ICL UT

INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY of TENNESSEE

2010/2011 REPORT

INNOVATIVE COMPUTING LABORATORY **2010/2011** REPORT

EDITED BY **Scott Wells** DESIGNED BY **David Rogers**

INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY of TENNESSEE

2010/2011 REPORT

FROM THE DIRECTOR

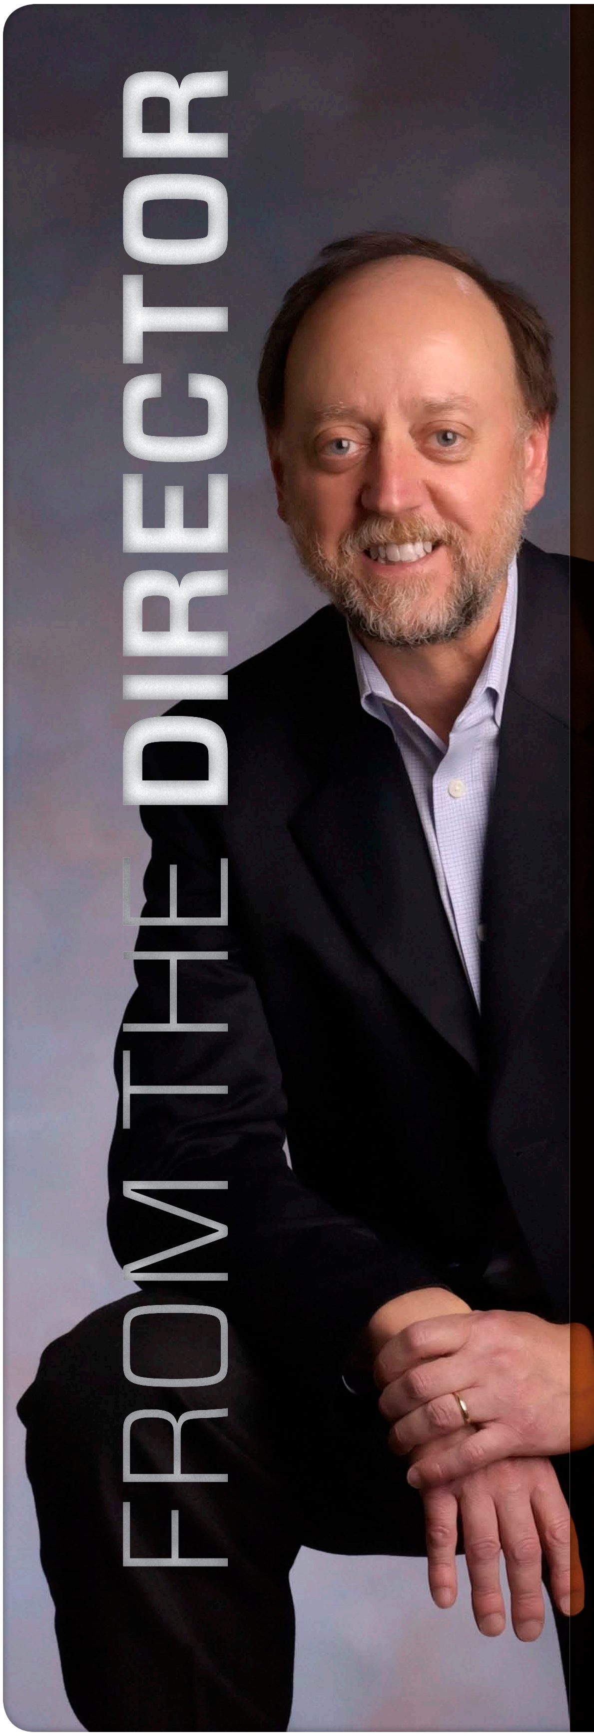

This is an exciting time for computing. This year marks the beginning of the road to exascale computing. 'Going to the exascale' will mean radical changes in computing architecture, software, and algorithms – basically, vastly increasing the levels of parallelism to the point of billions of threads working 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 life of Computational Science revolves around a multifaceted software ecosystem. But today there is (and should be) a real concern that this ecosystem of Computational Science, with all its complexities, is not ready for the major challenges that will soon confront the field. Domain scientists now want to create much larger, multi-dimensional applications in which a variety of previously independent models are coupled together, or even fully integrated. They hope to be able to run these applications on Peta and Exascale systems with millions to billions of threads of execution, to extract all performance that these platforms can deliver, to recover automatically from the processor failures that regularly occur at this scale, and to do all this without sacrificing good programmability. This vision of what Computational Science wants to become contains numerous unsolved and exciting problems for the software and algorithm research community. Unfortunately, it also highlights aspects of the current software environment that are either immature, under funded, or both.

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. 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 are helping to maintain the balanced software ecosystem. 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 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, perhaps a once in a life time 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

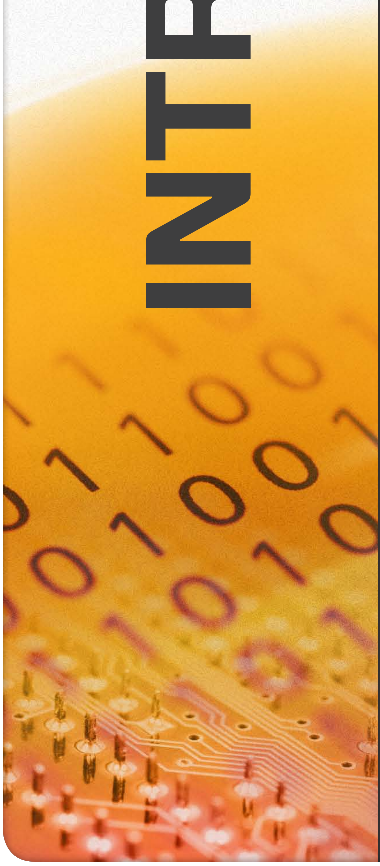

As one of University of Tennessee's oldest and largest research laboratories, the Innovative Computing Laboratory (ICL) has been unwavering in its mission since the beginning: being a world leader in enabling technologies and software for scientific computing. Our goals of providing leading edge tools 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.

ICL was established in 1989 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. 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 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.

