ICL UT

INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY of TENNESSEE

2011/2012 REPORT

INNOVATIVE COMPUTING LABORATORY **2011/2012** REPORT

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

EEO/TITLE IX/AA/SECTION 504 STATEMENT

All qualified applicants will receive equal consideration for employment and admissions without regard to race, color, national origin, religion, sex, pregnancy, marital status, sexual orientation, gender identity, age, physical or mental disability, or covered veteran status.

Eligibility and other terms and conditions of employment benefits at The University of Tennessee are governed by laws and regulations of the State of Tennessee, and this non-discrimination statement is intended to be consistent with those laws and regulations.

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

Inquiries and charges of violation of Title VI (race, color, national origin), Title IX (sex), Section 504 (disability), ADA (disability), Age Discrimination in Employment Act (age), sexual orientation, or veteran status should be directed to the Office of Equity and Diversity (OED), 1840 Melrose Avenue, Knoxville, TN 37996-3560, telephone (865) 974-2498 (V/TTY available) or 974-2440. Requests for accommodation of a disability should be directed to the ADA Coordinator at the Office of Equity and Diversity.

INNOVATIVE COMPUTING LABORATORY
2011/2012 REPORT

FROM THE DIRECTOR

IN 2011, the Innovative Computing Laboratory is celebrating 22 years of leadership in enabling technologies for high performance computing. Looking back over the 22-year period, the evolution and growth of the technology for computing has been truly astonishing. In an environment where technology is continuously changing, ICL cannot afford to stand still. In 1989 the speed of a supercomputer was measured in gigaflops and in gigabytes. Today our measures are petaflops for speed and petabytes for memory, a million-fold increase over the standards of two decades ago. The research that ICL has undertaken in the past decade has followed a natural progression and growth from our original tread of numerical linear algebra to performance evaluation, to software repositories, and to distributed computing.

This is an exciting time for computing as we begin the journey on the road to exascale computing. 'Going to the exascale' will mean radical changes in computing architecture, software, and algorithms – basically, vastly increasing the levels of parallelism to the point of billions of threads working in tandem – which will force radical changes in how hardware is designed and how we go about solving problems. There are many computational and technical challenges ahead that must be overcome. The challenges are great, different than the current set of challenges, and exciting research problems await us. ICL's research agenda has never been stagnant; we have always taken leadership roles in enabling technologies for high performance computing.

The Innovative Computing Laboratory is prepared to address some of the most important computational scientific issues of our time. Our plans for the future are founded on our accomplishments as well as our vision. That

vision challenges us to be a world leader in enabling technologies and software for scientific computing. We have been and will continue to be providers of high performance tools to tackle science's most challenging problems, and play a major role in the development of standards for scientific computing in general.

We are building from a firm foundation. Over the past 22 years, we have developed robust research projects, attracted some of the best and brightest students and researchers, and created leading-edge research programs. The ICL staff's ongoing ability to apply the latest technologies to provide advanced services and solutions for the scientific computing community underscores the ICL's leadership role. Standards and efforts such as PVM, MPI, LAPACK, ScaLAPACK, BLAS, ATLAS, Netlib, NHSE, TOP500, PAPI, NetSolve, Open-MPI, FT-MPI, the HPC Challenge Benchmark, and LINPACK Benchmark have all left their mark on the scientific community. We can be proud of the recognition and use our tools receive. We are continuing these efforts with IESP, PLASMA, MAGMA, DAGuE, and MuMMI as well as other innovative computing projects.

We continue to grow in terms of the resources we have at our disposal. We have ongoing efforts to strengthen our organization and to ensure the proper balance and integration of research and projects. The pace of change will continue to accelerate in the coming years.

Advancing to the next stage of growth for computational simulation and modeling will require us to solve basic research problems in Computer Science and Applied Mathematics at the same time as we create and promulgate a new paradigm for the development of scientific software. To make progress on both fronts simultaneously will require a level of sustained, interdisciplinary collaboration among the core research communities that, in the past, has only been achieved by forming and supporting research centers dedicated to such a common purpose. I believe that the time has come for the leaders of the Computational Science movement to focus their energies on creating such software research centers to carry out this indispensable part of the mission. I have every confidence that our community stands ready to step up again to this momentous new effort.

This is truly a time of great excitement in the design of software and algorithms for the next generation, perhaps a once in a lifetime opportunity, and we will be part of that continuing evolution of the high performance computing ecology.

During these exciting times, I am grateful to our sponsors for their continued endorsement of our efforts. My special thanks and congratulations go to the ICL staff and students for their skill, dedication, and tireless efforts in making the ICL one of the best centers in the world for enabling technologies.

JACK DONGARRA
DIRECTOR OF ICL

INTRODUCTION

IN 1989, The Innovative Computing Laboratory (ICL) was founded by Dr. Jack Dongarra, who came to the University of Tennessee (UT) from Argonne National Laboratory. At that time, Dr. Dongarra received a dual appointment as a Distinguished Professor in the Computer Science Department at UT and as a Distinguished Scientist at nearby Oak Ridge National Laboratory (ORNL). Since that date, ICL has grown from two grad students and two Post-docs to a fully functional research laboratory, with a staff of nearly 50 researchers, students, and administrators.

ICL began in a small office inside one of UT's oldest academic buildings, but now occupies a large portion of a 70,000 sq. ft. wing of the newer Claxton building located at the heart of the Knoxville campus. In 2007, ICL and our UT colleagues in Computer Science joined the faculty of Computer and Electrical Engineering to form the Electrical Engineering and Computer Science (EECS) department in the College of Engineering.

As one of UT's oldest and largest research laboratories, ICL has been unwavering in its mission since day one: be a world leader in enabling technologies and software for scientific computing. In keeping with this effort, we continue to provide leading edge tools to tackle science's most challenging high performance computing problems, and we play a major role in the development of standards for scientific computing in general.

Our commitment to excellence has been one of the keys to our success as we continually strive to make a substantial impact in the high performance computing community. As a result, we continue to lead the way as one of the most respected academic research centers in the world.

KEY ICL RESEARCH

1989 PVM, Level 3 BLAS
1992 LAPACK, BLACS
1993 TOP500
1994 MPI
1995 ScaLAPACK
1997 NetSolve, ATLAS, RIB
1999 PAPI, HARNESS
2000 HPL
2002 FT-MPI, GCO
2003 LAPACK for Clusters, HPC Challenge
2006 Open MPI, FT-LA
2008 PLASMA, MAGMA
2009 Blackjack, MuMMI
2010 DAGuE
2011 PULSAR

SELECTED ICL ALUMNI

1991-1992 Adam Beguelin — SENSR.NET
1995-1996 Sven Hammarling — NUMERICAL ALGORITHMS GROUP
2004-2007 Alfredo Buttari — CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE
1991-2001 Clint Whaley — UNIVERSITY OF TEXAS, SAN ANTONIO
1995-1998 Henri Casanova — UNIVERSITY OF HAWAII, MANOA
2003-2004 Edgar Gabriel — UNIVERSITY OF HOUSTON
1989-2001 Susan Blackford — MYRICOM
1994-2009 Keith Seymour — NATIONAL INSTITUTE FOR COMPUTATIONAL SCIENCES (NICS)
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
2003-2006 Julien Langou — UNIVERSITY OF COLORADO AT DENVER
1989-1991 Ed Anderson — EPA
1996-2006 Graham Fagg — MICROSOFT
1991-1994 Roldan Pozo — NIST
1995-2001 Erich Strohmaier — LAWRENCE BERKELEY NATIONAL LABORATORY
2003-2006 Felix Wolf — JÜLICH SUPERCOMPUTING CENTRE
1992-1994 Richard Barrett — ORNL
1996-1997 Yves Robert — ENS-LYON, FRANCE
2000-2001 Reed Wade — WETA DIGITAL
2009 Emmanuel Agullo — INRIA, FRANCE
1989-1990 Zhaojun Bai — UNIVERSITY OF CALIFORNIA, DAVIS
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
2008-2011 Hatem Ltaief — KAUST, SAUDI ARABIA
1993-1994 Bernard Tourancheau — UNIVERSITY OF LYON, FRANCE
1997 Francois Tisseur — MANCHESTER UNIVERSITY, ENGLAND
1994 Frederic Desprez — ENS-LYON, FRANCE
1999-2004 Ken Roche — UNIVERSITY OF WASHINGTON

RESEARCH

INCREASED EFFORTS to keep pace with the evolution in HPC hardware and software represent 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 more than 20 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 of the HPC community, which includes a focus on algorithms and libraries for multicore and hybrid computing. We also now include work in high performance parallel and distributed computing, with efforts focused on message passing and fault tolerance. As we have gained a solid understanding of the challenges presented in these domains, we have further expanded our research to include work in performance analysis and benchmarking for high-end computers.

Demonstrating the range and diversity of our research, we will be engaged in more than 20 significant research projects during 2011-2012 across our main areas of focus. On the following pages, we provide brief summaries of some of our efforts in these research areas. For more detailed information about our research, visit our website – http://icl.eecs.utk.edu/.

NUMERICAL LINEAR ALGEBRA

EASI

Extreme-scale Algorithms and Software Institute

The mission of the Extreme-scale Algorithms and Software Institute (EASI) 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 EASI 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, University of Illinois, 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.

