CELEBRATING 20 YEARS

INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY of TENNESSEE

2009/2010 REPORT

INNOVATIVE COMPUTING LABORATORY

INNOVATIVE COMPUTING LABORATORY 2009/2010 REPORT

EDITED BY Scott Wells 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.

ICL 2009/2010 REPORT

FROM THE DIRECTOR

This year marks the 20th anniversary of the Innovative Computing Laboratory (ICL).

Looking back over this 20-year period, the evolution and growth of the technology for computing has been truly astonishing. In an environment where technology changes every 18 months, ICL cannot afford to stand still. In the past 20 years we have seen Gigaflop computing with around 10 processors turn into Teraflop computing with 1,000 processors to Petaflop computing with greater than 100,000 processors. We see the horizon and it's clear in the next 10 years we will be at Exascale computing with greater than 10,000,000 processors. 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.

CELEBRATING 20 YEARS

1989
Jack Dongarra

Standards and efforts such as PVM, MPI, LAPACK, ScaLAPACK, BLAS, ATLAS, Netlib, Top500, PAPI, NetSolve, and the Linpack Benchmark have all left their mark on the scientific community. We can be proud of the recognition and use our software systems receive. We are continuing these efforts with PLASMA, MAGMA, HPC Challenge, BlackJack, Open MPI, and MuMI, as well as other innovative computing projects.

The Innovative Computing Laboratory continues 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 to play a major role in the development of standards for scientific computing in general.

We are building from a firm foundation. Over the past 20 years, we have developed robust research programs, 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. We continue to grow in terms of the resources we have at our disposal. We have two of the most powerful Petaflop computing systems at our doorstep with ORNL's Jaguar and UTK's Kraken systems. 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.

This is truly a time of great excitement in the design of software and algorithms for the next generation, and we will be part of the 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 for enabling technologies in the world.

- Jack Dongarra, Director of the ICL

CELEBRATING 20 YEARS

TWENTY YEARS AGO in the summer of 1989, the Innovative Computing Laboratory (ICL) was established 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 here 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 center, with a staff of more than 40 researchers, students, and administrators. What began in a small office in one of UT's oldest academic buildings, ICL now occupies a large portion of a 70,000 sq. ft. wing of the newer Claxton building located in the heart of the Knoxville campus.

1989
3 of the 5 original members of ICL are included in this group of researchers in front of Ayres Hall. L to R: Sven Hammarling, ICLer Zhaojun Bai, Anne Greenbaum, ICLer Ed Anderson, Alan McKenney, ICL founder Jack Dongarra, Jeremy DuCroz (standing), Jim Demmel (sitting)

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

As we celebrate 20 years, our mission has remained unchanged since day one: being a world leader in enabling technologies and software for scientific computing. Our goals of providing leading edge tools to tackle science's most challenging high performance computing problems and playing a major role in the development of standards for scientific computing in general continue.

Since 1989, we have attracted many post-doctoral researchers and professors from multi-disciplines such as mathematics, chemistry, etc. Many of these experts came to UT specifically to work with Dr. Dongarra, which began 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 the following page for the list of some of the experts who have passed through ICL on their way to distinguished careers at other organizations and academic institutions.

1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009

Original ICL Logo

1998

2000

2009

ICL Established
Ayres Hall

ICL LANDMARKS

Move from Ayres Hall to
new Claxton Building

2007
Merger of CS and ECE
depts. to form EECS

1989
PVM
Level 3 BLAS

1992
LAPACK
BLACS

1994
MPI

1997
NetSolve
ATLAS
RIB

2000
HPL

2003
LAPACK for Clusters
HPCC

2008
PLASMA
MAGMA

1993
TOP500

1995
ScaLAPACK

1999
PAPI
HARNESS

2002
FT-MPI
GCO

2006
Open MPI
FT-LA

PVM

ATLAS &
NetSolve

PAPI

ICL RESEARCH

1992-1995
Jaeyoung Choi
Soongsil University, Korea

1995-1996
Greg Henry
Intel Corp.

2003-2004
Edgar Gabriel
University of Houston

1991-1992
Adam Beguelin
Formerly of AOL (now retired)

1995-1996
Sven Hammarling
Numerical Algorithms Group

1999-2004
Ken Roche
University of Washington

1991-2001
Clint Whaley
University of Texas, San Antonio

1995-1998
Henri Casanova
University of Hawaii, Manoa

2004-2007
Alfredo Buttari
CNRS Toulouse, France

1989-2001
Susan Blackford
Myricom

1994-2009
Keith Seymour
University College Dublin

1992-2005
Victor Eijkhout
University of Texas, Austin

1996-2001
Dorian Arnold
University of New Mexico

1990-1991
Robert van de Geijn
University of Texas, Austin

1993-2001
Antoine Petitet
ESI Group, France

1996-1997
Yves Robert
ENS-Lyon, France

2003-2006
Julien Langou
University of Colorado at Denver

1989-1991
Ed Anderson
EPA

1993-1994
Bernard Tourancheau
University of Lyon, France

1996-2006
Graham Fagg
Microsoft

1991-1994
Roldan Pozo
NIST

1995-2001
Erich Strohmaier
Lawrence Berkeley National Laboratory

2003-2006
Felix Wolf
Juelich Supercomputing Centre

2000-2001
Reed Wade
Catalyst IT

1992-1994
Richard Barrett
ORNL

1989-1990
Zhaojun Bai
University of California, Davis

1994
Frederic Desprez
ENS-Lyon, France

1999-2003
Sathish Vadhiyar
Indian Institute of Science, India

1990-1996
Bob Manchek
Stratus Technologies

1997
Francois Tisseur
Manchester University, England

1995-1996
Andy Cleary
Amazon

KEY ICL ALUMNI

1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009

RESEARCH

Remaining at the forefront of computational research requires the ability to adapt to an ever changing landscape in high performance computing.

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 20 years ago as in-depth investigations of the numerical libraries that encode the use of linear algebra in software, our research portfolio has grown extensively. We have evolved and expanded our research agenda to accomodate the aforementioned evolution of the HPC community. We now include work in high performance paralell and distributed computing. 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 2009-2010 across our main areas of focus: numerical linear algebra, high performance distributed computing, and performance analysis and benchmarking. On the following pages, we provide brief summaries of some of our efforts in each of these areas. For more detailed information about our research, visit our website - http://icl.eecs.utk.edu/.

BLACKJACK

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

ICL TEAM
David Cronk
Anthony Danalis
Jack Dongarra
Piotr Luszczek

Modern computing architectures change rapidly and exhibit high levels of complexity and heterogeneity. Developing compilers that can boost productivity while producing efficient code for these rapidly evolving targets is a difficult challenge that most existing compiler systems are unable to address in a comprehensive way. Blackjack is addressing the problem of characterizing target architectures in a rigorous and systematic manner, providing a comprehensive list of benchmarks and applications that can stress and evaluate the compilers as well as assess the productivity benefits of such compiler systems. Our Blackjack project is part of the DARPA AACE (Architecture-Aware Compiler Environment) effort, which aims to produce modular compilers that can discover their environment and automatically adapt to it, so that it can efficiently optimize programs without user involvement or expertise.

FT-LA

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

ICL TEAM
George Bosilca
Jack Dongarra
Peng Du
Hatem Ltaief
Piotr Luszczek

Fault-tolerance has never been so paramount with distributed machines currently reaching up to 300K cores. The scientific community has to tackle the problem from two directions. First, efficient middleware needs to be designed to detect failures. Second, the numerical applications have to be flexible enough to permit the recovery of the lost data-structure. At ICL, we have successfully developed Fault Tolerant MPI (FT-MPI) middleware and, more recently, a Fault Tolerant Linear Algebra (FT-LA) library that will efficiently handle several process failures.

We have also fully integrated FT-LA in the Coordinated Infrastructure for Fault Tolerant Systems (CIFTS) environment (see p. 12) to provide better communication and fault management between the system's software components and scientific applications. Our 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 in the context of tile algorithms (PLASMA). In case of failures, the core idea would be to restart the computation by mostly using critical information already present in the directed acyclic graph generated by those factorizations, which will considerably decrease the checkpoint sizes.