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, HPCC
- **2006** Open MPI, FT-LA
- **2008** PLASMA, MAGMA
- **2009** Blackjack, MuMI

KEY ICL ALUMNI

- **1991-1992** Adam Beguelin — FORMERLY OF AOL (NOW RETIRED)
- **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 — UNIVERSITY COLLEGE, DUBLIN
- **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 — JUELICH SUPERCOMPUTING CENTRE
- **1992-1994** Richard Barrett — ORNL
- **1996-1997** Yves Robert — ENS-LYON, FRANCE
- **2000-2001** Reed Wade — CATALYST IT
- **2009** Emmanuel Agullo — INRIA, FRANCE
- **1989-1990** Zhaojun Bai — UNIVERSITY OF CALIFORNIA, DAVIS
- **1995-1996** Greg Henry — INTEL CORP.
- **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
- **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, our research portfolio has grown extensively. 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 2010-2011 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.

BLACKJACK

The 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 with minimal user involvement and expertise. Modern computing architectures change rapidly and exhibit high levels of complexity and heterogeneity. Developing compilers that can boost productivity while producing efficient, optimized code for these rapidly evolving targets is a difficult challenge.

As part of the Blackjack project we have been developing a system benchmark suite, Blackjack-Bench, that can automatically characterize target architectures in a rigorous and systematic manner. We are also working on providing a comprehensive list of benchmarks and applications that can stress and evaluate the compilers, as well as assessing the productivity and performance benefits of such compiler systems. The target release of the initial benchmark suite is slated for late 2010.

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

DAGUE

DAGuE (Directed Acyclic Graph Unified Environment) 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 with labeled edges designating data dependencies. DAGs are represented in a compact, problem-size independent format that 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.

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

RESEARCH PROJECTS

FT-LA

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 structures. 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 to provide better communication and fault management between the system's software components and scientific applications. This work has been released and is freely available on our website. 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 that are used in PLASMA (see p.11). In case of failures, the main idea is to restart the computation by mostly using critical information already present in the directed acyclic graph generated by these factorizations, which will considerably decrease the checkpoint sizes.

FIND OUT MORE AT http://icl.eecs.utk.edu/ft-la/

FT-MPI/Open MPI

Message Passing is 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. Some fault tolerance 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). 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.

FIND OUT MORE AT http://icl.eecs.utk.edu/ftmpi/
http://www.open-mpi.org/

HPC CHALLENGE

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 June 2010 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 the high end computing arena as well as commodity hardware for parallel computing.

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

HPL

HPL, short for High Performance Linpack, is a software package that solves a (random) dense linear system in double precision (64 bits) arithmetic on distributed-memory computers. HPL 2.0, released in 2008, includes 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. In addition, you can now run the LINPACK benchmark on iPhone and Android smart phones.

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

SCA/LAPACK

LAPACK (Linear Algebra PACKage) and ScalableLAPACK 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 3.2.2 was released in June 2010. Recent work on LAPACK has revolved around new functionalities (e.g., SYTRS with BLAS 3, CS Decomposition). The LAPACK team is planning new usability features and an autotuning framework. In 2010, a special effort was made to release a native C Interface to LAPACK in collaboration with INTEL.

The next major release for LAPACK, expected in late 2010, comes three years after the release of v. 1.8.0, which was released in April 2007. ScaLAPACK 1.9.0 is also scheduled to be released in late 2010, which will include the addition of the MRRR algorithm.

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

MAGMA

The goal of the Matrix Algebra on GPU and Multicore Architectures (MAGMA) project is to create a new generation of linear algebra libraries that achieve 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.

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 scheduling their execution over the available hardware components. MAGMA 1.0 is a release that concentrates on fundamental linear algebra algorithms for multicore enhanced with a single GPU. This includes both the basic one and two-sided matrix factorizations, as well as linear systems and eigen/singular-value solvers based on them. Routines are provided in four precisions - single, double, single complex, and double complex. Linear systems solvers are provided in both working precision and mixed-precision using iterative refinement.

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

MUMI

MuMI, or Multicore Application Modeling Infrastructure, is a project that 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. MuMI combines ICL'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. During the first year of the project, PAPI was 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. Facilities are being developed within the PAPI project to enable the user to define modeling metrics at a high level and have these metrics mapped to underlying hardware events and characteristics.

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

PAPI

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 architectures of greatest interest to the High Performance Computing community.

Architecturally, PAPI provides simultaneous access to both on-processor and off-processor counters and sensors. The latest versions of PAPI, also called Component PAPI, or PAPI-C, support components for network counters and system health monitoring, as well as disk subsystems. Additional components are in development by 3rd parties around the world. Development support documentation is available at the PAPI website.

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

PLASMA

PLASMA (Parallel Linear Algebra Software for Multicore Architectures) is ICL's flagship project for multicore and manycore computing. The software is designed to deliver high performance from heterogeneous multi-socket, multicore systems with GPU accelerators. PLASMA achieves this objective by combining state-of-the-art solutions in algorithms, scheduling and software engineering. Currently, PLASMA offers routines for solving linear systems and least square problems. This includes mixed-precision routines for taking advantage of single precision speed, while delivering double precision accuracy. Also included are routines for fast parallel processing of very tall and thin matrices based on tree reduction patterns, very fast routines for explicitly forming an inverse of a matrix, and a full set of Level 3 BLAS operations on matrices stored in a tile layout. PLASMA will soon offer GPU support and routines for solving symmetric eigenvalue problems and singular value problems. PLASMA is available with easily accessible documentation in the form of traditional manuals, as well as an array of user friendly online resources.

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

TOP 500

Since 1993, a ranking of the top 500 fastest computers in the world has been compiled biannually with published results released in June and November. 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.

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

AGENCY **RESEARCH**

DOE SciDAC PERI

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). 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.

FIND OUT MORE AT http://www.peri-scidac.org/

DOE SciDAC CIFTS

Coordinated Infrastructure for Fault Tolerant Systems (CIFTS) 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. Our focus here at ICL 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.

FIND OUT MORE AT

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

DOE SciDAC CScADS

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. Our effort in this project focuses on on re-engineering numerical libraries for future HPC systems, in particular 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. Highly asynchronous codes are a good fit for the massive amount of concurrency present in multicore machines.