FT-LA

Fault Tolerant Linear Algebra

http://icl.eecs.utk.edu/ft-la/

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 distributed 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 (FT-MPI) middleware and, more recently, an FT-LA library that will efficiently handle several process failures. The project team has also integrated FT-LA in the CIFTS (Coordinated Infrastructure for Fault Tolerant Systems) environment to provide better communication and fault management between the system's software components and scientific applications. Future work in this area involves the development of scalable fault-tolerant, one-sided (Cholesky, LU, and QR) and two-sided (Hessenberg, tri-diagonalization, and bi-diagonalization) factorizations, following the ABFT principles.

Keeneland

Enabling Heterogeneous Computing for the Open Science Community

http://keeneland.gatech.edu/

Keeneland is a five-year, $12 million cyberinfrastructure project, funded under the NSF's Track 2D program, designed to bring emerging hardware architectures to the open science community. ICL is partnering with project leader Georgia Tech, as well as Oak Ridge National Laboratory, UTK's National Institute for Computational Sciences, Hewlett-Packard, and NVIDIA, to develop and deploy Keeneland's innovative and experimental system.

As part of our contribution, ICL performed education and outreach activities, developed numerical libraries to leverage the power of NVIDIA's CUDA-based GPUs used in the Keeneland machine, and teamed up with early adopters to map their applications to the Keeneland architecture. In 2010, the Keeneland project procured and deployed its initial delivery system (KIDS): a 201 Teraflop, 120-node HP SL390 system with 240 Intel Xeon CPUs and 360 NVIDIA Fermi graphics processors, connected by an InfiniBand QDR network. The Keeneland team hopes to procure and deploy its full scale system in 2012.

LAPACK

Linear Algebra PACKage

http://www.netlib.org/lapack/

ScaLAPACK

Scalable LAPACK

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 3.3.0, released in November 2010, includes LAPACKE, a native C interface for LAPACK, developed in collaboration with Intel, which provides NAN check and automatic workspace allocation. LAPACK 3.3.0 also includes new level-3 BLAS symmetric indefinite solve and symmetric indefinite inversion routines, along with complete CS decomposition routines. The next major release, LAPACK 3.4.0, is expected in November 2011. ScaLAPACK 2.0, which will include the MRRR algorithm, is also scheduled for release in late 2011.

MAGMA

Matrix Algebra on GPU and Multicore Architectures

http://icl.eecs.utk.edu/magma/

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 to deliver the fastest possible time to accurate solution within given energy constraints.

MAGMA 1.0 features top performance and high accuracy LAPACK compliant routines for multicore CPUs enhanced with NVIDIA GPUs. The MAGMA 1.0 release includes more than 150 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. MAGMA provides multiple precision arithmetic support (S/D/C/Z, including mixed-precision). All algorithms are hybrid, using both multicore CPUs and GPUs.

PLASMA

Parallel Linear Algebra Software for Multicore Architectures

http://icl.eecs.utk.edu/plasma/

The Parallel Linear Algebra Software for Multicore Architectures (PLASMA) package, ICL's flagship project for multicore and many-core computing, is designed to deliver high performance from homogeneous multi-socket multicore systems by combining state-of-the-art solutions in algorithms, scheduling, and software engineering.

PLASMA includes uniform and mixed precision routines for solving linear systems and least square problems, and now offers separate routines for obtaining singular value decomposition, and solving symmetric eigenvalue and generalized eigenvalue problems. PLASMA also includes routines for fast parallel processing of very tall and narrow matrices, a full set of Level 3 BLAS operations, and fast routines for explicitly forming an inverse of a matrix. PLASMA 2.4.2, released in September 2011, contains several bug fixes, a new version of QUARK, and is available with an array of user-friendly online resources.

PULSAR

Parallel Unified Linear Algebra with Systolic Arrays

The objective of the Parallel Unified Linear Algebra with Systolic Arrays (PULSAR) project is to address the challenges of extreme scale computing by applying the dataflow principles of Systolic Array architectures. PULSAR, recently funded by the NSF, targets billion-fold parallelism, where the volume of communication is important, but the locality of communication is even more critical. At the same time, it is paramount to provide a simple execution model to aid in efficient programming.

PULSAR's solution to these problems is a Virtual Systolic Array (VSA) architecture, where multidimensional virtual systolic arrays are designed for various scientific workloads and successively mapped to the hardware architecture by a virtualization layer with a substantial runtime component. Initially PULSAR will deliver dense linear algebra codes, produced by mapping tile algorithms for matrix factorizations to VSAs, and then mapping those VSAs to the interconnection topology of the largest distributed memory machines (MPPs) currently at our disposal.

QUARK

QUeuing And Runtime for Kernels

http://icl.eecs.utk.edu/quark/

QUeuing And Runtime for Kernels (QUARK) provides a library that enables the dynamic 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 an asynchronous, dynamic fashion in order to achieve a high utilization of the available resources.

QUARK is designed to be easy to use, is intended to scale to large numbers of cores, and should enable the efficient expression and implementation of complex algorithms. The driving application behind the development of QUARK is the PLASMA linear algebra library, and the QUARK runtime contains several optimizations inspired by the algorithms in PLASMA. An early release of QUARK is available, and includes a well-stressed and robust implementation and an initial User's Guide and Reference Guide. Additional documentation will be provided in future releases.

PERFORMANCE EVALUATION AND BENCHMARKING

Blackjack

Compiler Metrics and Evaluation

http://icl.eecs.utk.edu/blackjack/

The Blackjack project is developing metrics and tools for evaluating compilers for scientific computing. Modern computing architectures change rapidly and exhibit high levels of complexity and heterogeneity. Developing compilers that boost productivity, while producing efficient optimized code for these rapidly evolving targets, is a difficult challenge. Blackjack evaluates compilers by implementing relevant micro-benchmarks and using representative applications to test and analyze the productivity, correctness, and performance of multiple commercially available and freely available compiler systems.

As part of the Blackjack project, the team has been developing a system benchmark suite, called BlackjackBench, which can automatically characterize target architectures in a rigorous and systematic manner. Given this potential, BlackjackBench could be used by compilers to adapt their optimizations to different underlying platforms.

CScADS

Center for Scalable Application Development Software

http://cscads.rice.edu/

The Center for Scalable Application Development Software (CScADS) for Advanced Architectures was created at Rice University to facilitate the scalability of applications to the petascale and beyond, while fostering the development of new tools by the computer science community through the support of common software infrastructures and standards.

CScADS is a collaborative effort between Argonne National Laboratory, Rice University, University of California at Berkeley, University of Wisconsin, and the University of Tennessee (ICL). ICL's effort in this project focuses on re-engineering numerical libraries for future HPC systems utilizing multicore processors. This work explores the use of multithreading to tolerate synchronization latency in the context of matrix factorization. The model relies on dynamic, dataflow-driven execution models and avoids both global synchronization and the implicit point-to-point synchronization of send/receive style message passing.

FutureGrid

Distributed Computing Infrastructure

https://portal.futuregrid.org/

FutureGrid is a distributed computing infrastructure that uses the nation's high performance research networks to create a test-bed for developing new approaches to parallel, grid, and cloud computing. FutureGrid uses virtualization technology to create a cloud computing environment that applications can access and utilize in a uniform way. FutureGrid partners are deploying high performance computing clusters at their sites and connecting them to the NSF's newest national cyberinfrastructure for scientific research, XSEDE — the Extreme Science and Engineering Discovery Environment.

Under leadership from Indiana University, the FutureGrid team includes the University of Tennessee (ICL), Purdue University, University of California San Diego, University of Chicago/Argonne National Labs, University of Florida, University of Southern California, University of Virginia, and Technische Universität Dresden. ICL is contributing in the areas of performance measurement and application benchmarking.

HPCC

HPC Challenge Benchmark

http://icl.eecs.utk.edu/hpcc/

The HPC Challenge (HPCC) benchmark suite is designed to assess the bounds of performance on many real-world applications. Included in the benchmark suite are tests for sustained floating point operations, memory bandwidth, rate of random memory updates, interconnect latency, and interconnect bandwidth. The main factor that differentiates the various components of the suite is the memory access patterns that, in a meaningful way, span the temporal and spatial locality space.

Each year, the HPCC competition features contestants who submit performance numbers from the world's largest supercomputer installations and other implementations that use a vast array of parallel programming environments. Results are announced at the annual Supercomputing Conference, and are available to the public through the HPCC website to help track the progress of both the high-end computing arena and the commodity hardware segment.

HPL

High Performance LINPACK

http://icl.eecs.utk.edu/hpl/

The High Performance LINPACK (HPL) benchmark is a software package that solves a (random) 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 TOP500 list of the world's most powerful supercomputers.

HPL 2.0 includes several major bug fixes and accuracy enhancements based on user feedback. The major focus of HPL 2.0 is to improve the accuracy of reported benchmark results, and ensure scalability of the code on large supercomputer installations. Development continues on a time-limited version of HPL to shorten the time required for tuning and running the benchmark. In 2011, the LINPACK benchmark app for iOS achieved performance of over 1 gigaflops on an Apple iPad 2.