FT-MPI AND OPEN MPI

URLS
http://icl.eecs.utk.edu/ftmpi/
http://www.open-mpi.org/

ICL TEAM
Wesley Bland
George Bosilca
Aurelien Bouteiller
Jack Dongarra
Thomas Herault
Pierre Lemarinier
Teng Ma
Brian Sheely

Message Passing has become the dominant programming paradigm for high performance parallel applications. ICL's expertise in this area has led to the development of a leading edge MPI library called FT-MPI, which allows for flexible new models of fault tolerance and recovery that were previously impossible. Since the release of the FT-MPI runtime library at SC 2003, 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. Many features from FT-MPI, such as point-to-point messaging, tuned collective communication algorithms and a heterogeneous data-type engine, have been integrated to the open source production quality MPI implementation known as Open MPI, which is part of a collaborative effort involving several institutions including ICL. Additionally, some fault tolerant mechanisms designed in the context of FT-MPI are currently considered by the MPI Forum for inclusion in the next version of the MPI standard (MPI 3.0).

GCO

URL
http://icl.eecs.utk.edu/gco/

ICL TEAM
Jack Dongarra
Shirley Moore

Empirical optimization (or auto-tuning) has emerged as an effective technique for addressing the limitations of traditional compiler technology while achieving performance levels competitive with hand tuning. Current empirical optimization techniques such as ATLAS and FFTW can achieve exceptional performance in part because the algorithms to be optimized are known ahead of time. We are addressing this limitation by extending the techniques used in ATLAS to the optimization of arbitrary code. Since the algorithm to be optimized is not known in advance, compiler technology is required to analyze the source code and generate the candidate implementations. We have developed an empirical tuning infrastructure that works in conjunction with existing code generators such as the ROSE LoopProcessor from Lawrence Livermore National Laboratory and POET from the University of Texas, San Antonio. Our most recent research has involved investigating different search heuristics and a parallel search strategy to improve the process of finding the best code implementation.

GRIDSOLVE

URL
http://icl.eecs.utk.edu/gridsolve/

ICL TEAM
Yuanshun (Shaun) Dai
Jack Dongarra
Don Fike
Asim YarKhan

The goal of the GridSolve project is to provide an easy-to-use Grid middleware infrastructure that will enable users to solve complex scientific problems on remote Grid resources while using simple interfaces from their preferred Scientific Computing Environments (SCEs). The GridSolve system is a GridRPC-based, client-agent-server system that includes service registration, service discovery, load balancing and service level fault tolerance. Recent work in GridSolve has focused on workflow management, where a workflow is inferred from a sequence of service calls, and the data movement is transparently managed between server machines. The current GridSolve software release includes automatic workflow inference and execution, improved scheduling capabilities, improved client interfaces (Matlab, Octave, IDL, C and Fortran), and various performance enancements.

HPC CHALLENGE

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

ICL TEAM
Jack Dongarra
Piotr Luszczek

The HPC Challenge (HPCC) benchmark suite has been designed to assess the bounds on the performance of many real applications. 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. The sustained floating point operation rate and memory bandwidth, the rate of random memory updates, and the interconnect latency and bandwidth are the major tests included in the suite. The most recent version of the code was released in 2007 and added a number of algorithmic variants of the tests. The additions prepare the code for the yearly HPCC competition where results are announced at the annual SC conference. The competition features contestants who submit performance numbers from the world's largest supercomputer installations and implementations of the benchmark suite that use a vast array of parallel programming environments. The performance results submitted through the HPCC web site and for the competition are publicly available to help track the progress of both the high end computing arena as well as commodity hardware for parallel computing.

HPL

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

ICL TEAM
Jack Dongarra
Piotr Luszczek

HPL is a software package that solves a (random) dense linear system in double precision (64 bits) arithmetic on distributed-memory computers. Nearly four years since the last stable release, HPL 2.0 was released in September 2008. Besides including major bug fixes and accuracy enhancements that have been reported by users since 2004, the major focus of this release was to improve accuracy of reported benchmark results and ensure scalability of the code on the largest supercomputer installations with hundreds of thousands of computational cores. 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 run on large computer installations to produce a result that can be submitted to the TOP500 list of the world's most powerful supercomputers.

LAPACK AND SCALAPACK

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

ICL TEAM
Jack Dongarra
Julie Langou

LAPACK and ScaLAPACK are libraries for solving dense linear algebra problems and are very widely used in the scientific community. ICL has been a major contributor to the development and maintenance of these two packages over the years. LAPACK is sequential, relies on the BLAS library, and benefits from the multicore BLAS library whereas ScaLAPACK is parallel distributed and relies on BLAS, LAPACK, MPI, and BLACS libraries. LAPACK 3.2.1 was released in April 2009. Recent work on LAPACK has revolved around the implementation of Rectangle Full Packed format, QR updating, LDLT with rook pivoting, variance of factorization, several iterative refinement routines and much more. ScaLAPACK 1.8.0 was released in April 2007. Recent work on ScaLAPACK has focused on the externalization of the LAPACK routines and the addition of new drivers. Since 2007, a special effort has been made to release and maintain LAPACK, CLAPACK and ScaLAPACK under Windows, natively. Maintenance releases for LAPACK are expected to be released by the end of 2009. ScaLAPACK 1.9.0 is also scheduled to be released in late 2009, which includes the addition of the MRRR algorithm.

MAGMA

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

ICL TEAM
Jack Dongarra
Peng Du
Rajib Nath
Stanimire Tomov

MAGMA's (Matrix Algebra on GPU and Multicore Architectures) goal is to create a new generation of linear algebra libraries that achieves the fastest possible time to an accurate solution on heterogeneous/hybrid systems, using all available processing power within given energy constraints. The main focus is the development of a dense linear algebra library for multicore+GPU systems. MAGMA is designed to be similar to LAPACK in functionality, data storage, and interface, in order to allow scientists to effortlessly port any of their LAPACK-relying software components to take advantage of the new architectures.

In addition, MAGMA is being designed to run on homogeneous x86-based multicores and to take advantage of GPU components (if available). This is achieved by developing a class of hybrid algorithms that split the computation into tasks of varying granularity (e.g., large for available GPUs) and dynamically scheduling their execution. In addition to increasing the performance proportionally to the computational characteristics of the GPUs used, MAGMA overcomes bottlenecks associated with memory-bound operations (as in the two-sided factorizations) leading to speedups exceeding even an order of magnitude.

MUMI

URL
Coming soon

ICL TEAM
Jack Dongarra
Heike Jagode
Shirley Moore
Fengguang Song
Dan Terpstra

MuMI, the Multicore Application Modeling Infrastructure, is an effort to develop a unifying framework within which an understanding of performance and performance-power tradeoffs can be developed for modern multicore architectures. This project is focused on developing a framework to facilitate systematic measurement, modeling, and prediction of performance, power consumption and performance-power tradeoffs for multicore systems. In addition to this development work, MuMI will be used to model, analyze and optimize performance and power consumption of key benchmarks and applications on multicore systems. MuMI combines the strengths of our current PAPI effort with Texas A&M's Prophesy performance modeling interface and Virginia Tech's Power-Pack, a power-performance measurement and analysis system.

PAPI

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

ICL TEAM
Jack Dongarra
Heike Jagode
Shirley Moore
Phil Mucci
James Ralph
Brian Sheely
Dan Terpstra

PAPI, the Performance API, 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 3rd party tools. PAPI continues to be ported to the most interesting new Cray and IBM architectures, as well as multicore offerings from Intel and AMD.

Architecturally, PAPI has also been restructured to allow simultaneous access to both on-processor and off-processor counters and sensors. Component PAPI, or PAPI-C, is provided with components for limited network counter and system health monitoring support. Additional components are in development by 3rd parties around the world. Development support documentation is available at the PAPI website.

PLASMA AND MULTICORE

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

ICL TEAM
Emmanuel Agullo
Jack Dongarra
Bilel Hadri
Josh Hoffman
Jakub Kurzak
Hatem Ltaief
Piotr Luszczek
Fengguang Song
Asim YarKhan

PLASMA (Parallel Linear Algebra Software for Multicore Architectures) is ICL's flagship project for multicore computing. PLASMA is designed to give high efficiency on homogeneous multicore processors and multi-socket systems of multicore processors. As of 2009, the majority of such systems are on-chip symmetric multiprocessors with classic super-scalar processors as their building blocks (x86 and alike) augmented with short-vector SIMD extensions (SSE and alike). PLASMA is being engineered with the objective to eventually supersede LAPACK and ScaLAPACK, principally by restructuring the software to achieve much greater efficiency, where possible, on modern computers based on multicore processors. PLASMA also relies on new or improved algorithms. PLASMA can solve dense systems of linear equations and linear least squares problems and associated computations such as matrix factorizations, matrix multiplication and the like. Unlike LAPACK, PLASMA currently does not solve eigenvalue or singular value problems and does not support band matrices. For all supported types of computation the same functionality is provided for real and complex matrices in single precision and double precision.

