastructure.

In 2015, ICL competed for and won a Scholarly Activity and Research Incentive Fund (SARIF). These SARIF funds, provided by the University's Office of Research and Engagement (ORE), allowed ICL to make a significant addition to the lab's hardware infrastructure: a 9 node cluster; each node hosting 2 sockets of Intel Xeon E5-2650 v3 10-core Haswell CPUs, 64GB of RAM, with accelerators from NVIDIA and Intel, all linked through an Infiniband interconnect.

This cluster is invaluable to our research scientists and members of the University's EECS department, with whom we share access, and bolsters our existing campus resources, which include other heterogeneous ICL machines and UTK's Newton Cluster.

REGIONAL RESOURCES


Titan, Cray XK7
OAK RIDGE NATIONAL LABORATORY


Darter System Integrated with XSEDE
NATIONAL INSTITUTE FOR COMPUTATIONAL SCIENCES

ICL has access to many local resources in East Tennessee to help keep us at the forefront of enabling technology research, including Titan, a Cray XK7 supercomputer, which currently holds the number two spot on the TOP500 at 17.59 petaflop/s.

In fact, the recent modernization of the DOE's National Center for Computational Sciences (NCCS), just 30 minutes away at the Oak Ridge National Laboratory (ORNL), has enabled us to leverage our ORNL collaborations to take advantage of what has become one of the world's fastest scientific computing facilities.

GRID RESOURCES

With the continuing trend of high performance grid and cloud computing, it is important for ICL to have access to these types of infrastructures in order to test and implement our software packages on grid hardware and virtualized environments. In keeping with this goal, research staff members at ICL have access to grid resources all over the US, as well as Europe, including XSEDE and Grid5000.

PARTNERSHIPS

Since 1989, ICL has fostered relationships with many other academic institutions and research centers. We have also aggressively sought to build lasting, collaborative partnerships with HPC vendors and industry research leaders, both here and abroad. In this section, we recognize many of the partners and collaborators that we have worked with over the years, most of whom we are still actively involved with.

GOVERNMENT AND ACADEMIC

Our relationships with academic and government research institutions also play a pivotal role in our success. By exchanging ideas, expertise, and personnel, ICL becomes more dynamic with each new collaboration. Our lab routinely develops relationships with researchers whose primary focus is on other scientific disciplines, such as biology, chemistry, or physics, which makes many of our collaborations truly multidisciplinary.

INDUSTRY

Vendors and industry research leaders are an integral part of our partnerships, contributing significantly to our efforts to be a world leader in computational science research. Many have utilized our work, including our linear algebra libraries and performance analysis tools. As a result of these exchanges, we maintain close working relationships with many industry leaders.

INTEL SCIENCE AND TECHNOLOGY CENTER FOR BIG DATA

ICL joined the SciDB project of the Intel Science and Technology Center (ISTC) for Big Data, one of a series of research collaborations that Intel is establishing with universities in the US to identify and prototype revolutionary technology opportunities, and exchange expertise in various fields of high performance computing.

In the case of ICL, the lab will help improve the efficiency of large scale data analytics by providing efficient codes for linear algebra on the Intel Xeon Phi coprocessor. The lab will also provide expertise on fault tolerance to help make the compute intensive portion of data management more resilient, which is essential given the large databases used in Big Data applications. Finally, the distributed nature of large data processing calls for optimal data distribution and redistribution operations, which has long been one of ICL's core strengths.