MuMMI

Multiple Metrics Modeling Infrastructure

http://icl.eecs.utk.edu/mummi/

The Multiple Metrics Modeling Infrastructure (MuMMI) project is developing a framework to facilitate systematic measurement, modeling, and prediction of performance, power consumption, and performance-power tradeoffs for applications running on multicore systems. MuMMI combines UTK's PAPI hardware performance monitoring capabilities with Texas A&M's Prophesy performance modeling interface and Virginia Tech's Power-Pack power-performance measurement and analysis system.

PAPI has been integrated with Power-Pack and Prophesy, and performance and power consumption data have been collected for a range of benchmarks and applications running on multicore systems. Recent enhancements to the PAPI project also enable the user to define high-level modeling metrics and map them to underlying hardware events and characteristics.

PAPI

Performance API

http://icl.eecs.utk.edu/papi/

The Performance API (PAPI) has become the *de facto* standard within the HPC community for providing access to the hardware performance counters found on modern high performance computing 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. Architecturally, PAPI provides simultaneous access to both on-processor and off-processor counters and sensors. PAPI continues to be ported to the architectures of greatest interest to the high performance computing community, including heterogeneous computing systems and virtual computing environments.

The latest version of PAPI, called Component PAPI, or PAPI-C, supports components for network counters and system health monitoring, as well as disk subsystems. Virtual PAPI, or PAPI-V—currently in development and supported by NSF and VMware—is intended to provide performance measurement standards in virtual environments, which are common in cloud computing. Prompted by the growing trend in heterogeneous computing, the PAPI team continues to work closely with NVIDIA in the ongoing development of a PAPI CUDA component that can monitor performance counters in a CUDA system at the kernel level to provide insights into the optimal performance of CUDA code.

PERI

Performance Engineering Research Institute

http://www.peri-scidac.org/

The Performance Engineering Research Institute (PERI) is conducting performance research designed to make the transition to petascale systems smoother, so that researchers can benefit quickly from these ultra-fast machines. The effort involves performance modeling, development of an automatic tuning system, and application engagement.

PERI is a collaborative effort between Argonne National Laboratory, Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, Oak Ridge National Laboratory, Rice University, University of California at San Diego, University of Maryland, University of North Carolina, University of Southern California, University of Utah, and the University of Tennessee (ICL). ICL is drawing upon its previous experience with empirical auto-tuning methodologies for numerical libraries to help generalize these methodologies to auto-tune performance critical portions of important scientific applications.

POINT

Productivity from Open INtegrated Tools

http://nic.uoregon.edu/point/

The Productivity from Open INtegrated Tools (POINT) project is integrating, hardening, and deploying an open, portable, robust performance tools environment for the NSF-funded high performance computing centers. Entry points for the tools are available for users of various levels of expertise, and the project has a comprehensive outreach and training component.

POINT is a collaborative effort between the University of Oregon, National Center for Supercomputing Applications, Pittsburgh Supercomputing Center, and the University of Tennessee (ICL). ICL is using this opportunity to extend the PAPI tool set, while ensuring continued and enhanced interoperability with the other tools in the POINT tool suite.

SUPER

Institute for Sustained Performance, Energy, and Resilience

http://super-scidac.org/

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 auto-tuning technology on 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. ICL's contributions to the SUPER project are in the areas of performance measurement, analysis, communication modeling, auto-tuning, resilience, and application engagement.

TOP500

Supercomputer Sites

http://top500.org/

Since 1993, a ranking of the 500 fastest computers in the world has been compiled biannually with published results released in June and November. Each machine on the TOP500 is ranked based on performance results from running the numerically intensive High Performance LINPACK (HPL) benchmark developed by ICL. While other benchmarks, including HPCC, have been developed to measure performance of HPC systems, the TOP500 still relies on the HPL benchmark and remains the *de facto* ranking relied upon by commercial, industrial, government, and academic institutions.

ICL continues to partner with NERSC/Lawrence Berkeley National Laboratory and the University of Mannheim, Germany to produce the TOP500 rankings. As of June 2011, the Japanese K Computer, built by Fujitsu, is the fastest supercomputer in the world, and is capable of 8.162 petaflops. Housed at the RIKEN Advanced Institute for Computational Science (AICS), the K Computer is more powerful than the next five systems on the TOP500 list combined. A list of the current TOP500 rankings, along with an archive of charts, statistics, and news, is available on the TOP500 website.

DISTRIBUTED COMPUTING

CIFTS

Coordinated Infrastructure for Fault Tolerant Systems

http://mcs.anl.gov/research/cifts/

The Coordinated Infrastructure for Fault Tolerant Systems (CIFTS) project is a multi-institution effort to enable collaboration between all levels of the HPC software stack, from the operating system to the application. Although many software components have the capability to recover from disruptive failures in modern HPC systems, more often than not, the lack of coordination leads to confusion and contradictory reactions from different entities, preventing true fault tolerance.

ICL's focus is to enable sturdy and collaborative fault tolerant linear algebra libraries and software. Partnering with us in this effort are Argonne National Laboratory, Lawrence Berkeley National Laboratory, Oak Ridge National Laboratory, Indiana University, and Ohio State University.

DAGuE

Directed Acyclic Graph Unified Environment

http://icl.eecs.utk.edu/dague/

The Directed Acyclic Graph Unified Environment (DAGuE) is a generic framework for architecture-aware scheduling and management of micro-tasks on distributed many-core heterogeneous architectures. Applications are represented as a Direct Acyclic Graph of tasks (DAG) with labeled edges designating data dependencies. DAGs are represented in a compact, problem size independent format, which can be queried on demand to discover data dependencies in a totally distributed fashion.

DAGuE assigns computation threads to the cores, overlaps communications and computations, and uses a dynamic, fully distributed scheduler based on cache awareness, data-locality, and task priority. DAGuE includes a set of tools to generate the DAGs and integrate them in legacy codes, a runtime library to schedule the micro-tasks, and tools to evaluate and visualize the efficiency of the scheduling. Many dense Linear Algebra computational kernels have been re-implemented using DAGuE, enabling better performance on distributed many-core systems.

FT-MPI

Fault Tolerant Message Passing Interface

http://icl.eecs.utk.edu/ftmpi/

Open MPI

High Performance Message Passing Library

http://icl.eecs.utk.edu/open-mpi/

Fault Tolerant Message Passing Interface (FT-MPI) is a leading edge, full MPI 1.2 specification implementation that provides process level fault tolerance at the MPI API level, which allows for flexible new models of fault tolerance and recovery that were previously impossible. Since the release of the FT-MPI runtime library, research in FT-MPI has mainly centered on system level software and environment management in order to enhance and improve its performance, robustness, and scalability. This research covers diverse topics, from self-healing networks, to the fundamental understanding and modeling of group communications in a fault enabled environment.

Some fault tolerance mechanisms designed in the context of FT-MPI are currently under consideration by the MPI Forum for inclusion in the next version of the MPI standard (MPI 3.0).

Many features from FT-MPI, such as point-to-point messaging, tuned collective communication algorithms, and the heterogeneous data-type engine, have been integrated into the open source production quality MPI implementation known as Open MPI, which is part of a collaborative effort involving several institutions, including ICL. In addition, our efforts in the context of Open MPI have significantly improved its scalability, performance on many-core environments, and architecture aware capabilities, making it ready for the next generation exascale challenges.

G8 ECS

Enabling Climate Simulation at Extreme Scale

The objective of the G8 Enabling Climate Simulation at Extreme Scale (ECS) project is to investigate how climate scientists can efficiently run climate simulations on future Exascale systems. Exascale supercomputers will appear in 2018–2020 featuring a hierarchical design, and will utilize hundreds of millions of computing cores. The numerical models of the physics, chemistry, and biology affecting the climate system need to be improved to run efficiently on these massive supercomputers.

This project gathers top minds in climate research and computer science to focus on the challenges in resilience, performance, and scalability when running these simulations at extreme scale. The project team includes personnel from the University of Illinois at Urbana-Champaign, University of Tennessee (ICL), University of Victoria, German Research School for Simulation Sciences, INRIA, Barcelona Supercomputing Center, Tokyo Institute of Technology, and the University of Tsukuba. ICL's main role in the ECS project is moving widely used climate model code toward Exascale, starting with node level performance and scalability, and then application resilience.

GridPAC

Grid with Power-Aware Computing

http://icl.eecs.utk.edu/gridpac/

Grid with Power-Aware Computing (GridPAC) is a middleware environment that will schedule multiple workflows across a distributed grid for system-wide optimization. GridPAC research focuses on supporting workflow execution using novel scheduling techniques on dynamic and heterogeneous resources. This project is a collaborative effort between the University of Texas Arlington, University of Florida, Virginia Tech, and the University of Tennessee at Knoxville (ICL).

At ICL, research has centered on the dynamic scheduling of DAG-structured workflows on a wide variety of computing resources, including multicore systems, systems utilizing multiple GPU resources, and grid-like distributed-memory, multicore systems using GPUs. The workflow applications are drawn from linear algebra, and thus they are likely to have a significant impact on multiple areas of scientific computing. GridPAC research is reflected in the PLASMA, QUARK, MAGMA, and DAGuE projects.