TOP500

URL
http://www.top500.org/

ICL TEAM
Jack Dongarra

For 15 years now, a ranking of the top 500 fastest computers in the world has been compiled biannually with published results released in June and November. The basis for this list is computer performance 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 rankings.

OTHER RESEARCH EFFORTS

DOE SciDAC PERI

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

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 University of Tennessee (ICL).

DOE SciDAC CIFTS

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

The Coordinated Infrastructure for Fault Tolerant Systems (CIFTS) is providing a coordinated infrastructure that will enable Fault Tolerant Systems to adapt to faults occurring in the operating environment in a holistic manner. CIFTS is a collaborative effort between Argonne National Laboratory, Lawrence Berkeley National Laboratory, Oak Ridge National Laboratory, Indiana University, Ohio State University, and University of Tennessee (ICL).

DOE SciDAC CSCADS

URL
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 support of common software infrastructures and standards. CScADS is a collaborative effort between the Argonne National Laboratory, Rice University, the University of California at Berkeley, the University of Tennessee (ICL), and the University of Wisconsin.

NSF KEENELAND

URL
http://keeneland.gatech.edu/

As part of a new NSF-funded effort, nicknamed Keeneland, ICL is partnering with Georgia Tech, Oak Ridge National Laboratory (ORNL), UTK's National Institute for Computational Sciences (NICS), Hewlett-Packard (HP) and NVIDIA Corporation in developing and deploying innovative and experimental high-performance computing (HPC) systems. Current HPC systems are designed using hundreds or thousands of central processing units (CPUs), the processor type used for performing calculations in modern day computers. These new HPC systems will be hybrids in which graphical processing units (GPUs) will work in tandem with CPUs to take advantage of two important GPU characteristics: next-generation GPU peak FLOP (floating point operation per second) rates are significantly higher than for CPUs and GPUs have proven to be more energy efficient than CPUs. A critical component of this project, which will be one of the foci of our efforts, is the development of mathematical libraries that leverage NVIDIA Corporation's CUDA-based (compute-unified device architecture) GPUs.

NSF POINT

URL
http://www.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 to the tools for users at different levels of expertise are available, and the project has a comprehensive outreach and training component. POINT is a collaborative effort between the University of Oregon, the University of Tennessee (ICL), the National Center for Supercomputing Applications, and the Pittsburgh Supercomputing Center.

NSF FUTUREGRID

URL
http://www.futuregrid.org/

FutureGrid is a high performance grid test bed that will allow scientists to work collaboratively to develop and test novel approaches to parallel, grid, and cloud computing. The FutureGrid test bed will be composed of a high-speed network connected to distributed clusters of high performance computers and will be linked to the TeraGrid—the NSF's national cyberinfrastructure of high performance computing resources for scientific research. Led by Indiana University, FutureGrid is a collaborative effort that includes the University of Tennessee (ICL), Purdue University, San Diego Supercomputer Center at University of California San Diego, University of Chicago/Argonne National Labs, University of Florida, University of Southern California Computing Center, University of Virginia, and the Center for Information Services and GWT-TUD from Technische Universtit Dresden.

PEOPLE

CELEBRATING 20 YEARS

Our most important asset, as with most organizations, is our staff.

Our success continues to hinge on our ability to skillfully apply our expertise to the computing challenges that confront the ever changing HPC landscape. With a staff of top researchers from all around the world, we can apply a variety of unique and innovative approaches to the challenges and problems inherent to world-class, scientific computing.

In addition, 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. We have been, and will continue to be, very proactive in securing internships and assistantships for those students who are hard working and willing to learn.


2001


2002


2003


2004


2005


2006


2007


2008

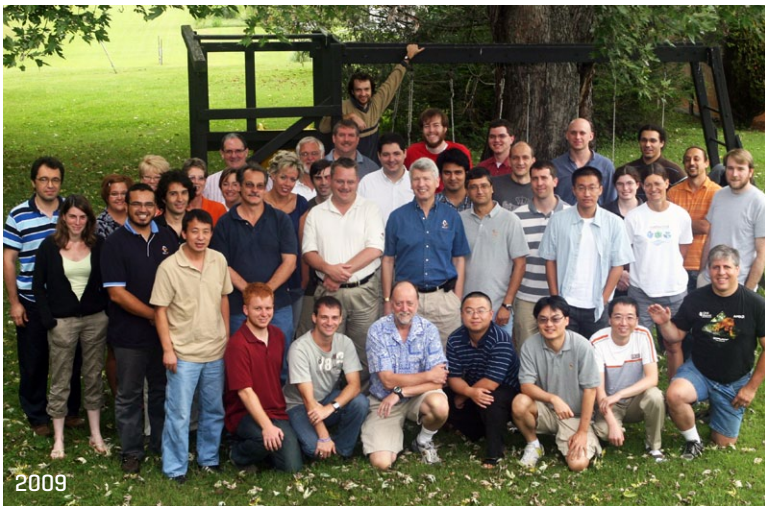
2009

CURRENT STAFF AND STUDENTS

Emmanuel Agullo
Senior Research Associate

Wesley Bland
Graduate Research Assistant

George Bosilca
Research Assistant Professor

Aurelien Bouteiller
Senior Research Associate

David Cronk
Research Director

Yuanshun (Shaun) Dai
Assistant Professor

Anthony Danalis
Senior Research Associate

Jack Dongarra
University Distinguished Professor

Peng Du
Graduate Research Assistant

Mathieu Faverge
Post-doc Research Associate

Don Fike
Research Assistant

Teresa Finchum
Administrative Services Assistant

Azzam Haidar
Post-doc Research Associate

Thomas Herault
Visiting Scientist

Josh Hoffman
Undergrad Student Assistant

Heike Jagode
Senior Research Associate

Jakub Kurzak
Senior Research Associate

Julie Langou
Research Associate

Tracy Lee
Accounting Specialist III

Pierre Lemarinier
Senior Research Associate

Hatem Ltaief
Senior Research Associate

Piotr Luszczek
Research Scientist

Teng Ma
Graduate Research Assistant

Leighanne Mills
Administrative Support Assistant

Shirley Moore
Research Consultant

Terry Moore
Associate Director

Phil Mucci
Research Consultant

Rajib Nath
Graduate Research Assistant

Paul Peltz
IT Administrator II

Tracy Rafferty
Coordinator II

James Ralph
Graduate Research Assistant

David Rogers
Graphic Designer

Brian Sheely
Research Associate II

Matt Skinner
Research Assistant

Gwang Son
Graduate Research Assistant

Fengguang Song
Graduate Research Assistant

Dan Terpstra
Research Leader

Stanimire Tomov
Research Scientist

Scott Wells
Assistant Director for Communications

Vince Weaver
Post-doc Research Associate

Asim YarKhan
Senior Research Associate

Brian Zachary
IT Administrator I

VISITORS

TWENTY-FIRST CENTURY RESEARCH agendas in the HPC landscape increasingly require collaborative initiatives to be successful. For this reason, our list of research collaborators and partners continues to grow. A byproduct of our relationship with multiple institutions is the enormous opportunities to host and work with researchers from all around the world. 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. In addition, it is not uncommon to have students (undergraduate as well as graduate) from various universities study with us for months on end, learning about our approaches and solutions to computing problems. We believe the experience of sharing expertise between our visitors and ourselves during these visits has been mutually beneficial and we will continue providing opportunities for visits from our national and international colleagues in research.

Phil Andrews
UT/National Institute for Computational Sciences

Richard Barrett
Oak Ridge National Laboratory

Swen Bohm
University of Reading, UK

Aaron Collier
General Atomics

Camille Coti
University of Paris-Sud, France

Gilles Fedak
Laboratoire de Recherche en Informatique
Universite Paris XI, France

Len Freeman
University of Manchester, UK

Erika Fuentes
Microsoft

Patrick Geoffray
Myricom

Akila Gothandaraman
UT/EECS

Sven Hammarling
Numerical Algorithms Group (NAG), UK

Thomas Herault
University of Paris-Sud, France

Sven Holmquist
Synective Labs AB, Sweden

Aurélie Hurault
Institut National Polytechnique de Toulouse, France

David Keffer
UT/Chemical & Biomolecular Engineering

Andreas Knüpfer
Center for Information Services and High Performance Computing, Dresden Germany

Frank Lauer
University of Reading, UK

Julien Langou
University of Denver at Denver

Jeff Larkin
Cray, Inc.