FIND OUT MORE AT http://cscads.rice.edu/

NSF Keeneland

Keeneland is a new, $12 million cyberinfrastructure project, funded under the NSF's Track 2D program. ICL is partnering with project leader Georgia Tech, as well as Oak Ridge National Laboratory (ORNL), UTK's National Institute for Computational Sciences (NICS), Hewlett-Packard (HP) and NVIDIA Corporation, in order to develop and deploy Keeneland's innovative and experimental system. Current high performance computing (HPC) systems are designed using hundreds or thousands of central processing units (CPUs), which are the type of general purpose processor used in all kinds of modern day computers, ranging from notebooks to supercomputers. As HPC systems scale up, there is good reason to believe that, like Keeneland, they 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. ICL will focus on the the critical area of mathematical libraries, which must be completely renovated in order to leverage the NVIDIA Corporation's CUDA-based (compute-unified device architecture) GPUs that Keeneland possses

FIND OUT MORE AT http://keeneland.gatech.edu/

NSF 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 (TAU), the University of Tennessee - ICL (PAPI), the National Center for Supercomputing Applications (PerfSuite), and the Pittsburgh Supercomputing Center (applications). 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.

FIND OUT MORE AT http://nic.uoregon.edu/point/

NSF FutureGrid

FutureGrid is a high performance test bed for scientific grids and clouds that will allow application writers to work collaboratively to develop and test novel approaches to parallel, grid, and cloud computing. The FutureGrid test bed will be composed of high-speed networks 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. As the project is led by the Indiana University, FutureGrid is also 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 in Germany.

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

PEOPLE

ICL Group, Fall 2010

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 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 students who are hard working and willing to learn.

STAFF & STUDENTS


Wes Alvaro
GRADUATE RESEARCH ASSISTANT


Dulceneia Becker
SENIOR RESEARCH ASSOCIATE


Wesley Bland
GRADUATE RESEARCH ASSISTANT


George Bosilca
RESEARCH ASSISTANT PROFESSOR


Aurelien Bouteiller
SENIOR RESEARCH ASSOCIATE


Bonnie Browne
NON-UT STUDENT ASSISTANT


David Cronk
RESEARCH DIRECTOR


Yuanshun Dai
ASSISTANT PROFESSOR


Anthony Danalis
SENIOR RESEARCH ASSOCIATE


Tingxing Dong
GRADUATE RESEARCH ASSISTANT


Jack Dongarra
UNIVERSITY DISTINGUISHED PROFESSOR


Peng Du
GRADUATE RESEARCH ASSISTANT


Mathieu Faverge
POST DOCTORAL RESEARCH ASSOCIATE


Don Fike
RESEARCH ASSISTANT


Teresa Finchum
ADMINISTRATIVE SERVICES ASSISTANT


Azzam Haidar
POST DOCTORAL RESEARCH ASSOCIATE


Blake Haugen
GRADUATE RESEARCH ASSISTANT


Thomas Herault
VISITING SCIENTIST


Josh Hoffman
GRADUATE RESEARCH ASSISTANT


Mitch Horton
POST DOCTORAL RESEARCH ASSOCIATE


Heike Jagode
SENIOR RESEARCH ASSOCIATE


Kiran Kumar Kasichayanula
GRADUATE RESEARCH ASSISTANT


Ajay Katta
GRADUATE RESEARCH ASSISTANT

PEOPLE

STAFF & STUDENTS
CONTINUED

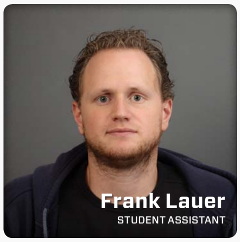

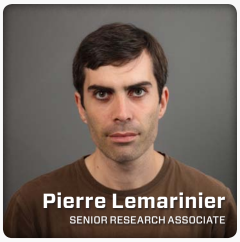

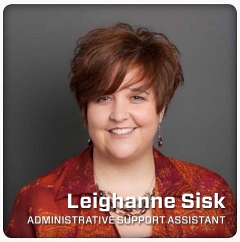

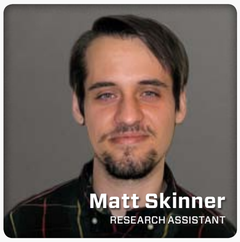

VISITORS

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. 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 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.

Emmanuel Agullo
INRIA Bordeaux, France

Cédric Augonnet
INRIA Bordeaux, France

Vincent Berthoux
France

John Drake
University of Tennessee, US

Ian Jones
UK

Guido Juckeland
Center for Information Services and High Performance Computing (ZIH), Germany

Krerkchai Kusolchu (Jom)
Suranaree University of Technology, Thailand

Tilman Kuestner
Technische Universtat Munchen, Germany

Jeff Larkin
Cray, US

Alexey Lastovetsky
University College Dublin, Ireland

John Levesque
Cray, US

Samuel "Sticks" Mabakane
Center for High Performance Computing (CHPC), South Africa