INTERNATIONAL COLLABORATORS

Barcelona Supercomputing Center	Barcelona, Spain
Central Institute for Applied Mathematics	Jülich, Germany
Danish Computing Center for Research and Education	Lyngby, Denmark
Doshisha University	Kyoto, Japan
Dresden University of Technology	Dresden, Germany
École Normale Supérieure de Lyon	Lyon, France
École Polytechnique Federale de Lausanne	Lausanne, Switzerland
European Centre for Research and Advanced Training in Scientific Computing	Toulouse, France
European Exascale Software Initiative	European Union
Forschungszentrum Jülich	Jülich, Germany
Hokkaido University	Sapporo, Japan
High Performance Computing Center Stuttgart	Stuttgart, Germany
INRIA	France
ETH Zurich	Zurich, Switzerland
Kasetsart University	Bangkok, Thailand
King Abdullah University of Science and Technology	Saudi Arabia
Laboratoire d'Informatique de Paris 6 (LIP6)	Paris, France
Moscow State University	Moscow, Russia
National Institute of Advanced Industrial Science and Technology (AIST)	Tsukuba, Japan
Parallel and HPC Application Software Exchange	Tsukuba, Japan
Prometeus	Mannheim, Germany
Regionales RechenZentrum Erlangen	Erlangen, Germany
RIKEN	Wako, Japan
Rutherford Appleton Laboratory	Oxford, England
Soongsil University	Seoul, South Korea
Technische Universitaet Wien	Vienna, Austria
Tokyo Institute of Technology	Tokyo, Japan
Université Claude Bernard de Lyon	Lyon, France
University of Bordeaux	Bordeaux, France
University of Capetown	Capetown, South Africa
University of Manchester	Manchester, England
University of Paris-Sud	Paris, France
University of Umeå	Umeå, Sweden

LEADERSHIP

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

BDEC BIG DATA AND EXTREME SCALE COMPUTING

BDEC Barcelona

In the past several years, the United States, the European Union, and Japan have each moved aggressively to develop their own plans for achieving exascale computing in the next decade. Such concerted planning by the traditional leaders of HPC speaks eloquently about both the substantial rewards that await the success of such efforts, and about the unprecedented technical obstacles that apparently block the path upward to get there. But while these exascale initiatives, including the International Exascale Software Project (IESP), have understandably focused on the big challenges of exascale for hardware and software architecture, the relatively recent emergence of the phenomena of "Big Data" in a wide variety of scientific fields represents a tectonic shift that is transforming the entire research landscape on which all plans for exascale computing must play out.

The workshop series on Big Data and Extreme-scale Computing (BDEC) marks a distinctly new phase for the work of the IESP community, and is premised on the idea that we must begin to systematically map out and account for the ways in which the major issues associated with Big Data intersect with, impinge upon, and potentially change the national (and international) plans that are now being laid for achieving exascale computing.

In 2015, ICL was instrumental in organizing and staging the third BDEC workshop in Barcelona, Spain. Along with Jack Dongarra, and following through with work they began with the IESP and the first two BDEC meetings in Charleston and Fukuoka, several members of ICL's CITR staff, including Terry Moore, Tracy Rafferty, Sam Crawford, and David Rogers, played essential roles in making the third BDEC workshop a major success. A follow-up BDEC meeting was held in Frankfurt, Germany during the International Supercomputing Conference (ISC 15). This 1-day "reporting meeting" focused on current progress in the US, Europe, and Asia, and opportunities for public engagement.

The next BDEC workshop is in the planning stages and is set for sometime in early 2016. Given the ever increasing emphasis that science, government, and industry continue to place on both big data and extreme-scale computing, this example of ICL's community leadership seems likely to become more and more prominent.

In addition to the development of tools and applications, ICL is regularly engaged in other activities and efforts that include our leadership at conferences and workshops, as well as our teaching and outreach. Having a leadership role in the HPC arena requires that ICL be engaged with the community, and actively share our vision for the future of high performance computing. This section contains some of the activities in which we are participating or have a taken a leadership role.

THE UNIVERSITY OF TENNESSEE KNOXVILLE

CENTER FOR INFORMATION TECHNOLOGY RESEARCH

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

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

CITR's second objective is to grow an interdisciplinary Computational Science program as part of the university curriculum. To this end, CITR helped establish the Interdisciplinary Graduate Minor in Computational Science (IGMCS) to offer UT graduate students an opportunity to acquire the balanced package of knowledge and skills required for today's computationally intensive research methods. CITR is also the sole provider of administrative support for the IGMCS.