DOMAIN COLLABORATION

PetaApps

A Petaflop Cyberinfrastructure for Computing Free Energy Landscapes of Macro-and Bio-molecular Systems

The objective of this project is to develop an exascale cyberinfrastructure for the efficient calculation of free energy landscapes for complex macro-and bio-molecular systems. We then plan to demonstrate its effectiveness by applying it to two outstanding science problems: conformations of a linker protein in solution and self-assembly of lipids. In the context of this project, we have begun investigating, implementing, and evaluating different approaches to improve the scalability and the resilience of the cyberinfrastructure software computing for free energy landscapes. Our efforts are aimed at two axes of research. On one side, we are analyzing the current scalability of the existing framework and potential ways to drastically improve it. On the other side, we are investigating the software needs in order to achieve a reasonable level of resilience in this cyberinfrastructure.

PetaApps

Multiscale Software for Quantum Simulations in Nano Science and Technology

This project, led by North Carolina State University, is concerned with ab initio methods for computing the properties of materials and molecular structures. Petascale hardware enables ab initio calculations at unprecedented scale. The real-space multigrid (RMG) method being developed by this project uses a real-space mesh to represent the wavefunctions, the charge density, and the ionic pseudopotentials. The real-space formulation is advantageous for parallelization, since each processor can be assigned a region of space; and for convergence acceleration, since multiple length scales can be dealt with separately. Recent work focuses on overcoming memory and communication bottlenecks to allow RMG to scale to tens of thousands of cores, enabling simulations for thousands of atoms.

PEOPLE

ICL GROUP FALL 2011

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 more than a dozen bright and motivated graduate and undergraduate students. During the fall of 2011, we recruited 8 graduate students from institutions all over the world, including the University of Tennessee's own Electrical Engineering and Computer Science department. We have been, and will continue to be, very proactive in securing internships and assistantships for students who are hardworking and willing to learn.

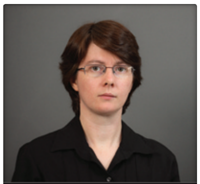
Dulceneia Becker
SENIOR RESEARCH ASSOCIATE


Wesley Bland
GRADUATE RESEARCH ASSISTANT


George Bosilca
RESEARCH ASSISTANT PROFESSOR

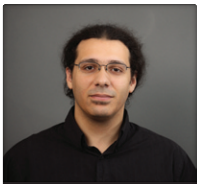
Aurelien Bouteiller
RESEARCH SCIENTIST II

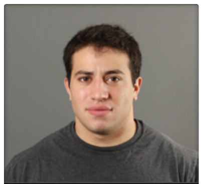
Anthony Canino
GRADUATE RESEARCH ASSISTANT

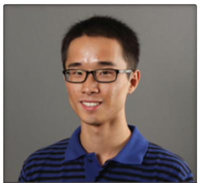
Chongxiao Cao
GRADUATE RESEARCH ASSISTANT


Charles Collins
IT ADMINSTRATOR I


Sam Crawford
INFORMATION SPECIALIST I


Yuanshun Dai
ASSISTANT PROFESSOR


Anthony Danalis
SENIOR RESEARCH ASSOCIATE

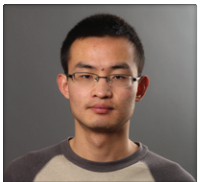
Tingxing Dong
GRADUATE RESEARCH ASSISTANT


Jack Dongarra
UNIV. DISTINGUISHED PROFESSOR

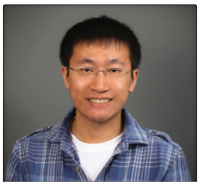
Peng Du
GRADUATE RESEARCH ASSISTANT


Mathieu Faverge
POST DOC. RESEARCH ASSOCIATE

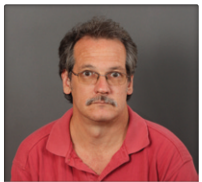
Don Fike
IT SPECIALIST III


Teresa Finchum
ADMIN. SERVICES ASSISTANT


Mark Gates
POST DOC. RESEARCH ASSOCIATE


Peter Gaultney
GRADUATE RESEARCH ASSISTANT


Azzam Haidar
RESEARCH SCIENTIST I


Blake Haugen
GRADUATE RESEARCH ASSISTANT


Thomas Herault
VISITING SCIENTIST


Mitch Horton
POST DOC. RESEARCH ASSOCIATE


Heike Jagode
SENIOR RESEARCH ASSOCIATE


Yulu Jia
GRADUATE RESEARCH ASSISTANT


Matt Johnson
GRADUATE RESEARCH ASSISTANT

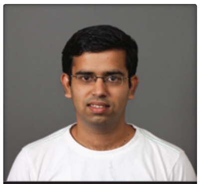

Vijay Joshi
GRADUATE RESEARCH ASSISTANT

Khairul Kabir
GRADUATE RESEARCH ASSISTANT

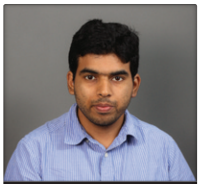

Kiran Kasichayanula
GRADUATE RESEARCH ASSISTANT

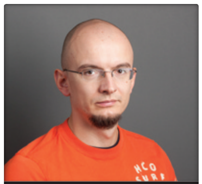

Jakub Kurzak
RESEARCH DIRECTOR

Julie Langou
RESEARCH LEADER

Tracy Lee
ACCOUNTING SPECIALIST III

Piotr Luszczek
RESEARCH DIRECTOR

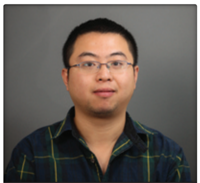

Teng Ma
GRADUATE RESEARCH ASSISTANT

Terry Moore
ASSOCIATE DIRECTOR

Shirley Moore
RESEARCH ASSOCIATE PROFESSOR

Phil Mucci
RESEARCH CONSULTANT

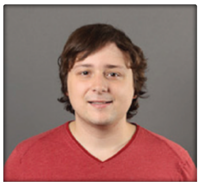

John Nelson
GRADUATE RESEARCH ASSISTANT

Paul Peltz
IT ADMINISTRATOR II

Tracy Rafferty
COORDINATOR II

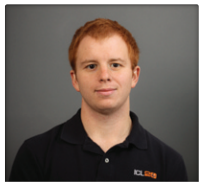

James Ralph
RESEARCH ASSISTANT

David Rogers
IT SPECIALIST III

Leighanne Sisk
ADMIN. SUPPORT ASSISTANT

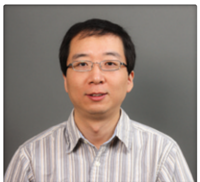

Fengguang Song
POST DOC. RESEARCH ASSOCIATE

Dan Terpstra
RESEARCH LEADER

Stanimire Tomov
RESEARCH DIRECTOR

Vince Weaver
POST DOC. RESEARCH ASSOCIATE

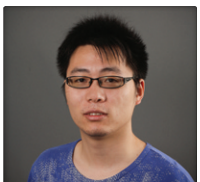

Wei Wu
GRADUATE RESEARCH ASSISTANT

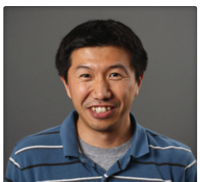

Ichitaro Yamazaki
SENIOR RESEARCH ASSOCIATE

Asim YarKhan
SENIOR RESEARCH ASSOCIATE

VISITORS

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. A byproduct of these relationships is the enormous opportunities to host and work with top minds within the global HPC community. 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.

It is also not uncommon to have students, both undergraduate and graduate, from various universities study with us for months on end, learning about our approaches and solutions to computing problems. We believe this mutual sharing of experience has been extremely beneficial, and we will continue providing opportunities for visits from our national and international colleagues in research.

Hartwig Anzt
KARLSRUHE INSTITUTE OF TECHNOLOGY, GERMANY

Marc Baboulin
UNIVERSITY OF PARIS-SUD, FRANCE

Micah Beck
EECS DEPARTMENT
UNIVERSITY OF TENNESSEE

Jim Browne
TEXAS ADVANCED COMPUTING CENTER
UNIVERSITY OF TEXAS

Yiannis Cotronis
UNIVERSITY OF ATHENS, GREECE

Skevos Evripidou
UNIVERSITY OF CYPRUS, GREECE

Franz Franchetti
DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING, CARNEGIE MELLON UNIVERSITY

Alexander Gaenko
IOWA STATE

Adriana Garties
OLIN COLLEGE OF ENGINEERING

Victor Gergel
NIZHNI NOVGOROD, RUSSIA

Ken Habgood
EECS DEPARTMENT
UNIVERSITY OF TENNESSEE

Julien Herrmann
ENS-LYON, FRANCE

Julien Langou
UNIVERSITY OF COLORADO AT DENVER

Joshua New
ORNL

Lynne Parker
CISML
UNIVERSITY OF TENNESSEE

Pierre Ramet
INRIA BORDEAUX, FRANCE

Ashay Rane
TEXAS ADVANCED COMPUTING CENTER
UNIVERSITY OF TEXAS

Ala Rezmerita
UNIVERSITY OF PARIS-SUD / INRIA SACLAY
FRANCE

Yves Robert
ENS-LYON, FRANCE

Abhinav Sarje
LAWRENCE BERKELEY NATIONAL LAB

Raffaele Solca
ETH ZURICH, SWITZERLAND

Lukas Steiblys
UNIVERSITY OF MANCHESTER, UK

Vendel Szeremi
UNIVERSITY OF MANCHESTER, UK

Vladimir Voevodin
MOSCOW, RUSSIA

ALUMNI

Since its inception, ICL has attracted many post-doctoral researchers and professors 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. See our timeline on page 5 for a list of some of the prominent experts who have passed through ICL on their way to distinguished careers at other organizations and academic institutions.