Dzung Phan
University of Florida, US

Adam Scarborough
UK

Trey White
ORNL, US

Songhua Xu
ORNL, US

Xiaohua Zhang
ORNL, US

PEOPLE ALUMNI

From the beginning, we have attracted many post-doctoral researchers and professors from diverse disciplines such as mathematics, chemistry, etc. 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 '90-'93
Sudesh Agrawal '01-'06
Bivek Agrawal '04-'06
Emmanuel Agullo '09
Jennifer Allgeyer '93
Ed Anderson '89-'91
Daniel Andrzejewski '07
Thara Angskun '03-'07
Papa Arkhurst '03
Dorian Arnold '99-'01
Cedric Augonnet '10
Marc Baboulin '08
Zhaojun Bai '90-'92
Ashwin Balakrishnan '01-'02
Richard Barrett '92-'94
Alex Bassi '00-'01
David Battle '90-'92
Micah Beck '00-'01
Daniel Becker '07
Adam Beguelin '91
Annamaria Benzoni '91
Tom Berry '91
Vincent Berthoux '10
Scott Betts '97-'98
Nikhil Bhatia '03-'05
Laura Black '96
Noel Black '02-'03
Susan Blackford '89-'01
Kartheek Bodanki '09
David Bolt '91
Fernando Bond '99-'00
Carolyn Bowers '92
Barry Britt '07-'09
Randy Brown '97-'99
Cynthia Browne '05
Murray Browne '98-'99
Antonin Bukovsky '98-'03
Greg Bunch '95
Alfredo Buttari '08
Domingo Gimenez Canovas '01
Henri Casanova '95-'98
Ramkrishna Chakrabarty '05
Sharon Chambers '98-'00
Zizhong Chen '01-'06
Jaeyoung Choi '94-'95
Wahid Chrabakh '99
Eric Clarkson '98
Andy Cleary '95-'97
Michelle Clinard '89-'91
Matthias Colin '04
Tom Cortese '09
Camille Coti '07
Jason Cox '93-'97
Javier Cuenca '03
Manoel Cunha '06
Cricket Deane '98-'99
Remi Delmas '06
Frederic Desprez '94-'95
Ying Ding '00-'01
Jun Ding '01-'03
Jin Ding '03
Martin Do '93-'94
Leon Dong '00-'01
Nick Dongarra '00
David Doolin '97
Andrew Downey '98-'03
Mary Drake '89-'92
Julio Driggs '02-'04
Brian Drum '01-'04
Eduardo Echavarria '05
Victor Eijkhout '92-'05
Brett Ellis '95-'05
Shawn Ericson '04
Zachary Eyler-Walker '97-'98
Lisa Ezzell '03-'04
Christoph Fabianek '03
Graham Fagg '96-'06
Shengzhog Feng '05-'06
Salvatore Filippone '04
Anna Finchum '10
Mike Finger '97
Markus Fischer '97-'98
Len Freeman '09
Xiaoquan Fu '03-'04
Erika Fuentes '07
Karl Fuerlinger '08
Megan Fuller '06
Edgar Gabriel '03-'04
Tracy Gangwer '92-'93
Lynn Gangwer '00-'01
Nathan Garner '01-'06
Kelley Garner '98
Tina Garrison '91
Adriana Garties '09
Christoph Geile '08
Jean Patrick Gelas '01
Boris Gelfend '93
Jonathan Gettler '96
Bruno Giuseppe '01
Eric Greaser '93
Stan Green '92-'96
Alice Gregory '04-'06
Jason Gurley '97-'98
Bilel Hadri '09
Hunter Hagewood '00-'01
Christian Halloy '96-'97
Sven Hammarling '96-'97
J. Mike Hammond '94-'95
Hidehiko Hasegawa '95-'96
Satomi Hasegawa '95-'96
Chris Hastings '96
David Henderson '99-'01
Greg Henry '96
Holly Hicks '93-'94
Alexandra Hicks-Hardiman '09
Sid Hill '96-'98
Tomoyuki Hiroyasu '02-'03
George Ho '98-'00
Jeff Horner '95-'99
Yan Huang '00-'01

Aurelie Hurault '09
Chris Hurt '02
Paul Jacobs '92-'95
Emmanuel Jeannot '06
Weizhong Ji '99-'00
Weicheng Jiang '92-'95
Song Jin '97-'98
Patrick Johansson '01
Aral Johnson '09
Sean Jolly '97-'98
Kim Jones '96-'97
Jan Jones '92-'08
Venkata Kakani '07
Ajay Kalhan '95
Balajee Kannan '01
Madhuri Kasam '07-'08
David Katz '02
Joshua Kelly '00-'01
Supriya Kilambi '08
Myung Ho Kim '05-'06
Youngbae Kim '92-'96
Jenya Kirshtein '08
Michael Kolatis '93-'96
Chandra Krintz '99-'01
Tilman Kuestner '10
Krerkchai Kusolchu '10
Coire Kyle '05
Amanda Laake '03-'04
Julien Langou '08
Jeff Larkin '03-'05
Brian LaRose '90-'92
DongWoo Lee '00-'02
Todd Letsche '93-'94
Sharon Lewis '92-'95
Weiran Li '02
Xiang Li '01
Yinan Li '09
Chaoyang Liu '00
Kevin London '96-'05
Matt Longley '99
Daniel Lucio '08
Richard Luczak '00-'01
Robert Manchek '90-'96
Tushti Marwah '04
Donald McCasland '94
Paul McMahan '94-'00
Eric Meek '03-'06
James Meyering '91-'92
Jeremy Millar '98-'02
Michelle Miller '99-'03
Cindy Mitchell '01-'02
Stuart Monty '93
Erik Moore '00
Keith Moore '87-'07
Robert Morgan '90-'91
Kishan Motheramgari '97
Steven Moulton '91-'93
Matthew Nabity '08
Shankar Narasimhaswami '04-'05
Rajib Nath '08-'10
Fernando Navarro '09
Donnie Newell '10
Peter Newton '94-'95
Jonas Nilsson '01
Jakob Oestergaard '00
Caroline Papadopoulos '97-'98
Leelinda Parker '02
Dilip Patlolla '08
Andy Pearson '89-'91
Theresa Pepin '94
Antoine Petitet '93-'01
Vlado Pjesivac '08
Jelena Pjesivac-Grbovic '03-'07
James S. Plank '91-'92
Tim Poore '09
Roldan Pozo '92-'94
Farzona Pulatova '05-'06
Martin Quinson '01
Tammy Race '99-'01
Ganapathy Raman '98-'00
Kamesh Ramani '03
Mei Ran '99-'04
Arun Rattan '97
Sheri Reagan '95-'96
Mike Reynolds '94
Jon Richardson '90-'91
Yves Robert '96-'97
Ken Roche '99-'04
Andrew Rogers '97-'99
Tom Rothrock '97-'98
Tom Rowan '93-'97
Kiran Sagi '01-'05
Evelyn Sams '98-'99
Ken Schwartz '92-'93
Keith Seymour '94-'09
Farial Shahnaz '01
Brian Sheely '09-'10
Zhiao Shi '01-'07
Sergei Shinkarev '05-'07
Majed Sidani '91-'92
Shilpa Singhal '96-'98
Peter Soendergaard '00
Gwang Son '07-'09
Thomas Spencer '99-'01
Erich Strohmaier '95-'01
Xiaobai Sun '95
Martin Swany '96-'99
Daisuke Takahashi '02
Judi Talley '93-'99
Ronald Tam '09
Yuan Tang '05-'06
Yusuke Tanimura '03
Keita Teranishi '98
Joe Thomas '09
John Thurman '98-'99
Francoise Tisseur '97
Jude Toth '93-'94
Bernard Tourancheau '93-'94
Lauren Vaca '04
Sathish Vadhiyar '99-'03
Robert van de Geijn '90-'91
Chad Vawter '95
Eugene Vecharynski '08
Scott Venckus '93-'95
Antoine Vernois '04
Reed Wade '90-'96
Michael Walters '01-'05
Mike Waltz '99
Robert Waltz '90-'91
Jerzy Wasniewski '00
David West '90-'92
R. Clint Whaley '91-'01
Jody Whisnant '97-'98
James White '99
Scotti Whitmire '95-'96
Susan Wo '00-'01
Felix Wolf '03-'05
Jiayi Wu '04-'07
Qiu Xia '04-'05
Tinghua Xu '98-'00
Tao Yang '99
Jin Yi '09-'10
Haihang You '04-'09
Lamia Youseff '07
Brian Zachary '09-'10
Yuanlei Zhang '01-'05
Junlong Zhao '02
Yong Zheng '01
Luke Zhou '00-'01
Min Zhou '02-'04

PARTNERSHIPS

Over the past 21 years, we have aggressively sought to build lasting, collaborative partnerships with both 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 significantly contributed to our efforts to be a world leader in 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, 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 follwing 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.