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

Computational Science integrates elements that are normally studied in different parts of the traditional curriculum, but which are not fully covered or combined by any one of them. As computational power continues to increase and data storage costs decrease, the potential for new discoveries using Computational Science is greater than ever. And as more academic disciplines begin to realize and exploit the incredible benefits Computational Science provides, the IGMCS program is expected to grow by adding new disciplines, new courses, and new faculty. As of late 2015, there are 18 departments from four UT colleges contributing more than 130 courses to the program.

ICL Group
FALL 2015

PEOPLE

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

As part of an engineering college at a top 50 public research university, we have a responsibility to combine exemplary teaching with cutting-edge research. As such, we regularly employ bright and motivated graduate and undergraduate students. We have been, and will continue to be, very proactive in securing internships and assistantships for highly motivated and hardworking students.

2015 STAFF AND STUDENTS

Ahmad Ahmad
POST DOCTORAL RESEARCH ASSOCIATE

Hartwig Anzt
RESEARCH SCIENTIST I

George Bosilca
RESEARCH DIRECTOR

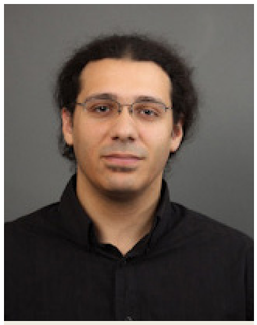

Aurelien Bouteiller
RESEARCH SCIENTIST II

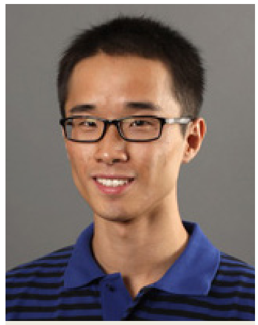

Chongxiao Cao
GRADUATE RESEARCH ASSISTANT

Sam Crawford
INFORMATION SPECIALIST I

Anthony Danalis
RESEARCH SCIENTIST II

Jack Dongarra
UNIVERSITY DISTINGUISHED PROFESSOR

Joe Dorris
GRADUATE RESEARCH ASSISTANT

David Eberius
GRADUATE RESEARCH ASSISTANT

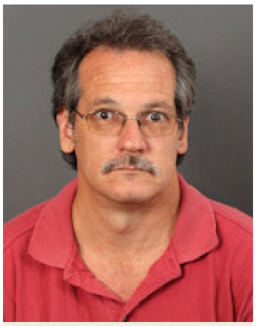

Don Fike
IT SPECIALIST III

Teresa Finchum
ADMINISTRATIVE SPECIALIST II

Mark Gates
RESEARCH SCIENTIST I

Damien Genet
POST DOCTORAL RESEARCH ASSOCIATE

Amina Guermouche
SENIOR RESEARCH ASSOCIATE

Azzam Haidar
RESEARCH SCIENTIST II

Blake Haugen
GRADUATE RESEARCH ASSISTANT

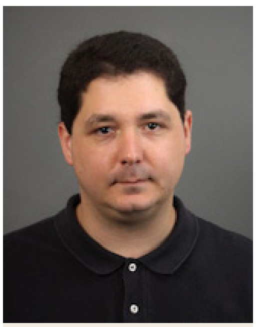

Thomas Herault
RESEARCH SCIENTIST II

Reazul Hoque
GRADUATE RESEARCH ASSISTANT

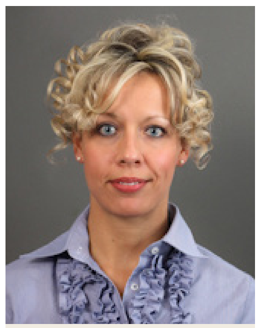

Heike Jagode
RESEARCH DIRECTOR


Khairul Kabir
GRADUATE RESEARCH ASSISTANT


Cindy Knisley
FINANCIAL SPECIALIST

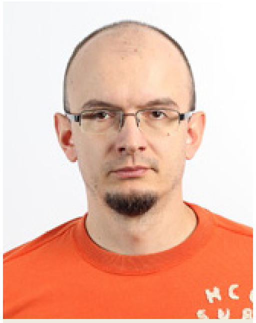
Jakub Kurzak
RESEARCH DIRECTOR


Julie Langou
RESEARCH LEADER


Xi Luo
GRADUATE RESEARCH ASSISTANT


Piotr Luszczek
RESEARCH DIRECTOR


Terry Moore
ASSOCIATE DIRECTOR

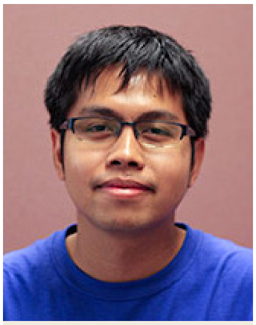
Thananon Patinyasakdikul
GRADUATE RESEARCH ASSISTANT


Tracy Rafferty
PROGRAM MANAGER


Sangamesh Ragate
GRADUATE RESEARCH ASSISTANT


Stephen Richmond
UNDERGRADUATE STUDENT ASSISTANT


Yves Robert
VISITING SCHOLAR


David Rogers
IT SPECIALIST III


Leighanne Sisk
ADMINISTRATIVE SPECIALIST I


Stanimire Tomov
RESEARCH DIRECTOR


Yaohung Tsai
GRADUATE STUDENT ASSISTANT


Phil Vaccaro
GRADUATE RESEARCH ASSISTANT


Wei Wu
GRADUATE RESEARCH ASSISTANT


Ichitaro Yamazaki
RESEARCH SCIENTIST I


Asim YarKhan
RESEARCH SCIENTIST II

2015 VISITORS

Since ICL was founded, we have routinely hosted many visitors, some who stay briefly to give seminars or presentations, and others who remain with us for as long as a year collaborating, teaching, and learning. By collaborating with researchers from around the globe, we are able to leverage an immense array of intellectual resources. For this reason, our list of research collaborators and partners continues to grow. These relationships present enormous opportunities to host and work with top minds within the global HPC community.

Marc Baboulin
INRIA
FRANCE

Michael Barton
UNITED STATES ARMY
RESEARCH LABORATORY (ARL)