Antonina Litvinova
University of Reading, UK

Osni Marques
University of California, Berkeley

Michael Resch
High Performance Computing Center Stuttgart (HLRS), Germany

Tabitha Samuel
UT/EECS

Olaf Storaasli
Oak Ridge National Laboratory

Volodymyr O. Turchenko
University of Calabria, Cosenza, Italy

Felix Wolf
Julich Supercomputing Centre, Germany

Jin Yi
Shanghai University, China

ALUMNI

OVER THE PAST 20 YEARS, ICL has employed many students and staff from multiple backgrounds and academic disciplines. From undergraduate students to experienced research scientists, nearly 200 individuals have passed through our doors, many of whom have gone on to both achieve and contribute great things in academia, government, and industry. Many of our former staff and students are now faculty members at some of the most prestigious academic institutions in the US and around the world.

In addition, many of our former students have gone on to apply their experience and knowledge to careers with many of the largest companies in the computing industry including Amazon, Cray, Google, Hewlett Packard, IBM, Intel, Microsoft, Myricom, and NEC just to name a few. And many have even gone on to form companies of their own. See our 20 year timeline on page 5 for more information about former ICLers.

Aebischer, Carolyn 1990-2005
Agrawal, Bivek 2003-2006
Agrawal, Sudesh 2000-2006
Allgeyer, Jennifer 1993
Anderson, Ed 1989-1991
Andrzejewski, Daniel 2007
Angskun, Thara 2003-2007
Arkhurst, Papa 2003
Arnold, Dorian 1996-2001
Baboulin, Marc 2006-2008
Bai, Zhaojun 1989-1990
Balakrishnan, Ashwin 2002
Barrett, Richard 1992- 1994
Bassi, Alex 2000-2001
Battle, David 1990-1992
Beck, Micah 1997-2001
Beguelin, Adam 1991-1992
Benzoni, Annamaria 1991
Berry, Tom 1991
Betts, Scott 1997-1998
Bhatia, Nikhil 2003-2005
Black, Laura 1996
Black, Noel 2002-2003
Blackford, Susan 1989-2001
Bodanki, Kartheek 2008-2009
Bolt, David 1991
Bond, Fernando 1993-1995
Bowers, Carolyn 1992
Britt, Barry 2007-2009
Brown, Randy 1998-1999
Browne, Bonnie 2009
Browne, Cynthia 2003-2008
Browne, Murray 1998-1999
Bukovsky, Antonin 1999-2003
Bunch, Greg 1994-1996
Buttari, Alfredo 2004-2007
Canon, Louis 2006
Casanova, Henri 1995-1998
Chakrabarty, Ramkrishna '05-'06
Chambers, Sharon 1998-2000
Chen, Zizhong 2000-2006
Choi, Jaeyoung 1992-1995
Chrabakh, Wahid 1999
Clarkson, Eric 1998-1999
Cleary, Andy 1995-1996
Clinard, Michelle 1989-1991
Cortese, Tom 2002-2009
Coti, Camille 2007-2009
Cox, Jason 1993-1997
Deane, Cricket 1998-1999
Delmas, Remi 2006
Desprez, Frederic 1994
Ding, Jun 2001-2003
Ding, Jin 2003
Ding, Ying 2000-2001
Do, Martin 1993-1994
Dong, Leon 2000-2001
Doolin, David 1997
Dongarra, Nick 2000
Downey, Andrew 2000-2003
Drake, Mary 1989-1992
Driggs, Julio 2002-2004
Drum, Brian 2001-2004
Echavarria, Eduardo 2005
Eijkhout, Victor 1992-2005
Ellis, Brett 1994-2005
Ericson, Shawn 2004-2005
Eyler-Walker, Zachary '97-'98
Ezzell, Lisa 2003-2004
Fagg, Graham 1996-2006
Feng, Shengzhog 2005-2006
Finger, Mike 1997
Fischer, Markus 1995-1998
Fu, Xiaoquan 2003-2004
Fuentes, Erika 1999-2007
Fuerlinger, Karl 2007-2008
Fuller, Megan 2006-2008
Gabriel, Edgar 2003-2004
Gangwer, Lynn 2000-2001
Gangwer, Tracy 1992-1993
Garner, Kelley 1998
Garner, Nathan 1996-2004
Garrison, Tina 1991
Gelfend, Boris 1993
Gettler, Jonathan 1994-1996
Greaser, Eric 1991-1993
Green, Stan 1992-1996
Gregory, Alice 2004-2006
Gurley, Jason 1997-1998
Hadri, Bilel 2008-2009
Hagewood, Hunter 2000-2001
Halloy, Christian 1996-1997
Hammarling, Sven 1995-1996
Hammond, J. Mike 1994-1995
Hastings, Chris 1994-1996
Henderson, David 1999-2001
Henry, Greg 1995-1996
Hicks, Holly 1993-1994
Hill, Sid 1994-2000
Ho, George 1998-2000
Horner, Jeff 1994-1999
Huang, Yan 2000-2001
Hurt, Chris 2002
Jacobs, Paul 1991-1992
Jeannot, Emmanuel 2006
Ji, Weizhong 2000
Jiang, Weicheng 1992-1995
Jin, Song 1997-1998
Jolly, Sean 1997-1998
Jones, Jan 1992-2008
Jones, Kim 1996-1997
Kakani, Venkata 2007-2009
Kalhan, Ajay 1995
Kannan, Balajee 2000-2001
Karnowski, Larry 1998
Kasam, Madhuri 2007-2008
Katz, David 2002
Kelly, Joshua 2000-2001
Kilambi, Supriya 2008
Kim, Youngbae 1992-1996