Carolyn Aebischer
1990-1993

Sudesh Agrawal
2001-2006

Bivek Agrawal
2004-2006

Emmanuel Agullo
2009

Jennifer Allgeyer
1993

Wes Alvaro
2011

Ed Anderson
1989-1991

Daniel Andrzejewski
2007

Thara Angskun
2003-2007

Papa Arkhurst
2003

Dorian Arnold
1999-2001

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

Adam Beguelin
1991

Annamaria Benzoni
1991

Tom Berry
1991

Vincent Berthoux
2010

Scott Betts
1997-1998

Nikhil Bhatia
2003-2005

Laura Black
1996

Noel Black
2002-2003

Susan Blackford
1989-2001

Kartheek Bodanki
2009

David Bolt
1991

Fernando Bond
1999-2000

Carolyn Bowers
1992

Barry Britt
2007-2009

Randy Brown
1997-1999

Bonnie Browne
2011

Cynthia Browne
2005

Murray Browne
1998-1999

Antonin Bukovsky
1998-2003

Greg Bunch
1995

Alfredo Buttari
2008

Domingo Gimenez Canovas
2001

Henri Casanova
1995-1998

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

Matthias Colin
2004

Tom Cortese
2009

Camille Coti
2007

Jason Cox
1993-1997

David Cronk
1999-2010

Javier Cuenca
2003

Manoel Cunha
2006

Cricket Deane
1998-1999

Remi Delmas
2006

Frederic Desprez
1994-1995

Ying Ding
2000-0201

Jun Ding
2001-2003

Jin Ding
2003

Martin Do
1993-1994

Leon Dong
2000-2001

Nick Dongarra
2000

David Doolin
1997

Andrew Downey
1998-2003

Mary Drake
1989-1992

Julio Driggs
2002-2004

Brian Drum
2001-2004

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

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
2007

Karl Fuerlinger
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
2009

Christoph Geile
2008

Jean Patrick Gelas
2001

Boris Gelfend
1993

Jonathan Gettler
1996

Bruno Giuseppe
2001

Eric Greaser
1993

Stan Green
1992-1996

Alice Gregory
2004-2006

Jason Gurley
1997-1998

Bilel Hadri
2009

Hunter Hagewood
2000-2001

Christian Halloy
1996-1997

Sven Hammarling
1996-1997

J. Mike Hammond
1994-1995

Hidehiko Hasegawa
1995-1996

Satomi Hasegawa
1995-1996

Chris Hastings
1996

David Henderson
1999-2001

Greg Henry
1996

Julien Hermann
2011

Holly Hicks
1993-1994

Alexandra Hicks-Hardiman
2009

Sid Hill
1996-1998

Tomoyuki Hiroyasu
2002-2003

George Ho
1998-2000

Josh Hoffman
2010

Jeff Horner
1995-1999

Yan Huang
2000-2001

Aurelie Hurault
2009

Chris Hurt
2002

Paul Jacobs
1992-1995

Emmanuel Jeannot
2006

Weizhong Ji
1999-2000

Weicheng Jiang
1992-1995

Song Jin
1997-1998

Patrick Johansson
2001

Aral Johnson
2009

Sean Jolly
1997-1998

Kim Jones
1996-1997

Jan Jones
1992-2008

Venkata Kakani
2007

Ajay Kalhan
1995

Balajee Kannan
2001

Madhuri Kasam
2007-2008

Ajay Katta
2011

David Katz
2002

Joshua Kelly
2000-2001

Supriya Kilambi
2008

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

Julien Langou
2008

Jeff Larkin
2003-2005

Brian LaRose
1990-1992

Frank Lauer
2010-2011

DongWoo Lee
2000-2002

Pierre Lemarinier
2011

Todd Letsche
1993-1994

Sharon Lewis
1992-1995

Weiran Li
2002

Xiang Li
2001

Yinan Li
2009

Chaoyang Liu
2000

Kevin London
1996-2005

Matt Longley
1999

Hatem Ltaief
2011

Daniel Lucio
2008

Richard Luczak
2000-2001

Robert Manchek
1990-1996

Tushti Marwah
2004

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

Erik Moore
2000

Keith Moore
1987-2007

Robert Morgan
1990-1991

Kishan Motheramgari
1997

Steven Moulton
1991-1993

Matthew Nabity
2008

Shankar Narasimhaswami
2004-2005

Rajib Nath
2008-2010

Fernando Navarro
2009

Donnie Newell
2010

Peter Newton
1994-1995

Jonas Nilsson
2001

Jakob Oestergaard
2000

Caroline Papadopoulos
1997-1998

Leelinda Parker
2002

Dilip Patlolla
2008

Andy Pearson
1989-1991

Theresa Pepin
1994

Antoine Petitet
1993-2001

Vlado Pjesivac
2008

Jelena Pjesivac-Grbovic
2003-2007

James S. Plank
1991-1992

Tim Poore
2009

Roldan Pozo
1992-1994

Farzona Pulatova
2005-2006

Martin Quinson
2001

Tammy Race
1999-2001

Ganapathy Raman
1998-2000

Kamesh Ramani
2003

Mei Ran
1999-2004

Arun Rattan
1997

Sheri Reagan
1995-1996

Mike Reynolds
1994

Jon Richardson
1990-1991

Yves Robert
1996-1997

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

Peter Soendergaard
2000

Raffaele Solca
2011

Gwang Son
2007-2009

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

Yuan Tang
2005-2006

Yusuke Tanimura
2003

Keita Teranishi
1998

Joe Thomas
2009

John Thurman
1998-1999

Francoise Tisseur
1997

Jude Toth
1993-1994

Bernard Tourancheau
1993-1994

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

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

Jin Yi
2009-2010

Haihang You
2004-2009

Lamia Youseff
2007

Brian Zachary
2009-2010

Yuanlei Zhang
2001-2005

Junlong Zhao
2002

Yong Zheng
2001

Luke Zhou
2000-2001

Min Zhou
2002-2004

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, industry research leaders, and academic institutions, both here and abroad.

These businesses and institutions have helped us build a solid foundation of meaningful and lasting relationships that have contributed significantly to our efforts to be a world leader in computational science research. We also routinely develop relationships with researchers whose primary focus is on other scientific disciplines, such as biology, chemistry, and physics, which makes many of our collaborations truly multidisciplinary. Together with these partners, we have built a strong portfolio of shared resources, both material and intellectual.

Many application and tool vendors, including Intel, MathWorks, SGI, and Cray, have utilized our work. In addition, Hewlett Packard, IBM, Intel, NVIDIA, SGI, and Sun have all utilized our linear algebra work. The dense linear algebra portions of their libraries have been based on the BLAS, LAPACK, and ScaLAPACK specifications and software developed by ICL. On the following page, 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. As our list of government and academic partners continues to grow, we also continue to search for opportunities to establish partnerships with HPC vendors.

- AMD
- Argonne National Laboratory
- Barcelona Supercomputing Center BARCELONA, SPAIN
- Central Institute for Applied Mathematics JÜLICH, GERMANY
- European Centre for Research and Advanced Training in Scientific Computing TOULOUSE, FRANCE
- Cray
- Defense Advanced Research Projects Agency
- Danish Computing Center for Research and Education LYNGBY, DENMARK
- Doshisha University KYOTO, JAPAN
- École Polytechnique Federale de Lausanne LAUSANNE, SWITZERLAND
- European Exascale Software Initiative EUROPEAN UNION
- Forschungszentrum Jülich JÜLICH, GERMANY
- Georgia Tech University
- High Performance Computing Center Stuttgart STUTTGART, GERMANY
- HP
- IBM
- Il Consiglio Nazionale delle Ricerche ROME, ITALY
- Indiana University
- Institut ETH Zentrum ZURICH, SWITZERLAND
- Intel
- International Exascale Software Project
- Information Sciences Institute
- The Joint Institute for Computational Science
- Kasetsart University BANGKOK, THAILAND
- Los Alamos National Laboratory
- Lawrence Livermore National Laboratory
- Microsoft Research
- MetaCenter Regional Alliance
- National Aeronautics and Space Administration
- National Center for Supercomputing Applications
- National Institute for Computational Sciences
- National Institute of Standards and Technology
- Normal Superieure de Lyon LYON, FRANCE
- The National Science Foundation
- Numerical Algorithms Group Ltd OXFORD, ENGLAND
- NVIDIA
- Oak Ridge National Laboratory
- Parallel and HPC Application Software Exchange TSUKUBA, JAPAN
- Rice University
- Rutherford Appleton Laboratory OXFORD, ENGLAND
- San Diego Supercomputing Center
- SGI
- Soongsil University SEOUL, SOUTH KOREA
- Sun Microsystems
- Technische Universitaet Wien VIENNA, AUSTRIA
- The MathWorks
- Tokyo Institute of Technology TOKYO, JAPAN
- United States Department of Defense
- United States Department of Energy
- Universität Mannheim MANNHEIM, GERMANY
- Université Claude Bernard de Lyon LYON, FRANCE
- University of California, Berkeley
- University of California, San Diego
- University of Manchester MANCHESTER, ENGLAND
- University of Umeå UMEÅ, SWEDEN
- University of Tennessee EECS
- Virtual Institute - High-Productivity Supercomputing

HARDWARE RESOURCES

AS THE new GPU hybrid computing paradigm leads 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 as 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.