Takeshi Fukaya
HOKKAIDO UNIVERSITY

Edmond Chow
GEORGIA TECH

Mathieu Faverge
UNIVERSITY OF BORDEAUX
FRANCE

Mike Guidry
ORNL

Tony Hey

Torsten Hoefler
ETH ZÜRICH

Toshiyuki Imamura
RIKEN AICS

Mike Jantz
EECS

Emmanuel Jeannot
INRIA
FRANCE

Moritz Kreutzer
FRIEDRICH-ALEXANDER
UNIVERSITY
ERLANGEN-NÜRNBERG

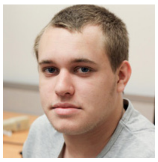

Valentin Le Fevre
ENS DE LYON

Samuel "Sticks" Mabakane
UNIVERSITY OF CAPETOWN

Ian Masliah
UNIVERSITY OF PARIS-SUD

Audris Mockus
EECS

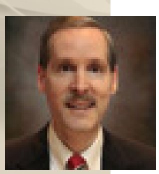

Bob Muenchen
UTK

Manish Parashar
RUTGERS

Kalyan Perumalla
ORNL

Eduardo Ponce
EECS

Nageswara Rao
ORNL

Joseph Schuchart
TU DRESDEN

Pierre Sens
LIP6

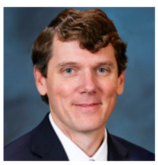

Chad Steed
ORNL

Ed Valeev
VIRGINIA TECH

ALUMNI

Carolyn Aebischer 1990-1993
Bivek Agrawal 2004-2006
Sudesh Agrawal 2004-2006
Emmanuel Agullo 2009
Ahmad Abdelfattah Ahmad 2012
Jennifer Allgeyer 1993
Wes Alvaro 2007-2011
Ed Anderson 1989-1991
Daniel Andrzejewski 2007
Thara Angskun 2003-2007
Papa Arkhurst 2003
Dorian Arnold 1996-2001
Cedric Augonnet 2010
Marc Baboulin 2008
Zhaojun Bai 1990-1992
Ashwin Balakrishnan 2001-2002
Richard Barrett 1992-1994
Alex Bassi 2000-2001
David Battle 1990-1992
Micah Beck 2000-2001
Daniel Becker 2007
Dulceneia Becker 2010-2012
Adam Beguelin 1991
Annamaria Benzoni 1991
Tom Berry 1991
Vincent Berthoux 2010
Scott Betts 1997-1998
Nikhil Bhatia 2003-2005
Noel Black 2002-2003
Laura Black 1996
Susan Blackford 1989-2001
Wesley Bland 2008-2013
Kartheek Bodanki 2009
David Bolt 1991
Fernando Bond 1999-2000
Carolyn Bowers 1992
Barry Britt 2007-2009
Randy Brown 1997-1999
Cynthia Browne 2005
Murray Browne 1998-1999
Bonnie Browne 2011-2012
Antonin Bukovsky 1998-2003
Greg Bunch 1995
Alfredo Buttari 2004-2007
Giuseppe Bruno 2001
Anthony Canino 2012
Domingo Gimenez Canovas 2001
Henri Casanova 1995-1998
Cedric Castagnede 2012
Ramkrishna Chakrabarty 2005
Sharon Chambers 1998-2000
Zizhong Chen 2001-2006
Jaeyoung Choi 1994-1995
Wahid Chrabakh 1999
Eric Clarkson 1998
Andy Cleary 1995-1997
Michelle Clinard 1989-1991
Vincent Cohen-Addad 2012
Matthias Colin 2004
Charles Collins 2012
Stephanie Cooper 2011-2013
Tom Cortese 2002-2009
Camille Coti 2007