INTERNATIONAL PARTNERS

Danish Computing Center for Research and Education LYNGBY, DENMARK

Dept of Mathematical and Computing Sciences TOKYO INSTITUTE OF TECHNOLOGY

École Polytechnique Fédérale de Lausanne LAUSANNE, SWITZERLAND

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

European Exascale Software Initiative EUROPEAN UNION

Fakultät für Mathematik und Informatik, Universität Mannheim MANNHEIM, GERMANY

Forschungszentrum Jülich Jülich Supercomputing Centre 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

University of Manchester MANCHESTER, ENGLAND

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 PARTNERS

ANL Argonne National Laboratory

Cray

DARPA Defense Advanced Research Projects Agency

DoD 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 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 National Center for Supercomputing Applications

NICS National Institute for Computational Sciences

NIST National Institute of Standards and Technology

NSF 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

As the new GPU hybrid computing paradigm leads the evolution of computational hardware toward Petascale computing, computing architectures are increasingly changing. 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
- 16 node (96 cores) Intel Core2 Quad cluster connected with 10G Infiniband
- 8 node (128 cores) AMD Opteron cluster connected with Myrinet 10G
- SiCortex (72 cores) SC072-PDS
- Several hybrid achitectures combining CPUs and GPUs


EECS/ICL Machine Room

Cray XT5 Kraken at Oak Ridge National Laboratory


EECS/ICL Machine Room

Cray XT5 Jaguar at Oak Ridge National Laboratory

In addition to these resources, we have access to several server-class machines and HPC clusters within the EECS department. These clusters consist of multiple architectures and comprise over 100 machines 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 has become 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 fourth fastest supercomputer in the world in mid-2010. We also have access to resources on the US 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
- SGI Altix

2010 PUBLICATIONS

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.

"**Proceedings of the International Conference on Computational Science**," *ICCS 2010*, Sloot, P., van Albada, D., Dongarra, J. eds. Elsevier, Amsterdam, May 31 - June 2, 2010.

"**8th International Conference on Parallel Processing and Applied Mathematics, Lecture Notes in Computer Science (LNCS)**," *PPAM 2009 Proceedings*, Wyrzykowski, R., Dongarra, J., Karczewski, K., Wsniewski, J. eds. Springer, Wroclaw, Poland, Vol. 6067, September, 2010.

"**Recent Advances in the Message Passing Interface, Lecture Notes in Computer Science (LNCS)**," *EuroMPI 2010 Proceedings*, Keller, R., Gabriel, E., Resch, M., Dongarra, J. eds. Springer, Stuttgart, Germany, Vol. 6305, September, 2010.

Agullo, E., Augonnet, C., Dongarra, J., Faverge, M., Ltaief, H., Thibault, S., Tomov, S. "**QR Factorization on a Multicore Node Enhanced with Multiple GPU Accelerators**," *University of Tennessee Computer Science Technical Report*, ICL-UT-10-04, October 1, 2010.

Agullo, E., Coti, C., Dongarra, J., Herault, T., Langou, J. "**QR Factorization of Tall and Skinny Matrices in a Grid Computing Environment**," *24th IEEE International Parallel and Distributed Processing Symposium* (also LAWN 224), Atlanta, GA, April 19-23, 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* (submitted), March, 2010.

Agullo, E., Giraud, L., Guermouche, A., Haidar, A., Roman, J. "**Towards a Complexity Analysis of Sparse Hybrid Linear Solvers**," *PARA 2010*, Reykjavik, Iceland, June 6-9, 2010.

Agullo, E., Giraud, L., Guermouche, A., Haidar, A., Roman, J., Lee-Tin-Yen, Y. "**MaPHyS or the Development of a Parallel Algebraic Domain Decomposition Solver in the Course of the Solstice Project**," *Sparse Days 2010 Meeting at CERFACS*, Toulouse, France, June 15-17, 2010.

Angskun, T., Fagg, G., Bosilca, G., Pjesivac-Grbovic, J., Dongarra, J. "**Self-Healing Network for Scalable Fault-Tolerant Runtime Environments**," *Future Generation Computer Systems*, Vol. 26, No. 3, pp. 479-485, March, 2010.

Bernholc, J., Hodak, M., Lu, W., Moore, S., Tomov, S. "**Scalability Study of a Quantum Simulation Code**," *PARA 2010*, Reykjavik, Iceland, June 6-9, 2010.

Bosilca, G., Bouteiller, A., Danalis, A., Herault, T., Lemarinier, P., Dongarra, J. "**DAGuE: A generic distributed DAG engine for high performance computing**," *Innovative Computing Laboratory Technical Report*, ICL-UT-10-01, April 11, 2010.

Bosilca, G., Bouteiller, A., Danalis, A, Faverge, M., Haidar, H., Herault, T., Kurzak, J., Langou, J., Lemarinier, P., Ltaief, H., Luszczek, P., YarKhan, A., Dongarra, J. "**Distributed-Memory Task Execution and Dependence Tracking within DAGuE and the DPLASMA Project**," *Innovative Computing Laboratory Technical Report*, ICL-UT-10-02, 2010, 2010.

Bosilca, G., Bouteiller, A., Danalis, A., Faverge, M., Haidar, H., Herault, T., Kurzak, J., Langou, J., Lemariner, P., Ltaief, H., Luszczek, P., YarKhan, A., Dongarra, J. "**Distibuted Dense Numerical Linear Algebra Algorithms on Massively Parallel Architectures: DPLASMA**," *University of Tennessee Computer Science Technical Report*, UT-CS-10-660, Sept. 15, 2010.

Bosilca, G., Bouteiller, A., Herault, T., Lemarinier, P., Dongarra, J. "**Dodging the Cost of Unavoidable Memory Copies in Message Logging Protocols**," *Proceedings of EuroMPI 2010*, Jack Dongarra, Michael Resch, Rainer Keller, Edgar Gabriel, eds. eds. Springer, Stuttgart, Germany, September, 2010.