GROWTH OF ICL

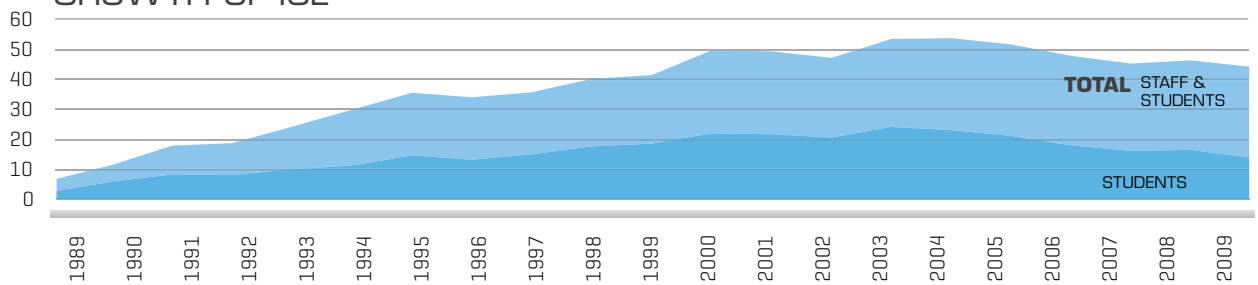

ICL STAFF/STUDENT NATIONALITIES

TOTAL COUNTRIES REPRESENTED **28**

DOMESTIC **55%**

INTERNATIONAL **45%**

Kolatis, Michael 1995-1996
Krintz, Chandra 1999-2001
Kyle, Coire 2005
Laake, Amanda 2003-2004
Langou, Julien 2003-2006
Larkin, Jeff 2002-2005
LaRose, Brian 1990-1992
Lee, DongWoo 2001-2002
Letsche, Todd 1993-1994
Lewis, Sharon 1994-1995
Li, Yinan 2006-2008
Li, Weiran 2002
Li, Xiang 2001-2003
Liu, Chaoyang 2000-2001
London, Kevin 1996-2005
Longley, Matt 1999
Lucio, Daniel 2005-2008
Luczak, Richard 2000-2001
Manchek, Robert 1990-1996
Marwah, Tushti 2004
McCasland, Donald 1994
McMahan, Paul 1994-2000
Meek, Eric 2002-2006
Meyering, James 1991-1992
Millar, Jeremy 1999-2002
Miller, Michelle 1999-2003
Mitchell, Cindy 2001-2002
Monty, Stuart 1993
Moore, Erik 2000
Moore, Keith 1991-2007
Morgan, Robert 1990-1991
Motheramgari, Kishan 1997
Moulton, Steven 1989 -1993
Nabity, Matthew 1998-2008
Narasimhaswami, Shankar '04-'05
Newton, Peter 1994-1995
Oestergaard, Jakob 2000
Papadopoulos, Caroline 1997-1998
Parker, Leelinda 2002
Patlolla, Dilip 2007-2008
Pearson, Andy 1989-1991
Pepin, Theresa 1994
Petitet, Antoine 1993-2001
Pjesivac, Vlado 2008
Pjesivac-Grbovic, Jelena '03-'07
Plank, James S. 1991
Poore, Tim 2009
Pozo, Roldan 1991-1994
Pulatova, Farzona 2005-2006
Race, Tammy 1999-2001
Raman, Ganapathy 1998-2000
Ramani, Kamesh 2003
Ran, Mei 1999-2004
Rattan, Arun 1997
Reagan, Sheri 1995-1996
Reynolds, Mike 1994
Richardson, Jon 1990-1991
Robert, Yves 1996-1997
Roche, Ken 1999-2004
Rogers, Andrew 1997-1999
Rothrock, Tom 1997-1998
Rowan, Tom 1993-1995
Sagi, Kiran 2000-2005
Samanta, Supratik 1998-1999
Sams, Evelyn 1998-1999
Schwartz, Ken 1992-1993
Seymour, Keith 1994-2009
Shahnaz, Farial 2001
Shi, Zhiao 2001-2007
Shinkarev, Sergei 2005-2007
Sidani, Majed 1991-1992
Singhal, Shilpa 1998
Skinner, Matt 2008-2009
Soendergaard, Peter 2000
Spencer, Thomas 1999-2001
Strohmaier, Erich 1995-2001
Sun, Xiaobai 1995
Swany, Martin 1996
Talley, Judi 1993-1999
Tang, Yuan 2005-2006
Teranishi, Keita 1997-1998
Thomas, Joe 2002-2009
Thurman, John 1997-1999
Tisseur, Francoise 1997
Toth, Jude 1993-1994
Tourancheau, Bernard 1993-1994
Vaca, Lauren 2004
Vadhiyar, Sathish 1999-2003
van de Geijn, Robert 1990-1991
Vawter, Chad 1995
Venckus, Scott 1994-1995
Wade, Reed 1990-1996
Wade, Reed 2000-2001
Walters, Michael 2000-2005
Waltz, Mike 1999
Waltz, Robert 1990-1991
West, David 1990-1992
Whaley, R. Clint 1991-2001
Whisnant, Jody 1997-1998
White, James 1999
Whitmire, Scotti 1995-1996
Wo, Ling 2000-2001
Wolf, Felix 2003-2005
Wu, Jiayi 2004-2007
Xia, Qiu 2003-2005
Xu, Tinghua 1998-2000
Yang, Tao 1999
You, Haihang 2004-2009
Zhang, Yuanlei 2001-2005
Zhao, Junlong 2002
Zheng, Yong 2001
Zhou, Long 2000-2001
Zhou, Min 2002-2004

PARTNERSHIPS

From the beginning, ICL has fostered relationships with many other academic institutions and research centers. In addition, we have aggressively sought to build lasting, collaborative partnerships with both HPC vendors and industry leaders here and abroad.

These businesses and institutions have helped us build a solid foundation of meaningful and lasting relationships that have significantly contributed to our efforts to be a world leader in enabling technology research. Leveraging the incredible growth of computational science research, we also routinely develop relationships with researchers whose primary focus is 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. In addition, many application and tool vendors have utilized our work. These include Intel, MathWorks, Etnus, SGI, and Cray. In addition, Hewlett Packard, IBM, Intel, 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 follwing page, we recognize many of the partners and collaborators that we have worked with over the years, most of which 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.

INTERNATIONAL COLLABORATORS

- **Danish Computing Center for Research and Education** Lyngby, Denmark
- **Dept of Mathematical and Computing Sciences** Tokyo Institute of Technology
- **University of Manchester** Manchester, England
- **European Centre for Research and Advanced Training in Scientific Computing** Toulouse, France
- **École Polytechnique Fédérale de Lausanne** Lausanne, Switzerland
- **Fakultät für Mathematik und Informatik, Universität Mannheim** Mannheim, Germany
- **Forschungszentrum Jülich Central Institute for Applied Mathematics** Jülich, Germany
- **High Performance Computing Center Stuttgart (HLRS)** Stuttgart, Germany
- **Institut ETH Zentrum** Zurich, Switzerland
- **International Exascale Software Project**
- **Intelligent Systems Design Laboratory, Doshisha University** Kyoto, Japan
- **Istituto per le Applicazioni del Calcolo "Mauro Picone"** Rome, Italy
- **Kasetsart University** Bangkok, Thailand
- **Laboratoire de L'Informatique du Parallelisme, École Normal Superieure de Lyon** Lyon, France
- **Mathematical Insitute, Utrecht University** Netherlands
- **Numerical Algorithms Group Ltd** Oxford, England
- **Parallel and HPC Application Software Exchange** Tsukuba, Japan
- **Laboratoire RESAM, Université Claude Bernard de Lyon** Lyon, France
- **Rutherford Appleton Laboratory** Oxford, England
- **Soongsil University** Seoul, South Korea
- **Technische Universitaet Wien** Vienna, Austria
- **University of Umeå** Umeå, Sweden
- **Virtual Institute High-Productivity Supercomputing** Germany

DOMESTIC COLLABORATORS

- **ANL** Argonne National Laboratory
- **Cray**
- **DARPA** Defense Advanced Research Projects Agency
- **DoD** The United States Department of Defense
- **DOE** United States Department of Energy
- **Emory University**
- **Georgia Institute of Technology**
- **Hewlett Packard**
- **IBM** International Business Machines
- **Indiana University**
- **Intel Corporation**
- **ISI** Information Sciences Institute
- **JICS** The Joint Institute for Computational Science
- **LANL** Los Alamos National Laboratory
- **LLNL** Lawrence Livermore National Laboratory
- **The MathWorks**
- **Microsoft Research**
- **MRA** MetaCenter Regional Alliance
- **NASA** National Aeronautics and Space Administration
- **NCSA** The National Computational Science Alliance
- **NICS** National Institute for Computational Sciences
- **NIST** National Institute of Standards and Technology
- **NSF** The National Science Foundation
- **NVIDIA Corporation**
- **ORNL CSMD** Computer Science and Mathematics Division, Oak Ridge National Lab.
- **Rice University**
- **SDSC** San Diego Supercomputing Center
- **SGI** Silicon Graphics Incorporated
- **Sun Microsystems**
- **University of California, Berkeley**
- **University of California, San Diego**
- **UTK-EECS** Univ. of Tennessee Department of Electrical Engineering and Computer Science

HARDWARE RESOURCES

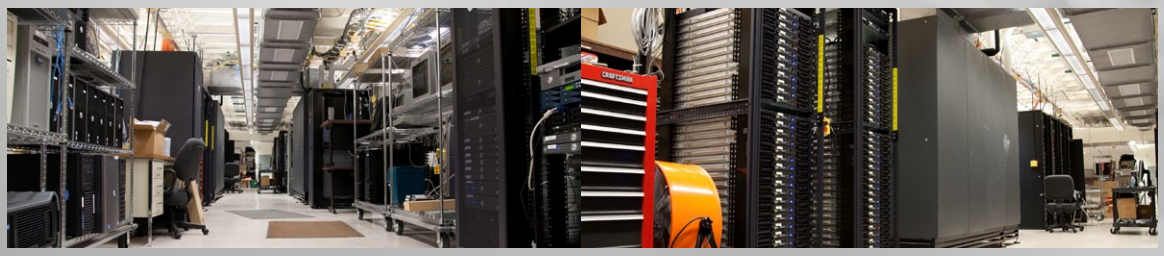

EECS/ICL Machine Room

Changes in the landscape of high performance computing reflect constant new developments in hardware and technologies that give HPC machines their computational power.