Jason Cox 1993-1997
David Cronk 1999 - 2010
Javier Cuenca 2003
Manoel Cunha 2006
Yuanshun Dai 2007-2013
Cricket Deane 1998-1999
Remi Delmas 2006
Frederic Desprez 1994-1995
Jin Ding 2003
Jun Ding 2001-2003
Ying Ding 2000-2001
Martin Do 1993-1994
Simplice Donfack 2014
Leon Dong 2000-2001
Tingxing Dong 2010-2015
Nick Dongarra 2000
David Doolin 1997
Andrew Downey 1998-2003
Mary Drake 1989-1992
Julio Driggs 2002-2004
Brian Drum 2001-2004
Peng Du 2005-2012
Eduardo Echavarria 2005
Victor Eijkhout 1992-2005
Brett Ellis 1995-2005
Shawn Ericson 2004
Zachary Eyler-Walker 1997-1998
Lisa Ezzell 2003-2004
Christoph Fabianek 2003
Graham Fagg 1996-2006
Mathieu Faverge 2010-2012
Shengzhog Feng 2005-2006
Salvatore Filippone 2004
Anna Finchum 2010
Mike Finger 1997
Markus Fischer 1997-1998
Len Freeman 2009
Xiaoquan Fu 2003-2004
Erika Fuentes 2003-2007
Karl Fuerlinger 2006-2008
Megan Fuller 2006
Edgar Gabriel 2003-2004
Tracy Gangwer 1992-1993
Lynn Gangwer 2000-2001
Nathan Garner 2001-2006
Kelley Garner 1998
Tina Garrison 1991
Adriana Garties 2011
Peter Gaultney 2011-2014
Christoph Geile 2008
Jean Patrick Gelas 2001
Boris Gelfend 1993
Jonathan Gettler 1996
Eric Greaser 1993
Stan Green 1992-1996
Alice Gregory 2004-2006
Jason Gurley 1997-1998
Bilel Hadri 2008-2009
Hunter Hagewood 2000-2001
Christian Halloy 1996-1997
Sven Hammarling 1996-1997
J. Mike Hammond 1994-1995
Satomi Hasegawa 1995-1996
Hidehiko Hasegawa 1995-1996
Chris Hastings 1996
David Henderson 1999-2001
Greg Henry 1996
Julien Herrmann 2011-2012
Holly Hicks 1993-1994
Alexandra Hicks-Hardiman 2009
Sid Hill 1996-1998
Tomoyuki Hiroyasu 2002-2003
George Ho 1998-2000
Josh Hoffman 2008-2010
Jeff Horner 1995-1999
Mitch Horton 2010-2012
Yan Huang 2000-2001
Aurelie Hurault 2009
Chris Hurt 2002
Paul Jacobs 1992-1995
Emmanuel Jeannot 2001-2006
Weizhong Ji 1999-2000
Yulu Jia 2011-2015
Weicheng Jiang 1992-1995
Song Jin 1997-1998
Patrick Johansson 2001
Aral Johnson 2009
Matt Johnson 2011-2013
Sean Jolly 1997-1998
Jan Jones 1992-2008
Kim Jones 1996-1997
Vijay Joshi 2011-2013
Venkata Kakani 2007
Ajay Kalhan 1995
Balajee Kannan 2001
Madhuri Kasam 2007-2008
Kiran Kasichayanula 2010-2012
Ajay Katta 2010
David Katz 2002
Joshua Kelly 2000-2001
Supriya Kilambi 2008