Bosilca, G., Coti, C., Herault, T., Lemarinier, P., Dongarra, J. "**Constructing Resiliant Communication Infrastructure for Runtime Environments in Advances in Parallel Computing**," in *Advances in Parallel Computing - Parallel Computing: From Multicores and GPU's to Petascale*, Chapman, B., Desprez, F., Joubert, G., Lichnewsky, A., Peters, F., Priol, T. Eds. eds. Vol. 19, pp. 441-451, 2010.

Bouteiller, A., Bosilca, G., Dongarra, J. "**Redesigning the Message Logging Model for High Performance**," *Concurrency and Computation: Practice and Experience (online version)*, June 27, 2010.

Brady, T., Lastovetsky, A., Seymour, K., Guidolin, M., Dongarra, J. "**SmartGridRPC: The new RPC model for high performance Grid Computing and Its Implementation in SmartGridSolve**," *Concurrency and Computation: Practice and Experience* (to appear), January, 2010.

Buttari, A., Langou, J., Kurzak, J., Dongarra, J. "**A Class of Parallel Tiled Linear Algebra Algorithms for Multicore Architectures**," *Parallel Computing* (to appear), 2010.

Dongarra, J. "**LINPACK on Future Manycore and GPu Based Systems**," *PARA 2010*, Reykjavik, Iceland, June 6-9, 2010.

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, 2010, 2010.

Dongarra, J., Beckman, P. et al. "**International Exascale Software Project Roadmap v1.0**," *University of Tennessee Computer Science Technical Report*, UT-CS-10-654, May 30, 2010.

Dongarra, J., Beckman, P. et al., "**The International Exascale Software Project Roadmap**," *University of Tennessee Computer Science Technical Report*, UT-CS-10-662, October 18, 2010.

Dongarra, J., Luszczek, P. "**Reducing the time to tune parallel dense linear algebra routines with partial execution and performance modelling**," *University of Tennessee Computer Science Technical Report*, UT-CS-10-661, October 8, 2010.

Dongarra, J., Moore, S. "**Empirical Performance Tuning of Dense Linear Algebra Software**," in *Performance Tuning of Scientific Applications* (to appear), Bailey, D., Lucas, R., Williams, S. eds. 2011, 2010.

Du, P., Luszczek, P., Dongarra, J. "**OpenCL Evaluation for Numerical Linear Algebra Library Development**," *Symposium on Application Accelerators in High-Performance Computing (SAAHPC '10)*, Knoxville, TN, July 13-15, 2010.

Du, P., Luszczek, P., Tomov, S., Dongarra, J. "**Mixed-Tool Performance Analysis on Hybrid Multicore Architectures**," *First International Workshop on Parallel Software Tools and Tool Infrastructures (PSTI 2010)*, San Diego, CA, Sept. 13-16, 2010.

Du, P., Parsons, M., Fuentes, E., Shaw, S., Dongarra, J. "**Tuning Principal Component Analysis for GRASS GIS on Multi-core and GPU Architectures**," *FOSS4G 2010*, Barcelona, Spain, Sept. 6-9, 2010.

Du, P., Weber, R., Luszczek, P., Tomov, S., Peterson, G., Dongarra, J. "**From CUDA to OpenCL: Towards a Performance-portable Solution for Multiplatform GPU Programming**," *Parallel Computing* (submitted), August, 2010.

Giraud, L., Haidar, A., Pralet, S. "**Using multiple levels of parallelism to enhance the performance of domain decomposition solvers**," *Parallel Computing*, Bekas, C., d'Ambra, P., Grama, A., Saad, Y., Yanev, P. eds. Elsevier journals, Vol. 36, No. 5-6, p285-296, 2010.

Giraud, L., Haidar, A., Saad, Y. "**Sparse approximations of the Schur complement for parallel algebraic hybrid solvers in 3D**," *Numerical Mathematics: Theory, Methods and Applications*, Zhiming, C. eds. Golbal Science Press, Beijing, Vol. 3, No. 3, pp. 64-82, 2010.

Gustavson, F., Wasniewski, J., Dongarra, J. "**Level-3 Cholesky Kernel Subroutine of a Fully Portable High Performance Minimal Storage Hybrid Format Cholesky Algorithm**," *ACM TOMS* (submitted), also LAPACK Working Note (LAWN) 211, 2010.

Gustavson, F., Wasniewski, J., Dongarra, J., Langou, J. "**Rectangular Full Packed Format for Cholesky's Algorithm: Factorization, Solution, and Inversion**," *ACM Transactions on Mathematical Software (TOMS)*, Vol. 37, No. 2, April, 2010.

Hadri, B., Agullo, E., Dongarra, J. "**Tile QR Factorization with Parallel Panel Processing for Multicore Architectures**," *24th IEEE International Parallel and Distributed Processing Symposium* (submitted), 2010.

Haidar, A., Ltaief, H., YarKhan, A., Dongarra, J. "**Analysis of Dynamically Scheduled Tile Algorithms for Dense Linear Algebra on Multicore Architectures**," Submitted to *the First International Workshop on Parallel Software Tools and Tool Infrastructures (PSTI 2010)*, San Diego, CA, April 26, 2010.

Hurault, A., YarKhan, A. "**Intelligent Service Trading and Brokering for Distributed Network Services in GridSolve**," *VECPAR 2010, 9th International Meeting on High Performance Computing for Computational Science*, Berkeley, CA, June 22-25, 2010.

Jagode, H., Knuepfer, A., Dongarra, J., Jurenz, M., Mueller, M. S., Nagel, W. E. "**Trace-based Performance Analysis for the Petascale Simulation Code FLASH**," *International Journal of High Performance Computing Applications* (to appear), 2010.

Kurzak J., Dongarra, J. "**QR Factorization for the CELL Processor**," *Scientific Programming*, Vol. 17, Nos. 1-2, pp. 31-42, 2010.

Kurzak, J., Ltaief, H., Dongarra, J., Badia, R. "**Scheduling Dense Linear Algebra Operations on Multicore Processors**," *Concurrency and Computation: Practice and Experience*, Vol. 22, no. 1, pp. 15-44, January, 2010.

Li, Y., YarKhan, A., Dongarra, J., Seymour, K., Hurault, A. "**Enabling Workflows in GridSolve: Request Sequencing and Service Trading**," *Journal of Supercomputing* (submitted), 2010.