ICL has multiple heterogeneous systems in house, and access to numerous architectures around the country, due in large part to our many partners and collaborators. Locally, we maintain systems ranging from individual desktops to large, networked clusters.

HYBRID SYSTEMS

8x AMD Opteron Processor 8439 SE
(48 cores)
128G RAM
GeForce GTX 480
NVIDIA Tesla S1070

4x AMD Opteron Processor 6172
(48 cores)
128G RAM
NVIDIA Tesla S2050

4x AMD Opteron Processor 6180 SE
(48 cores)
256G RAM
NVIDIA Tesla S2050

4x AMD Opteron Processor 6180 SE
(48 cores)
256G RAM
NVIDIA Tesla S2050

AMD Phenom II X6 1100T Processor
8G RAM
AMD Radeon HD 5870
AMD Radeon HD 6970

CLUSTER SYSTEMS

16 Node (2x Intel E5520) (128 cores) cluster with NVIDIA Fermi/AMD FirePro V7800 and OCZ 50G SSD connected with Mellanox Infiniband 20G

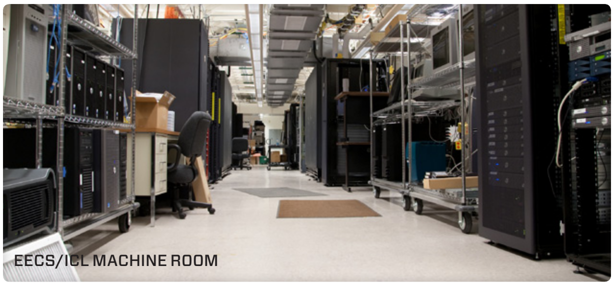
EECS/ICL MACHINE ROOM

CRAY XT5 KRAKEN AT OAK RIDGE NATIONAL LABORATORY


EECS/ICL MACHINE ROOM

CRAY XT5 JAGUAR AT OAK RIDGE NATIONAL LABORATORY

ICL also has access to many remote resources to help keep us at the forefront of enabling technology research, including some machines that are regularly found on the TOP500 list of the world's fastest supercomputers. 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. The NCCS houses Jaguar, a Cray XT5 that was the third fastest supercomputer in the world in mid-2011. The National Institute for Computational Sciences (NICS), another computing facility at ORNL, houses Kraken, UT's Cray XT5 system which is one of the world's fastest open-science supercomputers. We also have access to resources on XSEDE—the successor to TeraGrid—and France's Grid5000.

The following are some of the remote systems and architectures that we utilize:

- Cray X2, XT4, XT5, HP XC System
- IBM Power 6 & 7, BlueGene/P and 2nd Generation Cell
- Many large (512+ proc) Linux Clusters

PUBLICATIONS

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

Agullo, E., Augonnet, C., Dongarra, J., Ltaief, H., Namyst, R., Thibault, S., Tomov, S. "**A Hybridization Methodology for High-Performance Linear Algebra Software for GPUs**," in *GPU Computing Gems, Jade Edition*, Hwu, W. eds. Elsevier, 2, 473-484, 2011.

Agullo, E., Augonnet, C., Dongarra, J., Faverge, M., Langou, J., Ltaief, H., Tomov, S. "**LU Factorization for Accelerator-based Systems**," *ICL Technical Report*, Submitted to AICCSA 2011, ICL-UT-10-05, December 27, 2010.

Agullo, E., Coti, C., Herault, T., Langou, J., Peyronnet, S., Rezmerita, A., Cappello, F., Dongarra, J. "**QCG-OMPI: MPI Applications on Grids**," *Future Generation Computer Systems*, Vol. 27, No. 4, pp. 357-369, April 2011.

Agullo, E., Giraud, L., Guermouche, A., Haidar, A., Lanteri, S., Roman, J. "**Algebraic Schwarz Preconditioning for the Schur Complement: Application to the Time-Harmonic Maxwell Equations Discretized by a Discontinuous Galerkin Method**," *20th International Conference on Domain Decomposition Methods*, UC San Diego, in La Jolla, California, February 7-11, 2011.

Agullo, E., Giraud, L., Guermouche, A., Haidar, A., Roman, J. "**Parallel algebraic domain decomposition solver for the solution of augmented systems**," *Parallel, Distributed, Grid and Cloud Computing for Engineering*, Ajaccio, Corsica, France, 12-15 April, 2011.

Baboulin, M., Becker, D., Dongarra, J. "**A parallel tiled solver for dense symmetric indefinite systems on multicore architectures**," *University of Tennessee Computer Science Technical Report*, ICL-UT-11-07, October 12, 2011.

Baboulin, M., Dongarra, J., Herrmann, J., Tomov, S. "**Accelerating Linear System Solutions Using Randomization Techniques**," INRIA RR-7616 / LAWN 246 (presented at *International AMMCS'11*), Waterloo, Ontario, Canada, July 25-29, 2011.

Becker, D., Baboulin, M., Dongarra, J. "**Reducing the Amount of Pivoting in Symmetric Indefinite Systems**," *University of Tennessee Innovative Computing Laboratory Technical Report*, Submitted to *PPAM 2011*, Knoxville, TN, ICL-UT-11-06, May 14, 2011.

Becker, D., Faverge, M., Dongarra, J. "**Towards a Parallel Tile LDL Factorization for Multicore Architectures**," *ICL Technical Report*, Submitted to *SC11*, Seattle, WA, ICL-UT-11-03, April 15, 2011.

Bosilca, G., Bouteiller, A., Herault, T., Lemarier, P., Saengpatsa, N., Tomov, S., Dongarra, J. "**Performance Portability of a GPU Enabled Factorization with the DAGuE Framework**," *IEEE Cluster: workshop on Parallel Programming on Accelerator Clusters* (PPAC), June 24, 2011.

Bosilca, G., Bouteiller, A., Herault, T., Lemarinier, P., Saengpatsa, N., Tomov, S., Dongarra, J. "**A Unified HPC Environment for Hybrid Manycore/GPU Distributed Systems**," *IEEE International Parallel and Distributed Processing Symposium* (submitted), Anchorage, AK, May 16-20, 2011.

Bosilca, G., Herault, T., Rezmerita, A., Dongarra, J. "**On Scalability for MPI Runtime Systems**," *University of Tennessee Computer Science Technical Report*, Knoxville, TN, ICL-UT-11-05, May 1, 2011.

Chaarawi, M., Gabriel, E., Keller, R., Graham, R., Bosilca, G., Dongarra, J. "**OMPIO: A Modular Software Architecture for MPI I/O**," *18th EuroMPI*, Cotronis, Y., Danalis, A., Nikolopoulos, D., Dongarra, J. eds. Springer, Santorini, Greece, pp. 81-89, September, 2011.

Danalis, A., Luszczek, P., Marin, G., Vetter, J., Dongarra, J. "**BlackjackBench: Hardware Characterization with Portable Micro-Benchmarks and Automatic Statistical Analysis of Results**," *IEEE International Parallel and Distributed Processing Symposium* (submitted), Anchorage, AK, May 16-20, 2011.

Dongarra, J. "**Performance of Various Computers Using Standard Linear Equations Software (Linpack Benchmark Report)**," *University of Tennessee Computer Science Technical Report*, UT-CS-89-85, 2011.

Dongarra, J., Beckman, P., et al. "**The International Exascale Software Roadmap**," *International Journal of High Performance Computing*, Vol. 25, No. 1, pp. 3-60, 2011, ISSN 1094-3420.

Dongarra, J., Faverge, M., Herault, T., Langou, J., Robert, Y. "**Hierarchical QR factorization algorithms for multi-core cluster systems**," University of Tennessee Computer Science Technical Report (also LAWN 257), UT-CS-11-684, Oct 4, 2011.

Dongarra, J., Faverge, M., Ishikawa, Y., Namyst, R., Rue, F., Trahay, F. "**EZTrace: a generic framework for performance analysis**," *ICL Technical Report*, Submitted to CCGrid 2011, ICL-UT-11-01, December 7, 2010.

Dongarra, J., Faverge, M., Ltaief, H., Luszczek, P. "**Exploiting Fine-Grain Parallelism in Recursive LU Factorization**," *University of Tennessee Computer Science Technical Report* submitted to *PARCO'11*, Gent, Belgium, ICL-UT-11-04, April, 2011.

Du, P., Bouteiller, A., Bosilca, G., Herault, T., Dongarra, J. "**Algorithm-based Fault Tolerance for Dense Matrix Factorizations**," *University of Tennessee Computer Science Technical Report*, Knoxville, TN, UT-CS-11-676, August 05, 2011.