But 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 configuratrions present unique challenges that require testing and development of applications that are often quite unique to the platform on which they reside. For this reason, it is imperative that we have access to a very wide range of computing resources in order to conduct our cutting-edge research. On this front, we have multiple, heterogeneous systems in-house. But we also have access to multiple architectures around the country, due in large part to our many collaborators and partners. Locally, we maintain systems ranging from individual desktops to large, networked clusters. Below is a summary of the many computing resources used by ICL.

The following are the local systems that we use on a daily basis to test our work:

64 node (128 cores) Intel EM64T cluster connected with Myrinet 2000
62 node (124 cores) AMD Opteron cluster connected with Myrinet 2000
24 node (96 cores) AMD Opteron cluster connected with Silverstorm and Mellanox Infiniband
8 node (16 cores) Intel Core2 Duo cluster with GigE
8 node (128 cores) AMD Opteron cluster connected with Myrinet 10G
5 node (10 cores) Intel Itanium cluster
SiCortex (72 cores) SC072-PDS

In addition to these resources, we have access to several server class machines and several HPC clusters within the EECS department. These clusters consist of multiple architectures and comprise over 100 machines

CELEBRATING 20 YEARS

TORC Cluster, 2000
ICL's first cluster:
8 Dual PIII 550Mhz/512MB RAM
100Mbit, Myrinet, Giganet

Cray XT5 Jaguar at Oak Ridge National Laboratory

Cray XT5 Kraken at Oak Ridge National Laboratory

with various architectures. All of our clusters are arranged in the classic Beowulf configuration in which machines are connected by low latency, high-speed network switches.

Also, access to many remote resources, some that are regularly found in the TOP500 list of the world's fastest supercomputers, help keep us at the forefront of enabling technology research. The recent modernization of the DOE's Center for Computational Sciences, just 30 minutes away at the Oak Ridge National Laboratory (ORNL), has enabled us to leverage our ORNL collaborations to take advantage of what is becoming the world's fastest scientific computing facility. UT's National Center for Computational Sciences (NCCS) at ORNL houses Kraken, UT's Cray XT5 system that is now the fastest open-science supercomputer in the world, as well as Jaguar, another Cray XT5 that was the second fastest supercomputer in the world in 2009. The following are some of the remote systems and architectures that we utilize:

Cray X1E, XT3, XT4, and XT5
HP XC system
IBM Power 5, 5+, 6, Cluster 1600, BlueGene/L, BlueGene/P and the Cell
Several large (512+ proc) Linux Clusters
SGI Altix

2009 PUBLICATIONS

"Computational Science – ICCS 2009, Proceedings of the 9th International Conference," *Lecture Notes in Computer Science: Theoretical Computer Science and General Issues*, Allen, G., Nabrzyski, J., Seidel, E., van Albada, G.D., Dongarra, J., Sloot, P.M.A. eds. Baton Rouge, LA, Parts I-II, Vols. 5544-5545, May 25-27, 2009.

Agullo, E., Demmel, J., Dongarra, J., Hadri, B., Kurzak, J., Langou, J., Ltaief, H., Luszczek, P., Tomov, S. **"Numerical linear algebra on emerging architectures: The PLASMA and MAGMA projects,"** *Journal of Physics: Conference Series*, Vol. 180, 2009.

Agullo, E., Hadri, B., Ltaief, H., Dongarra, J. **"Comparative Study of One-Sided Factorizations with Multiple Software Packages on Multi-Core Hardware,"** *University of Tennessee Computer Science Technical Report* (also *LAPACK Working Note* 217), UT-CS-09-640, April, 2009.

Agullo, E., Hadri, B., Ltaief, H., Dongarra, J. **"Comparative Study of One-Sided Factorizations with Multiple Software Packages on Multi-Core Hardware,"** *2009 International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '09)* (to appear), 2009.

Alam, S. R., Barrett, R. F., Jagode, H., Kuehn, J. A., Poole, S. W. and Sankaran, R. **"Impact of Quad-core Cray XT4 System and Software Stack on Scientific Computation,"** *Euro-Par 2009, Lecture Notes in Computer Science*, Delft, The Netherlands, Springer Berlin / Heidelberg, Volume 5704/2009, pp. 334-344, August 25-28, 2009.

Alvaro, W., Kurzak, J., Dongarra, J. **"Optimizing Matrix Multiplication for a Short-Vector SIMD Architecture - CELL Processor,"** *Parallel Computing*, Volume 35, pp. 138-150, 2009.

Baboulin, M., Buttari, A., Dongarra, J., Kurzak, J., Langou, J., Langou, J., Luszczek, P., Tomov, S. **"Accelerating Scientific Computations with Mixed Precision Algorithms,"** *Computational Physics Communication* (to appear), 2009.

Baboulin, M., Dongarra, J., Gratton, S., Langou, J. **"Computing the Conditioning of the Components of a Linear Least-squares Solution,"** *Numerical Linear Algebra with Applications*, Vol. 16, No. 7, pp. 517-533, 2009.

Bosilca, G., Coti, C., Herault, T., Lemarinier, P., Dongarra, J. **"Constructing resiliant communication infrastructure for runtime environments,"** *Innovative Computing Laboratory Technical Report*, ICL-UT-09-02, July 31, 2009.

Buttari, A., Dongarra, J., Kurzak, J., Langou, J. **"Parallel Dense Linear Algebra Software in the Multicore Era,"** in *Cyberinfrastructure Technologies and Applications*, Cao, J. eds. Nova Science Publishers, Inc., pp. 9-24, 2009, 2009.

Buttari, A., Langou, J., Kurzak, J., Dongarra, J. **"A Class of Parallel Tiled Linear Algebra Algorithms for Multicore Architectures,"** *Parallel Computing* (to appear), 2009.

Cunha, M., Telles, J., YarKhan, A., Dongarra, J. **"Grid Computing applied to the Boundary Element Method,"** *Proceedings of the First International Conference on Parallel, Distributed and Grid Computing for Engineering*, Topping, B.H.V., Ivanyi, P. eds. Stirlingshire, UK, Civil-Comp Press, Paper 27, 2009.

Dai, Y.-S., Dongarra, J. **"Reliability and Performance Modeling and Analysis for Grid Computing,"** in *Handbook of Research on Scalable Computing Technologies* (to appear), Li, K.-C., Hsu, C.-H., Yang, L. T., Dongarra, J., Zima, H. eds. IGI Global, pp. 219-245, 2009.

Danalis, A., Pollock, L., Swany, M., Cavazos, J. **"MPI-aware Compiler Optimizations for Improving Communication-Computation Overlap,"** *Proceedings of the 23rd annual International Conference on Supercomputing (ICS '09)*, Yorktown Heights, NY, USA, ACM, pp. 316-325, June, 2009.

de Supinski, B., Alam, S., Bailey, D., Carrington, L., Daley, C., Dubey, A., Gamblin, T., Gunter, D., Hovland, P., Jagode, H., Karavanic, K., Marin, G., Mellor-Crummey, J., Moore, S., Norris, B., Oliker, L., Olschanowsky, C., Roth, P., Schulz, M., Shende, S., Snavely, A. **"Modeling the Office of Science Ten Year Facilities Plan: The PERI Architecture Tiger Team,"** *SciDAC 2009, Journal of Physics: Conference Series*, San Diego, California, IOP Publishing, Vol. 180 (2009) 012039, June 14-18, 2009.

Dongarra, J., Beckman, P., Aerts, P., Cappello, F., Lippert, T., Matsuoka, S., Messina, P., Moore, T., Stevens, R., Trefethen, A., Valero, M. **"The International Exascale Software Project: A Call to Cooperative Action by the Global High Performance Community,"** *International Journal of High Performance Computing Applications* (to appear), July, 2009.

Dongarra, J., Bosilca, G., Delmas, R., Langou, J. **"Algorithmic Based Fault Tolerance Applied to High Performance Computing,"** *Journal of Parallel and Distributed Computing*, Volume 69, pp. 410-416, 2009.

Dongarra, J., Meuer, H., Simon, H., Strohmaier, E. **"Recent Trends in High Performance Computing,"** in *Birth of Numerical Analysis* (to appear), 2009.

Gustavson, F., Wasniewski, J., Dongarra, J., Langou, J. **"Rectangular Full Packed Format for Cholesky's Algorithm: Factorization, Solution and Inversion,"** *ACM TOMS* (to appear), 2009.