Ltaief, H., Kurzak, J., Dongarra, J. "**Parallel Band Two-Sided Matrix Bidiagonalization for Multicore Architectures**," *IEEE Transactions on Parallel and Distributed Systems*, pp. 417-423, April, 2010.

Ltaief, H., Kurzak, J., Dongarra, J., M. Badia, R. "**Scheduling Two-sided Transformations using Tile Algorithms on Multicore Architectures**," *Journal of Scientific Computing*, Vol. 18, No. 1, pp. 33-50, 2010.

Ltaief, H., Tomov, S., Nath, R., Dongarra, J. "**Hybrid Multicore Cholesky Factorization with Multiple GPU Accelerators**," *IEEE Transaction on Parallel and Distributed Systems* (submitted), March 26, 2010.

Ltaief, H., Tomov, S., Nath, R., Du, P., Dongarra, J. "**A Scalable High Performant Cholesky Factorization for Multicore with GPU Accelerators**," *Proc. of VECPAR'10* (to appear), Berkeley, CA, June 22-25, 2010.

Luszczek, P., Dongarra, J. "**Analysis of Various Scalar, Vector, and Parallel Implementations of RandomAccess**," *Innovative Computing Laboratory (ICL) Technical Report*, ICL-UT-10-03, June, 2010.

Ma, T., Bouteiller, A., Bosicla, G., Dongarra, J. "**Locality and Topology aware Intra-node Communication Among Multicore CPUs**," *Proceedings of the 17th EuroMPI conference*, LNCS, Stuttgart, Germany, September, 2010.

Nath, R., Tomov, S., Dongarra, J. "**An Improved MAGMA GEMM for Fermi GPUs**," *International Journal of High Performance Computing* (to appear), 2010.

Nath, R., Tomov, S., Dongarra, J. "**Accelerating GPU Kernels for Dense Linear Algebra**," *Proc. of VECPAR'10*, Berkeley, CA, June 22-25, 2010.

Nath, R., Tomov, S., Dongarra, J. "**An Improved MAGMA GEMM for Fermi GPUs**," *University of Tennessee Computer Science Technical Report*, UT-CS-10-655 (also LAPACK working note 227), July 29, 2010.

Song, F., Ltaief, H., Hadri, B., Dongarra, J. "**Scalable Tile Communication-Avoiding QR Factorization on Multicore Cluster Systems**," *University of Tennessee Computer Science Technical Report*, UT-CS-10-653, April, 2010.

Song, F., Ltaief, H., Hadri, B., Dongarra, J. "**Scalable Tile Communication-Avoiding QR Factorization on Multicore Cluster Systems**," *SC'10, ACM SIGARCH/ IEEE Computer Society*, New Orleans, LA, November 13-19, 2010.

Tomov, S., Dongarra, J., Baboulin, M. "**Towards Dense Linear Algebra for Hybrid GPU Accelerated Manycore Systems**," *Parallel Computing*, Vol. 36, No. 5-6, pp. 232-240, 2010.

Tomov, S., Ltaief, H., Nath, R., Dongarra, J. "**Faster, Cheaper, Better - a Hybridization Methodology to Develop Linear Algebra Software for GPUs**," *GPU Computing Gems* (to appear), 2010.

Tomov, S., Lu., W., Bernholc, J., Moore, S., Dongarra, J. "**Performance Evaluation for Petascale Quantum Simulation Tools**," *Proceedings of the Cray Users' Group Meeting*, Atlanta, GA, May 4, 2010.

Tomov, S., Nath, R., Dongarra, J. "**Accelerating the reduction to upper Hessenberg, tridiagonal, and bidiagonal forms through hybrid GPU-based computing**," *Parallel Computing* (to appear), June 19, 2010.

Tomov, S., Nath, R., Ltaief, H., Dongarra, J. "**Dense Linear Algebra Solvers for Multicore with GPU Accelerators**," *Proc. of IPDPS'10*, Atlanta, GA, January 15, 2010.

Turchenko, V., Grandinetti, L., Bosilca, G., Dongarra, J. "**Improvement of parallelization efficiency of batch pattern BP training algorithm using Open MPI**," *Proceedings of International Conference on Computational Science, ICCS 2010* (to appear), Elsevier, Amsterdam The Netherlands, June, 2010.

CONFERENCES

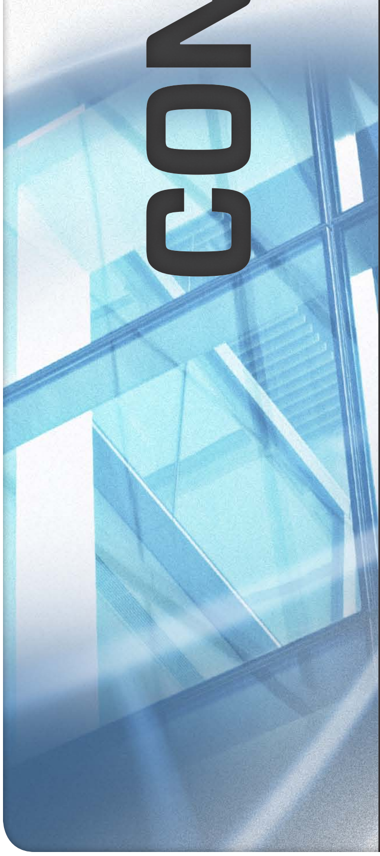

Every year, our research staff regularly attend national and international conferences, workshops, and seminars. These meetings provide opportunities to present our research, share our knowledge, and exchange ideas with leading computational science researchers from around the world. On these pages is a listing of events we have participated in over the past year.