Du, P., Luszczek, P., Dongarra, J. "**High Performance Linear System Solver with Resilience to Multiple Soft Errors**," *University of Tennessee Computer Science Technical Report* (also LAWN 256), UT-CS-11-683, Oct 4, 2011.

Du, P., Luszczek, P., Dongarra, J. "**High Performance Dense Linear System Solver with Soft Error Resilience**," *IEEE Cluster 2011*, Austin, TX, September 26-30, 2011.

Du, P., Luszczek, P., Tomov S., Dongarra, J. "**Soft Error Resilient QR Factorization for Hybrid System**," *University of Tennessee Computer Science Technical Report*, Knoxville, TN, UT-CS-11-675, July 1, 2011.

Fengguang, S., Tomov, S., Dongarra, J. "**Efficient Support for Matrix Computations on Heterogeneous Multi-core and Multi-GPU Architectures**," *University of Tennessee Computer Science Technical Report*, UT-CS-11-668, (also LAWN 250), June 16, 2011.

Haidar, A., Giraud, L., Ben Hadj Ali, H., Sourbier, F., Operto, S., Virieux, J. "**3-D parallel frequency-domain visco-acoustic wave modelling based on a hybrid direct/iterative solver**," *73rd EAGE Conference & Exhibition incorporating SPE EUROPEC 2011*, Vienna, Austria, 23-26 May, 2011.

Haidar, A., Ltaief, H., Dongarra, J. "**Toward High Performance Divide and Conquer Eigensolver for Dense Symmetric Matrices**," Submitted to *SIAM Journal on Scientific Computing* (SISC), 2011.

Haidar, A., Ltaief, H., Dongarra, J. "**Parallel Reduction to Condensed Forms for Symmetric Eigenvalue Problems using Aggregated Fine-Grained and Memory-Aware Kernels**," *University of Tennessee Computer Science Technical Report*, UT-CS-11-677, (also LAWN 254), August 5, 2011.

Haidar, A., Ltaief, H., YarKhan, A., Dongarra, J. "**Analysis of Dynamically Scheduled Tile Algorithms for Dense Linear Algebra on Multicore Architectures**," *University of Tennessee Computer Science Technical Report*, UT-CS-11-666, (also LAWN 243), March 10, 2011.

Horton, M., Tomov, S., Dongarra, J. "**A Class of Hybrid LAPACK Algorithms for Multicore and GPU Architectures**," *Symposium for Application Accelerators in High Performance Computing* (SAAHPC'11), Knoxville, TN, July 19-20, 2011.

Kurzak, J., Tomov, S., Dongarra, J. "**Autotuning GEMMs for Fermi**," University of Tennessee Computer Science Technical Report, UT-CS-11-671, (also Lawn 245), April 18, 2011.

Lively, C., Wu, X., Taylor, V., Moore, S., Chang, H.-C., Cameron, K. "**Energy and performance characteristics of different parallel implementations of scientific applications on multicore systems**," International Journal of High Performance Computing Applications, Vol. 25, No. 3, pp. 342-350, 2011.

Lively, C., Wu, X., Taylor, V., Moore, S., Chang, H.-C., Su, C.-H., Cameron, K. "**Power-Aware Prediction Models of Hybrid (MPI/OpenMP) Scientific Applications**," International Conference on Energy-Aware High Performance Computing (EnA-HPC 2011), Hamburg, Germany, September 7-9, 2011.

Ltaief, H., Luszczek, P., Dongarra, J. "**Profiling High Performance Dense Linear Algebra Algorithms on Multicore Architectures for Power and Energy Efficiency**," *International Conference on Energy-Aware High Performance Computing* (EnA-HPC 2011), Hamburg, Germany, September 7-9, 2011.

Ltaief, H., Luszczek, P., Dongarra, J. "**High Performance Bidiagonal Reduction using Tile Algorithms on Homogeneous Multicore Architectures**," *University of Tennessee Computer Science Technical Report*, UT-CS-11-673, (also LAWN 247), May 18, 2011.

Luszczek, P., Kurzak, J., Dongarra, J. "**Changes in Dense Linear Algebra Kernels - Decades Long Perspective**," in *Solving the Schrodinger Equation: Has everything been tried?* (to appear), Popular, P. eds. Imperial College Press, 2011.

Luszczek, P., Ltaief, H., Dongarra, J. "**Two-stage Tridiagonal Reduction for Dense Symmetric Matrices using Tile Algorithms on Multicore Architectures**," *IEEE International Parallel and Distributed Processing Symposium* (submitted), Anchorage, AK, May 16-20, 2011.

Luszczek, P., Meek, E., Moore, S., Terpstra, D., Weaver, V., Dongarra, J. "**Evaluation of the HPC Challenge Benchmarks in Virtualized Environments**," *6th Workshop on Virtualization in High-Performance Cloud Computing*, Bordeaux, France, August 30, 2011.

Ma, T., Bosilca, G., Bouteiller, A., Goglin, B., Squyres, J., Dongarra, J. "**Kernel Assisted Collective Intra-node MPI Communication Among Multi-core and Many-core CPUs**," *Int'l Conference on Parallel Processing* (ICPP '11), Taipei, Taiwan, September, 2011.

Ma, T., Bouteiller, A., Bosilca, G., Dongarra, J. "**Impact of Kernel-Assisted MPI Communication over Scientific Applications: CPMD and FFTW**," *18th EuroMPI*, Cotronis, Y., Danalis, A., Nikolopoulos, D., Dongarra, J. eds. Springer, Santorini, Greece, pp. 247-254, September, 2011.

Ma, T., Herault, T., Bosilca, G., Dongarra, J. "**Process Distance-aware Adaptive MPI Collective Communications**," *IEEE Int'l Conference on Cluster Computing* (Cluster 2011), Austin, Texas, September, 2011.

Malony, A., Biersdorff, S., Shende, S., Jagode, H., Tomov, S., Juckeland, G., Dietrich, R., Duncan Poole, P., Lamb, C. "**Parallel Performance Measurement of Heterogeneous Parallel Systems with GPUs**," *International Conference on Parallel Processing* (ICPP 2011), Taipei, Taiwan, September 13-16, 2011.

Moore, S., Ralph, J. "**User-defined Events for Hardware Performance Monitoring**," *ICCS 2011 Workshop: Tools for Program Development and Analysis in Computational Science*, www.sciencedirect.com, Singapore, June 1, 2011.

Nath, R., Tomov, S., Dong, T., Dongarra, J. "**Optimizing Symmetric Dense Matrix-Vector Multiplication on GPUs**," *ACM/IEEE Conference on Supercomputing* (SC11), Seattle, WA, November 12-18, 2011.

Sourbier, F., Haidar, A., Giraud, L., Ben Hadj Ali, H., Operto, S., Virieux, J. "**Three-dimensional parallel frequency-domain visco-acoustic wave modelling based on a hybrid direct/iterative solver**," To appear in *Geophysical Prospecting Journal*, 2011.

Weaver, V., Dongarra, J. "**Can Hardware Performance Counters Produce Expected, Deterministic Results?**," *3rd Workshop on Functionality of Hardware Performance Monitoring*, Atlanta, GA, December 4, 2010.

White, T., Dongarra, J. "**Overlapping Computation and Communication for Advection on a Hybrid Parallel Computer**," *IEEE International Parallel and Distributed Processing Symposium* (submitted), Anchorage, AK, May 16-20, 2011.

YarKhan, A., Kurzak, J., Dongarra, J. "**QUARK Users' Guide: QUeueing And Runtime for Kernels**," *University of Tennessee Innovative Computing Laboratory Technical Report*, ICL-UT-11-02, 2011.

You, H., Liu, Q., Li, Z., Moore, S. "**The Design of an Auto-tuning I/O Framework on Cray XT5 System**," *Cray Users Group Conference* (CUG'11) (Best Paper Finalist), Fairbanks, Alaska, May 23-26, 2011.

You, H., Rekapalli, B., Liu, Q., Moore, S. "**Autotuned Parallel I/O for Highly Scalable Biosequence Analysis**," *TeraGrid'11*, Salt Lake City, Utah, July 18-21, 2011.

CONFERENCES

EVERY YEAR, our research staff regularly attends 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. The following pages contain a list of events we have participated in over the past year.