Hadri, B., Ltaief, H., Agullo, E., Dongarra, J. **"Tall and Skinny QR Matrix Factorization Using Tile Algorithms on Multicore Architectures,"** *Innovative Computing Laboratory Technical Report* (also *LAPACK Working Note* 222 and *CS Tech Report* UT-CS-09-645), ICL-UT-09-03, September 4, 2009.

Hoefler, T., Dai, Y-S., Dongarra, J. **"Towards Efficient MapReduce Using MPI,"** *Lecture Notes in Computer Science, Recent Advances in Parallel Virtual Machine and Message Passing Interface - 16th European PVM/MPI Users' Group Meeting*, Ropo, M., Westerholm, J., Dongarra, J. eds. Espoo, Finland, Springer Berlin / Heidelberg, Volume 5759, pp. 240-249, 2009.

Jagode, H., Dongarra, J., Alam, S., Vetter, J., Spear, W., Malony, A. **"A Holistic Approach for Performance Measurement and Analysis for Petascale Applications,"** *ICCS 2009 Joint Workshop: Tools for Program Development and Analysis in Computational Science and Software Engineering for Large-Scale Computing*, G. Allen et al. eds. Baton Rouge, Louisiana, Springer-Verlag Berlin Heidelberg 2009, ICCS 2009, Part II, LNCS 5545, pp. 686-695, May 25-27, 2009.

Evidence of our research and our contributions to the HPC community might be best exemplified by the numerous publications we annually produce.

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.

Jagode, H., Knuepfer, A., Dongarra, J., Jurenz, M., Mueller, M. S., Nagel, W. E. "**Trace-based Performance Analysis for the Petascale Simulation Code FLASH**," *Innovative Computing Laboratory Technical Report*, ICL-UT-09-01, April 15, 2009.

Jagode, H., Moore S., Terpstra, D., Dongarra, J., Knuepfer, A., Jurenz, M., Mueller, M. S., Nagel, W. E. "**I/O Performance Analysis for the Petascale Simulation Code FLASH**," *ISC'09*, Hamburg, Germany, June 23 - 26, 2009.

Kurzak, J., Dongarra, J. "**Fully Dynamic Scheduler for Numerical Computing on Multicore Processors**," *University of Tennessee Computer Science Department Technical Report*, UT-CS-09-643 (Also *LAPACK Working Note* 220), 2009.

Kurzak, J., Dongarra, J. "**QR Factorization for the CELL Processor**," *Scientific Programming* (to appear), 2009.

Kurzak, J., Ltaief, H., Dongarra, J., Badia, R. "**Scheduling Linear Algebra Operations on Multicore Processors**," *University of Tennessee Computer Science Department Technical Report*, UT-CS-09-636 (Also *LAPACK Working Note* 213), 2009.

Kurzak, J., Ltaief, H., Dongarra, J., Badia, R. "**Scheduling Linear Algebra Operations on Multicore Processors**," *Concurrency Practice and Experience* (to appear), 2009.

Kurzak, J., Ltaief, H., Dongarra, J., Badia, R. "**Scheduling Linear Algebra Operations on Multicore Processors**," in *Trends in High Performance and Large Scale Computing* (to appear), Grandinetti, L., Joubert, G., Gentzsch, W. eds. eds. IOP Press, 2009.

Kurzak, J., Ltaief, H., Dongarra, J., Badia, R. "**Dependency-Driven Scheduling of Dense Matrix Factorizations on Shared-Memory Systems**," *PPAM 2009*, Poland, September, 2009.

Langou, J., Dongarra, J. "**The Problem with the Linpack Benchmark Matrix Generator**," *International Journal of High Performance Computing Applications*, Volume 23, No. 1, pp. 5-14, Spring, 2009.

Lemarinier, P., Bosilca, G., Coti, C., Herault, T., Dongarra, J. "**Constructing Resilient Communication Infrastructure for Runtime Environments**," *ParCo 2009*, Lyon France, September, 2009.

Li, Y., Dongarra, J., Tomov, S. "**A note on auto-tuning GEMM for GPUs**," *University of Tennessee Computer Science Technical Report*, UT-CS-09-635 (also *LAPACK Working Note* 212), January 12, 2009.

Li, Y., Dongarra, J., Tomov, S. "**A Note on Auto-tuning GEMM for GPUs**," *Computational Science – ICCS 2009, Proceedings of the 9th International Conference, Lecture Notes in Computer Science: Theoretical Computer Science and General Issues*, Allen, G., Nabrzyski, J., Seidel, E., van Albada, G.D., Dongarra, J., Sloot, P.M.A. eds. Baton Rouge, LA, Parts I-II, Vols. 5544-5545, pp. 884-892, May 25-27, 2009.

Ltaief, H., Kurzak, J., Dongarra, J. "**Scheduling Two-sided Transformations using Algorithms-by-Tiles on Multicore Architectures**," *University of Tennessee Computer Science Technical Report*, UT-CS-09-637 (also *LAPACK Working Note* 214), February 11, 2009.

Ltaief, H., Kurzak, J., Dongarra, J. "**Parallel Band Two-Sided Matrix Bidiagonalization for Multicore Architectures**," *IEEE Transactions on Parallel and Distributed Systems* (to appear), May, 2009.

Luszczek, P. "**Parallel Programming in MATLAB**," *The International Journal of High Performance Computing Applications*, Vol. 23, No. 3, pp 277-283, July, 2009.

Portillo, R., Teller, P., Cronk, D., Moore, S., et. al. "**Making Performance Analysis and Tuning Part of the Software Development Cycle**," *Proceedings of DoD HPCMP UGC 2009*, San Diego, CA, IEEE, June, 2009.

Ramakrishan, L., Nurmi, D., Mandal, A., Koelbel, C., Gannon, D., Huang, M., Kee, Y-S., Obertelli. G., Thyagaraja, K., Wolski, R., YarKhan, A., Zagorodnov, D. "**VGrADS: Enabling e-Science Workflows on Grids and Clouds with Fault Tolerance**," *SC'09 The International Conference for High Performance Computing, Networking, Storage and Analysis* (to appear), Portland, OR, 2009.

Seymour, K., YarKhan, A., Dongarra, J. "**Transparent Cross-Platform Access to Software Services using GridSolve and GridRPC**," in *Cloud Computing and Software Services: Theory and Techniques* (to appear), Ahson, S., Ilyas, M. eds. eds. CRC Press, 2009.

Song, F., Moore, S., Dongarra, J. "**Analytical Modeling and Optimization for Affinity Based Thread Scheduling on Multicore Systems**," *IEEE Cluster 2009*, New Orleans, Aug. 31 - Sept. 4, 2009.

Song, F., Moore, S., Dongarra, J. "**A Scalable Non-blocking Multicast Scheme for Distributed DAG Scheduling**," *The International Conference on Computational Science 2009 (ICCS 2009)*, Baton Rouge, LA, LNCS 5544, pp. 195-204, May, 2009.

Song, F., YarKhan, A., Dongarra, J. "**Dynamic Task Scheduling for Linear Algebra Algorithms on Distributed-Memory Multicore Systems**," *University of Tennessee Computer Science Technical Report*, UT-CS-09-638, April 13, 2009.

Song, F., YarKhan, A., Dongarra, J. "**Dynamic Task Scheduling for Linear Algebra Algorithms on Distributed-Memory Multicore Systems**," *International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '09)* (to appear), Portland, OR, November 14-20, 2009.

Tomov, S., Dongarra, J. "**Accelerating the Reduction to Upper Hessenberg Form Through Hybrid GPU-based Computing**," *University of Tennessee Computer Science Technical Report*, UT-CS-09-642 (also *LAPACK Working Note* 219), May 24, 2009.

Tomov, S., Lu, W., Bernholc, J., Moore, S., Dongarra, J. "**Performance evaluation for petascale quantum simulation tools**," *Proceedings of CUG09*, Atlanta, GA, May 4-7, 2009.

Youseff, L., Seymour, K., You, H., Zagorodnov, D., Dongarra, J., Wolski, R. "**Paravirtualization Effect on Single- and Multi-threaded Memory-Intensive Linear Algebra Software**," *Cluster Computing Journal: Special Issue on High Performance Distributed Computing*, Springer Netherlands, Volume 12, No. 2, pp. 101-122, June 2009.