- **NOV 11–13** 2009 — **MPI Forum**, Portland, OR
- **NOV 12** 2009 — **AACE/Blackjack Characterization Meeting**, Washington, DC
- **NOV 14–20** 2009 — **Supercomputing 2009**, Portland, OR
- **DEC 3–4** 2009 — **GPU Computing Seminar**, College Station, TX
- **DEC 7–9** 2009 — **Fall 2009 Cray XT5 Hex-core Workshop**, Oak Ridge, TN
- **DEC 8–10** 2009 — **Architectures and Technology for Extreme Scale Computing**, San Diego, CA
- **JAN 4–6** — **AACE (Architecture Aware Compiler Environment) PI Meeting**, Houston, TX
- **JAN 5–7** — **Sciences et Technologies de l'information et de la Communication (STIC) Review**, Paris, France
- **JAN 19–21** — **MPI Forum**, Atlanta, GA
- **JAN 20–21** — **Hybrid Multicore Consortium, First Annual Workshop**, San Francisco, CA
- **JAN 26** — **FAST-OS Meeting**, Oak Ridge, TN
- **JAN 27–29** — **SDCI/STCI as the Software Supply Chain of the National Cyberinfrastructure Workshop**, Arlington, VA
- **FEB 4** — **Scalable Tools Communication Infrastructure (STCI)**, Oak Ridge, TN
- **FEB 24–26** — **SIAM Conference on Parallel Processing and Scientific Computing (PP10)**, Seattle, WA
- **MAR 5** — **SBIR Reviews**, Washington, DC
- **MAR 8–10** — **MPI Forum**, San Jose, CA
- **MAR 8–11** — **LCI Conference on High Performance Cluster Computing**, Pittsburgh, PA
- **MAR 16** — **Blackjack Quarterly Progress Review**, Portland, OR
- **MAR 23–24** — **PERI All-Hands Meeting**, Knoxville, TN
- **MAR 26** — **ICL 20th Anniversary Workshop**, Knoxville, TN
- **APR 7** — **AACE/Blackjack Characterization Meeting**, Arlington, VA
- **APR 13–14** — **International Exascale Software Project (IESP) Workshop**, Oxford, UK

APR 19-23	**IPDPS 2010**	Atlanta, GA
APR 21	**Vampir Workshop 2010**	Bloomington (Indiana University), IN
APR 19-23	**HIPS 2010**	Atlanta, GA
APR 23-24	**CRA-W Grad Cohort Workshop**	Bellevue, WA
APR 24-25	**Chic Tech 2010 Retreat**	Urbana-Champaign, IL
MAY 3-5	**MPI Forum**	Chicago, IL
MAY 7	**Microwave Data Analysis for PetaScale Computers (MIDAS) Kickoff Meeting**	Paris, France
MAY 30 - JUN 3	**International Supercomputing Conference (ISC'10)**	Hamburg, Germany
JAN 2-4	**3rd 'Scheduling in Aussois' Workshop**	Aussois, French Alps
JUN 6-9	**Para 2010: State of the Art in Scientific and Parallel Computing**	Reykjavik, Iceland
JAN 12-13	**CUDA Center of Excellence 2010**	Beijing, China
JUN 14-16	**MPI Forum**	San Jose, CA
JUN 16	**MAGMA - a New Generation of Linear Algebra Libraries for GPU and Multicore Architectures**	ORNL
JAN 21	**HPC 2010 Advanced Workshop**	Cosenza, Italy
JUN 22-25	**VECPAR**	Berkeley, CA
JUN 29-30	**DARPA Blackjack Characterization Meeting - AACE PI Meeting**	Princeton, NJ
JUL 1	**SAFE-OS Project Review**	Paris, France
JUL 11-15	**SciDAC 2010**	Chattanooga, TN
JUL 12-15	**SAAHPC'10**	Knoxville, TN
AUG 2	**PetaApps Meeting**	Nashville, TN
AUG 2-5	**CScADS Workshop on Performance Tools for Petascale Computing**	Snowbird, UT
AUG 9-11	**CScADS Workshop on Automatic Tuning**	Snowbird, UT
AUG 31	**Euro-Par 2010**	Ischia - Naples, Italy
AUG 31 - SEP 3	**PETAL Workshop**	Orsay, France
SEP 1-2	**Performance Engineering Research Institute (PERI) All-Hands Meeting**	Salt Lake City, UT
SEP 7-10	**Clusters, Clouds, and Grids for Scientific Computing**	Flat Rock, NC
SEP 12-15	**EuroMPI 2010**	Stuttgart, Germany
SEP 13-16	**Parallel Software Tools and Tool Infrastructures (PSTI 2010)**	San Diego, CA
SEP 15-16	**HPEC 2010**	Lexington, MA
SEP 20-23	**GPU Technology Conference (GTC 2010)**	San Jose, CA
SEP 21-23	**NSF Workshop on Software Development Environments for Science & Engineering Applications**	Arlington, VA
SEP 28 - OCT 2	**2010 Grace Hopper Celebration of Women in Computing**	Atlanta, GA
OCT 17-19	**International Exascale Software Project (IESP) Workshop**	Maui, HI
OCT 18-19	**ALCF Early Science Workshop 2010**	Argonne, IL
OCT 21	**Final Review SAFE-OS**	Paris, France
NOV 13-19	**Supercomputing 2010**	New Orleans, LA

OTHER EFFORTS & ACTIVITIES

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 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 in which we are participating or have a taken a leadership role.

CITR

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.

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 the scientists and engineers will need to make them useful for extreme-scale research. But the challenge of creating an entirely new software stack for high performance computing is quite 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 and help orchestrate 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 year and a half of work, the IESP organized five workshops at different locations around the globe: Sante Fe, NM (USA); Paris, France; Tsukuba, Japan; Oxford, UK; and Maui, HI (USA). 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://exascale.org/

IGMCS

Interdisciplinary Graduate Minor in Computational Science

Addressing the need for any academic research strategy, the Center fo Information Technology Research 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 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 2010, 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/

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.

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,

"During the past year, the virtual institute continued its successful series of training activities aimed at teaching application developers the effective use of our programming tools. Tutorials and tuning workshops with our participation took place in Amsterdam, Bloomington, Bremen, Munich, and Portland – in some of these locations even multiple times. Soon we will have debuts in new countries including Greece and Saudi Arabia. With such a busy training schedule, we are very glad to have experienced instructors from the ICL on our team, many of whom tour the HPC landscape giving high-quality tutorials and have done so long before the virtual institute came into existence."

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

VI-HPS PARTNERS

Forschungszentrum Jülich
Jülich Supercomputing Centre

Technische Universität Dresden
Center for Information Services and High Performance Computing

RWTH Aachen University
Center for Computing and Communication

University of Tennessee
Innovative Computing Laboratory

ASSOCIATED PARTNERS

German Research School for Simulation Sciences
Laboratory of Parallel Programming

Technische Universität München
Chair for Computer Architecture

University of Stuttgart
High Performance Computing Centre

SPONSORS

INNOVATIVE COMPUTING LABORATORY

For more than 21 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

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

INNOVATIVE
COMPUTING LABORATORY
THE UNIVERSITY of TENNESSEE

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