2011 CONFERENCES

JANUARY 6	PetaApps Meeting	ATLANTA, GEORGIA
JANUARY 17-18	Future Grid NSF Review	BLOOMINGTON, IN
JANUARY 26-28	ESC Planning Meeting	CHICAGO, IL
FEBRUARY 7-9	MPI Forum	SAN JOSE, CA
FEBRUARY 28 - MARCH 4	SIAM Conference On Computational Science and Engineering (CSE11)	RENO, NEVADA
MARCH 1-2	PERI Biannual Review	HOUSTON, TX
MARCH 28-30	Virtual Institute-High Productivity Supercomputing 7th Tuning Workshop	STUTTGART, GERMANY
MARCH 29-31	HPCC Conference	NEWPORT, RI
APRIL 1-2	2011 CRA-W Grad Cohort Workshop	BOSTON, MA
APRIL 5-7	International Exascale Software Project	SAN FRANCISCO, CA
APRIL 14-15	Keeneland Workshop	ATLANTA, GEORGIA
MAY 9-11	MPI Forum	SAN JOSE, CA
MAY 11-13	Workshop on GPU-enabled Numerical Libaries	BASEL, SWITZERLAND
MAY 16-20	IEEE International Parallel & Distributed Processing Symposium	ANCHORAGE, AK
MAY 23-26	IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing	NEWPORT BEACH, CA
MAY 29 - JUNE 1	4th Scheduling Workshop	AUSSOIS, FRANCE
JUNE 1-3	International Conference on Computational Science	SINGAPORE, SINGAPORE

JUNE 3-4	**CodeStock**	KNOXVILLE, TN
JUNE 13-15	**An Event Apart**	ATLANTA, GA
JUNE 19-23	**ISC'11**	HAMBURG, GERMANY
JUNE 27-29	**International Research Workshop Advanced High Performance Computing Systems**	CETRARO, ITALY
JULY 10-14	**Scientific Discovery through Advanced Computing (SciDAC)**	DENVER, CO
JULY 10-15	**Seminar ETH Zurich**	ZURICH, SWITZERLAND
JULY 18-20	**MPI Forum**	CHICAGO, IL
JULY 18-21	**TeraGrid 2011**	SALT LAKE CITY, UT
JULY 19-21	**2011 Symposium on Application Accelerators in High Performance Computing**	KNOXVILLE, TN
AUGUST 1-4	**CScADS Performance Tools Workshop**	TAHOE CITY, CA
AUGUST 7-10	**Keeneland NSF Review**	ATLANTA, GA
AUGUST 8-9	**MuMMI Annual Review**	ARGONNE, IL
AUGUST 16	**Titan Summit**	OAK RIDGE, TN
AUGUST 16-19	**DOE Advanced Computational Software Collection Workshop**	BERKELEY, CA
AUGUST 17	**PetaApps Meeting**	ATLANTA, GEORGIA
AUGUST 29 - SEPTEMBER 2	**EuroPar 2011**	BORDEAUX, FRANCE
AUGUST 30-31	**Intel MIC Workshop**	PORTLAND, OR
AUGUST 30 - SEPTEMBER 2	**International Conference on Parallel Computing**	GHENT, BELGIUM
SEPTEMBER 5-7	**PetaApps Meeting**	ATLANTA, GEORGIA
SEPTEMBER 7-9	**Ena-HPC 2011**	HAMBURG, GERMANY
SEPTEMBER 11-14	**Parallel Processing and Applied Mathematics**	TORUN, POLAND
SEPTEMBER 13-16	**Fall Creek Falls 2011**	GATLINBURG, TN
SEPTEMBER 18-21	**EuroMPI**	SANTORINI, GREECE
SEPTEMBER 20-21	**PERI Biannual Review**	EUGENE, OR
SEPTEMBER 25-30	**IEEE Cluster 2011**	AUSTIN, TX
OCTOBER 5-7	**International Exascale Software Project**	COLOGNE, GERMANY
OCTOBER 10	**PetaApps Meeting**	NASHVILLE, TN
OCTOBER 10-11	**European Exascale Software Initiative Final International Conference**	BARCELONA, SPAIN
OCTOBER 24-26	**MPI Forum**	CHICAGO, IL
OCTOBER 26-29	**International Symposium on Computer Architecture and High Performance Computing**	VITÓRIA, ESPÍRITO SANTO, BRAZIL
NOVEMBER 3	**Exascale Supercomputing**	SEOUL, KOREA
NOVEMBER 12-18	**SC11**	SEATTLE, WA

OTHER EFFORTS

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 exciting future of high performance computing. This section contains some of the activities in which we are participating or have a taken a leadership role.

CITR
CENTER FOR INFORMATION TECHNOLOGY RESEARCH

As one of the nine Centers of Excellence at the University of Tennessee, the Center for Information Technology Research (CITR) was established in the spring of 2001 to drive the growth and development of leading edge Information Technology Research (ITR) at the University. ITR is a broad, cross-disciplinary research area that investigates ways in which fundamental innovations in Information Technology affect and are affected by the research process.

The mission of CITR is twofold: first, to build up a thriving, well-funded community in basic and applied ITR at UT in order to help the university capitalize on the rich supply of research opportunities that now exist in this area; and second, to grow an interdisciplinary Computational Science program as part of the University curriculum that enables graduate students to augment their degree with computational knowledge and skills from disciplines outside their major.

FIND OUT MORE AT **http://citr.eecs.utk.edu/**

10^18 INTERNATIONAL EXASCALE SOFTWARE PROJECT

Several recent high profile studies from the computational science community make it clear that the radical new design properties of future extreme-scale platforms—massive concurrency, processor heterogeneity, constrained power budgets, complex memory architectures, unprecedented data I/O requirements—will require equally radical innovations in the software infrastructure that scientists and engineers will need to enable their extreme-scale research. But the challenge of creating an entirely new software stack for high performance computing is daunting. Because it was clear to many community leaders that, for various reasons, this challenge would demand an unprecedented level of coordination and cooperation within the worldwide open source software R&D community, the International Exascale Software Project (IESP) was created in 2009 to catalyze the collective effort necessary to meet it. The IESP's guiding purpose is to empower ultrahigh resolution and data intensive science and engineering research through the year 2020 by developing a plan for 1) a common, high quality computational environment for peta/exascale systems and for 2) catalyzing, coordinating, and sustaining the effort of the international open source software community to create that environment as quickly as possible.

During its first two years of work, the IESP has organized seven workshops at different locations around the globe: Santa Fe, NM (USA); Paris, France; Tsukuba, Japan; Oxford, UK; Maui, HI (USA); San Francisco, CA (USA); and Cologne, Germany. The agendas for each workshop were structured to provide progressively greater definition for the components of the IESP plan, with each successive meeting building on the results of the previous meeting. The goal of the first year's workshops was to conduct an application needs assessment and then develop a coordinated roadmap to guide open source HPC development with better coordination and fewer missing components. Version 1.1 of the IESP Roadmap was published electronically on October 18, 2010. The work of the IESP is also credited with helping to stimulate major new government initiatives in the US, the EU, and Japan focused on (and working together toward) a new HPC software infrastructure for extreme-scale science. More information about the IESP, including the latest version of the Roadmap, meeting notes, white papers, and presentations, can be found by visiting the project website.

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

IGMCS

INTERDISCIPLINARY GRADUATE MINOR IN COMPUTATIONAL SCIENCE

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 2011, there were 15 departments from four UT colleges contributing more than 100 courses to the program.

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

In late 2009, the Innovative Computing Laboratory was designated a CUDA Center of Excellence (CCOE) by NVIDIA Corporation, a world-wide leader in technologies for visual computing and inventor of the graphics processing unit (GPU). This award led to the establishment of a productive long-term collaboration between ICL and NVIDIA. As part of the collaboration and CCOE designation, ICL has continuously received hardware, financial support, and other resources from NVIDIA. Joining a very small and select group of CUDA CCOEs such as labs at Harvard University, the University of Utah, and the University of Illinois at Urbana-Champaign, UTK's CCOE focuses on the development of numerical linear algebra libraries for CUDA-based hybrid architectures. Our 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 general area of high-performance scientific computing.

VI-HPS

In mid-2007, ICL became part of a new collaboration for HPC research called the Virtual Institute – High Productivity Supercomputing (VI-HPS), whose mission is "to improve the quality and accelerate the development process of complex simulation programs in science and engineering that are being designed for the most advanced parallel computer systems." The new institute, comprised of institutions in Germany and the US, including ICL, unites some of the brightest minds in HPC research who are committed to helping engineers and domain scientists become more efficient and effective users of HPC applications.

ICL's membership and contributions have already proven invaluable to the success of the institute, and we look forward to working with the other partners in the development of leading-edge tools. According to Felix Wolf, spokesman and member of the VI-HPS Steering Board, "During the past year, the virtual institute made significant progress towards closer integration of our performance-analysis tool suite. With funding from the German Ministry of Education and Research and the US Department of Energy, Office of Science, the joint measurement system Score-P, which will serve as a common basis for the tools Periscope, Scalasca, TAU, and Vampir, is approaching its beta release at SC11. In this system, PAPI, the hardware counter library from ICL, will provide access to a wealth of performance-relevant hardware events on the most recent architectures. We are glad that we can benefit from the high-quality products the team at ICL delivers to the HPC community."

VI-HPS PARTNERS

Jülich Forschungszentrum
RWTH Rheinisch-Westfälische Technische Hochschule Aachen
TUM Technische Universität München
Universität Stuttgart
Université de Versailles Saint-Quentin-en-Yvelines
German Research School for Simulation Sciences
Technische Universität Dresden
University of Oregon
The University of Tennessee

FIND OUT MORE AT **http://www.vi-hps.org/**

SPONSORS

FOR MORE than 22 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:

INNOVATIVE
COMPUTING LABORATORY
THE UNIVERSITY of TENNESSEE
Department of Electrical Engineering and Computer Science

Suite 413 Claxton
1122 Volunteer Blvd
Knoxville TN 37996-3450

(865) 974-8295
icl@eecs.utk.edu
http://icl.eecs.utk.edu/

INNOVATIVE
COMPUTING LABORATORY