2009 CONFERENCES/ SEMINARS/WORKSHOPS

JAN 13 Columbus, OH Air Force Research Lab (AFRL) CELL/Multicore Programming Tutorial	JAN 21 Argonne, IL Coordinated Infrastructure for Fault Tolerance Systems (CIFTS) Meeting	JAN 29-30 Oak Ridge, TN FastOS Meeting
FEB 6 Atlanta, GA HARNESS Meeting	FEB 9-11 San Jose, CA MPI Forum	FEB 16-20 Juelich, Germany 3rd VI-HPS Tuning Workshop
FEB 25-27 Louisville, KY ORTE and FastOS Meeting	MAR 2-6 Miami, FL SIAM Conference on Computational Science and Engineering (CSE09)	MAR 9-12 Boulder, CO 10th LCI International Conference on High-Performance Clustered Computing
MAR 19-20 Albuquerque, NM 2009 Fault Tolerance Workshop	MAR 23 Blacksburg, VA Virginia Tech Network Dynamics and Simulation Sciences Laboratory Seminar	MAR 23-26 Pittsburgh, PA PSC/Intel Multi-core Programming and Performance Tuning Workshop
MAR 25-27 Marina del Rey (Los Angeles), CA PERI All Hands Meeting, Spring 2009	APR 7-8 Santa Fe, NM International Exascale Software Project	APR 4-6 Chicago, IL MPI Forum and STCI Meeting
APR 13-16 Oak Ridge, TN ORNL NCCS/NICS 2009 Cray XT5 Quad-core Workshop	MAY 3-8 Dagstuhl, Germany Fault Tolerance in High-Performance Computing and Grids	MAY 4-7 Atlanta, GA Compute the Future
MAY 12-14 Austin, TX 2009 PET Technical Overview	MAY 13-15 Knoxville, TN Scheduling for Large-scale Systems	MAY 19-21 Monterey and Mountain View, CA Baseline Configuration Team / parCFD
MAY 25-27 Baton Rouge, LA International Conference on Computational Science 2009	JUN 1-5 Portland, OR SC09 June	JUN 8-10 Menlo Park, CA MPI Forum

Every year, our research staff regularly attend national and international conferences, workshops, and seminars.

These meetings provide ample opportunities to present our research, share our knowledge, and exchange ideas with leading computational science researchers from around the world. On these pages is a listing of the most recent events we have participated in.

JUN 9 Arlington, VA DARPA AACE Characterization Meeting	JUN 14-18 San Diego, CA Scientific Discover Through Advanced Computing (SciDAC 2009)	JUN 15-18 San Diego, CA DoD Users Group Conference (USG) 2009
JUN 22-25 Arlington, VA TeraGrid 09	JUN 23-26 Hamburg, Germany International Supercomputing Conference ISC '09	JUN 28-29 Paris, France International Exascale Software Project (IESP) Workshop
JUL 27-29 Chicago, IL MPI Forum	AUG 3-5 Athens, GA International Association for Mathematics and Computers in Simulation	AUG 6 Nashville, TN DoD MOS/HPTi PETTT Transition Meeting
AUG 10-12 Houston, TX CScADS 2009, Workshop on Libraries and Autotuning for Petascale Applications	AUG 20 Woodbridge, VA DoD HPCMP Utility Server Requirements Gathering Meeting and Engineering (CSE09)	AUG 31-SEP 4 New Orleans, LA IEEE Cluster Conference
SEP 1-4 Lyon, France International Conference on Parallel Computing 2009	SEP 2-10 Espoo, Finland Euro PVM/MPI 2009	SEP 9 Edinburgh, UK Scottish Computational Mathematics Symposium
SEP 14-15 Dresden, Germany 3rd Parallel Tools Workshop	SEP 13-16 Wroclaw, Poland 8th International Conference on Parallel Processing and Applied Mathematics (PPAM)	SEP 24 Oak Ridge, TN MPI I/O Meeting
SEP 29-30 College Park, MD DOE PERI All Hands Meeting	OCT 2-3 Indianapolis, IN FutureGrid All-Hands Meeting	OCT 13-14 Santa Fe, NM Los Alamos Computer Science Symposium
OCT 17-21 Tsukuba, Japan International Exascale Software Project (IESP) Workshop	NOV 11-13 Portland, OR MPI-Forum	NOV 14-20 Portland, OR SC09

OTHER EFFORTS AND ACTIVITES

http://www.exascale.org/

As next generation computing ushers in new hardware architectural designs that include hundreds of thousands of nodes and millions of cores, the demands on the software infrastructure necessary to address the computationally and data intensive work demanded by tomorrow's scientists and engineers are certain to increase. Closing the gap between the existing software infrastructure and these new demands requires an unprecedented level of coordination and cooperation within the worldwide open source software community. To address this issue, the IESP was created in 2009 with the guiding purpose of empowering 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.

In its inaugural year, the IESP held three workshops at different locations around the globe, which ICL helped organize: San Antonio, TX (USA); Paris, France; and Tsukuba, Japan. The agendas for each workshop were structured in order 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 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. More information about the IESP including meeting notes, whitepapers, and presentations can be found by visiting the project website at http://www.exascale.org/.

In addition to the development of tools and applications, ICL is regularly engaged in other activites and efforts that include leading conferences and workshops as well as teaching and outreach.

Having a leadership role in the HPC community requires that we be engaged with the community and actively share our vision for the next generations of computing that lie just beyond the horizon. In this section are some of the activities that we are participating in or have a leadership role in.

CITR
CENTER FOR INFORMATION TECHNOLOGY RESEARCH

http://citr.cs.utk.edu/

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 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. One, it is 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 two, it is 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.

IGMCS

http://igmcs.utk.edu/

Interdisciplinary Graduate Minor in Computational Science

Since first rate students and staff are indispensable to the success of any academic research strategy, CITR worked with faculty and administrators from several departments and colleges in 2007 to help establish a new university-wide program in Computational Science that supports advanced degree concentrations in this critical new area 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 will be 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 2009, there were 15 different departments from four university colleges contributing nearly 100 courses to the program. In addition, the IGMCS student enrollment has doubled over the past year.

NVIDIA CUDA Center of Excellence

In late 2009, ICL was awarded a CUDA Center of Excellence (CoE) by NVIDIA Corporation, a world-wide leader in technologies for visual computing and inventor of the graphical processing unit (GPU). As part of the award and CoE designation, ICL received hardware, financial support, and other resources from NVIDIA. Joining a very small and select group of CUDA CoEs such as labs at Harvard University, the University of Utah, and the University of Illinois at Urbana-Champaign, UT's CoE will focus on the development of linear algebra libraries for CUDA-based hybrid architectures. Our work on matrix algebra on GPU and multicore architectures (MAGMA) will further enable and expand our CUDA-based software library efforts, especially in the general area of high-performance scientific computing.

RWTH Aachen University,
Center for Computing and Communication

Forschungszentrum Jülich,
Jülich Supercomputing Centre

Technische Universität München,
Chair for Computer Architecture

Technische Universität Dresden, Center for Information Services and High Performance Computing

University of Tennessee
Innovative Computing Laboratory

University of Stuttgart,
High Performance Computing Centre

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 plus 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,

"The VI-HPS is proud to have had ICL's partnership from the very beginning. ICL stands for best-of-breed HPC software, satisfying the highest quality standards. Collaborating with ICL for many years now on the Scalasca project, a performance tool specifically designed for large-scale parallel systems, has been a very enjoyable and productive experience. Moreover, with PAPI, a well-recognized and widely-used performance-counter library, ICL contributes one of their flagship products to the VI-HPS tool suite, complementing the remaining tools in an ideal way. Finally, with their involvement in many leading HPC projects in the US, ICL represents an important link in the transatlantic cooperation needed to address today's challenges in scalable computing."

http://www.vi-hps.org/

SPONSORS

FOR 20 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 for their support of our efforts past and present:

Defense Advanced Research Projects Agency (DARPA)
Department of Defense (DoD)
Department of Energy (DOE)
National Aeronautics and Space Administration (NASA)
National Science Foundation (NSF)
National Institutes of Health (NIH)
Office of Naval Research (ONR)

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 for their generosity and their significance to our success:

INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY of TENNESSEE

Department of Electrical Engineering and Computer Science

Suite 413 Claxton
1122 Volunteer Blvd
Knoxville, TN 37996-3450

PHONE (865) 974-8295

E-MAIL icl@eecs.utk.edu

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


INNOVATIVE
COMPUTING LABORATORY
THE UNIVERSITY of TENNESSEE
http://icl.eecs.utk.edu/