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

Myung Ho Kim 2005-2006
Youngbae Kim 1992-1996
Jenya Kirshtein 2008
Michael Kolatis 1993-1996
Chandra Krintz 1999-2001
Tilman Kuestner 2010
Krerkchai Kusolchu 2010
Coire Kyle 2005
Amanda Laake 2003-2004
Xavier Lacoste 2012
Julien Langou 2003-2006
Jeff Larkin 2003-2005
Brian LaRose 1990-1992
Frank Lauer 2010
DongWoo Lee 2000-2002
Tracy Lee 1996-2012
Klaudia Leja 2008
Pierre Lemarinier 2008-2010
Todd Letsche 1993-1994
Sharon Lewis 1992-1995
Xiang Li 2001
Yinan Li 2006-2008
Weiran Li 2002
Chaoyang Liu 2000
Kevin London 1996-2005
Matt Longley 1999
Florent Lopez 2014
Hatem Ltaief 2008-2011
Daniel Lucio 2008
Richard Luczak 2000-2001
Teng Ma 2006-2012
Robert Manchek 1990-1996
Gabriel Marin 2013-2014
Tushti Marwah 2004
Theo Mary 2014
Donald McCasland 1994
Paul McMahan 1994-2000
Eric Meek 2003-2006
James Meyering 1991-1992
Jeremy Millar 1998-2002
Michelle Miller 1999-2003

Cindy Mitchell 2001-2002
Stuart Monty 1993
Eric Moore 2000
Keith Moore 1987-2007
Shirley Moore 1993-2012
Robert Morgan 1990-1991
Kishan Motheramgari 1997
Steven Moulton 1991-1993
Phil Mucci 1994-2004
Daichi Mukunoki 2012
Matthew Nabity 2008
Shankar Narasimhaswami 2004-2005
Rajib Nath 2008-2010
Fernando Navarro 2009
Donnie Newell 2010
John Nelson 2011-2013
Peter Newton 1994-1995
Jonas Nilsson 2001
Jakob Oestergaard 2000
Caroline Papadopoulos 1997-1998
Leelinda Parker 2002
Thibault Parpaite 2014
Dilip Patlolla 2007-2008
Andy Pearson 1989-1991
Paul Peltz 2003-2013
Theresa Pepin 1994
Antoine Petitet 1993-2001
Peter Pham 2012
Gregoire Pichon 2014
Vlado Pjesivac 2008
Jelena Pjesivac-Grbovic 2003-2007
James S. Plank 1991-1992
Ciara Proctor 2008
Tim Poore 2009
Roldan Pozo 1992-1994
Farzona Pulatova 2005-2006
Martin Quinson 2001
Tammy Race 1999-2001
James Ralph 2006-2014
Ganapathy Raman 1998-2000
Kamesh Ramani 2003

Mei Ran 1999-2004
Arun Rattan 1997
Sheri Reagan 1995-1996
Mike Reynolds 1994
George Rhinehart 2012
Jon Richardson 1990-1991
Ken Roche 1999-2004
Andrew Rogers 1997-1999
Tom Rothrock 1997-1998
Tom Rowan 1993-1997
Narapat (Ohm) Saengpatsa 2011
Kiran Sagi 2001-2005
Evelyn Sams 1998-1999
Ken Schwartz 1992-1993
Keith Seymour 1994-2009
Farial Shahnaz 2001
Brian Sheely 2009-2010
Zhiao Shi 2001-2007
Sergei Shinkarev 2005-2007
Majed Sidani 1991-1992
Shilpa Singhal 1996-1998
Matt Skinner 2008
Jonte Smith 2008
Peter Soendergaard 2000
Raffaele Solca 2012
Gwang Son 2007-2009
Fengguang Song 2003-2012
Thomas Spencer 1999-2001
Erich Strohmaier 1995-2001
Xiaobai Sun 1995
Martin Swany 1996-1999
Daisuke Takahashi 2002
Judi Talley 1993-1999
Ronald Tam 2009
Chunyan Tang 2014-2015
Yuan Tang 2005-2006
Yusuke Tanimura 2003
Keita Teranishi 1998
Dan Terpstra 2001-2014
Joe Thomas 2002-2009
John Thurman 1998-1999

Francoise Tisseur 1997
Jude Toth 1993-1994
Bernard Tourancheau 1993-1994
Malcolm Truss 2008
Lauren Vaca 2004
Sathish Vadhiyar 1999-2003
Robert van de Geijn 1990-1991
Chad Vawter 1995
Eugene Vecharynski 2008
Scott Venckus 1993-1995
Antoine Vernois 2004
Reed Wade 1990-1996
Michael Walters 2001-2005
Mike Waltz 1999
Robert Waltz 1990-1991
Jerzy Wasniewski 2000
Vince Weaver 2010-2012
Scott Wells 1997-2010
David West 1990-1992
R. Clint Whaley 1991-2001
Jody Whisnant 1997-1998
James White 1999
Scotti Whitmire 1995-1996
Susan Wo 2000-2001
Felix Wolf 2003-2005
Jiayi Wu 2004-2007
Qiu Xia 2004-2005
Tinghua Xu 1998-2000
Tao Yang 1999
Erlin Yao 2012-2013
Kevin Ye 2015
Jin Yi 2009-2010
Haihang You 2002-2009
Lamia Youseff 2007
Brian Zachary 2009-2010
Omar Zenati 2012
Yuanlei Zhang 2001-2005
Junlong Zhao 2002
Yong Zheng 2001
Luke Zhou 2000-2001
Min Zhou 2002-2004

ACKNOWLEDGMENTS

For 25 years, our knowledge and hard work have earned the trust and support of many agencies and organizations that have funded, and continue to fund, our efforts. Without them we simply would not be able to conduct cutting-edge research.

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

In addition to the support of the federal government, we have solicited strong support from private industry, which has also played a significant role in our success and growth. Some organizations have targeted specific ICL projects, while others have made contributions to our work that are more general and open-ended. We gratefully acknowledge the following vendors for their generosity and support:

THE UNIVERSITY OF TENNESSEE
KNOXVILLE

CENTER FOR INFORMATION TECHNOLOGY RESEARCH

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

Financial support for CITR is provided by the University of Tennessee.

INNOVATIVE COMPUTING LABORATORY

Suite 203 Claxton
1122 Volunteer Blvd
Knoxville, TN 37996-3450

INNOVATIVE COMPUTING LABORATORY

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