ICL UT INNOVATIVE COMPUTING LABORATORY
2008/2009 REPORT

THE UNIVERSITY of TENNESSEE

INNOVATIVE COMPUTING LABORATORY **2008/2009** REPORT

EDITED BY Scott Wells DESIGNED BY David Rogers

INNOVATIVE COMPUTING LABORATORY
2008/2009 REPORT

CONTENTS

FROM THE DIRECTOR

JACK DONGARRA, DIRECTOR OF ICL

"ICL is prepared to address some of the most important computational scientific issues of our time."

The Innovative Computing Laboratory now has 19 years under its wings and has seen many changes here at UTK and ORNL. The past year has been particularly exciting for high performance computing with the NSF Track 2 award coming to UT/JICS and the DOE Leadership class system at ORNL becoming the largest open DOE computing facility. In addition, ICL has transitioned to the College of Engineering and is experiencing the merger of the Computer Science Department and the Electrical and Computer Engineering Department into the Electrical Engineering and Computer Science Department. With these changes our long term goal of leadership in enabling technologies for high performance computing is still of great interest and importance.

ICL is prepared to address some of the most important computational scientific issues of our time. Our plans for the future are founded on our accomplishments as well as our vision. That vision challenges us to be a world leader in enabling technologies and software for scientific computing. We have been and will continue to be providers of high performance tools to tackle science's most challenging problems and to play a major role in the development of standards for scientific computing in general.

Looking at the current state and the future of ICL, we continue to collaborate with leading researchers around the country and world to address the challenges of high performance computing in the manycore and petascale era. 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 stimulating 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 ICL one of the best centers for enabling technology in the world.

- Jack Dongarra

BACKGROUND

Dr. Jack Dongarra founded ICL in 1989. Coming from Argonne National Laboratory, Dr. Dongarra received a dual appointment as a Distinguished Professor in the Computer Science Department at the University of Tennessee and as a Distinguished Scientist at nearby Oak Ridge National Laboratory (ORNL). Since that date, ICL has grown from a single graduate assistant to a fully functional center, with a staff of more than 40 researchers, students, and administrators.

AYRES HALL
HOME OF ICL 1989-2000

Our mission at ICL is simple: being a world leader in enabling technologies and software for scientific computing. Our goals are to provide leading edge tools to tackle science's most challenging high performance computing problems and to play a major role in the development of standards for scientific computing in general.

As part of the newly formed department of Electrical Engineering and Computer Science (EECS) in College of Engineering at the University of Tennessee Knoxville, ICL continues to lead the way as one of the most respected academic research centers in the world. 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.

Over the past 19 years, 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

Our mission at ICL is simple: being a world leader in enabling technologies and software for scientific computing.

to make exciting contributions at other institutions and organizations. Below is a short list of experts who have passed through ICL on their way to distinguished careers at other organizations and academic institutions.

Zhaojun Bai	The University of California, Davis
Richard Barrett	Oak Ridge National Laboratory
Adam Beguelin	Formerly of AOL, now retired
Susan Blackford	Myricom
Henri Casanova	University of Hawaii, Manoa
Jaeyoung Choi	Soongsil University, Korea
Andy Cleary	Lawrence Livermore National Laboratory
Frederic Desprez	ENS-Lyon, France
Victor Eijkhout	University of Texas, Austin
Graham Fagg	Microsoft
Edgar Gabriel	University of Houston
Robert van de Geijn	University of Texas, Austin
Julien Langou	University of Colorado at Denver
Antoine Petitet	ESI Group, France
Roldan Pozo	NIST
Erich Strohmaier	Lawrence Berkeley National Laboratory
Francoise Tisseur	Manchester University, England
Bernard Tourancheau	University of Lyon, France
Sathish Vadhiyar	Indian Institute of Science (IISC), India
Clint Whaley	University of Texas, San Antonio
Felix Wolf	Juelich Supercomputing Centre

Dozens of tools and applications of high value to the HPC community have been produced at ICL since its inception. Many of these now form the basic fabric of HPC, scientific computing. Below is just a brief list of some of the technologies that we have produced or been a part of over the past 19 years:

Active Netlib	ATLAS
BLAS	FT-MPI
HARNESS	KOJAK
LAPACK	LAPACK for Clusters (LFC)
LINPACK Benchmark	MPI
NetBuild	Netlib
NetSolve	PAPI
PVM	RIB
ScaLAPACK	TOP500

Today, some of these successes continue along with current efforts such as Fault Tolerant Linear Algebra (FT-LA), Generic Code Optimization (GCO), HPC Challenge benchmark suite (HPCC), High Performance Linpack (HPL), PLASMA, GridSolve, Open MPI and PAPI. Many of our efforts have been recognized nationally and internationally, which includes many awards such as four R&D 100 awards; PVM in 1994, ATLAS and NetSolve in 1999, and PAPI in 2001.

RESEARCH

Remaining at the forefront of computational research requires an ability to adapt to the ever changing landscape in high performance computing. Ever expanding 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 19 years ago as in-depth investigation 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 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, as well as asset management, for high-end computers.

Demonstrating the range and diversity of our research, we will be engaged in more than 20 significant research projects during 2008-2009 across our main areas of focus: numerical linear algebra, high performance distributed computing, performance analysis and benchmarking, asset management. On the following pages, we provide brief summaries of some of our efforts in each of these areas. For more detailed information about each research project, visit our website - http://icl.cs.utk.edu.

NUMERICAL LINEAR ALGEBRA

LA for Hybrid Systems

If multicore is a disruptive technology, try to imagine hybrid multicore systems enhanced with accelerators! This is really happening as accelerators, and in particular Graphical Processing Units (GPUs), are steadily making their way into the HPC community. ICL is involved in the development and evaluation of innovative linear algebra (LA) algorithms for hybrid systems, focusing on multi/manycore architectures enhanced with Graphical Processing Units (GPUs). By considering hybrid architectures, we can design algorithms that fully exploit the power that each of the hybrid components offer, and as a result deliver performance that is unprecedented for either current multicores or GPUs taken separately. The targets of our efforts are both dense linear algebra (DLA) and sparse iterative solvers and eigensolvers, as they form the core of many scientific computing applications.

FT-LA

Addressing fault-tolerance for applications running on large processor counts requires multiple levels of effort. Besides the middleware efforts that enable applications to transparently survive process failures, there is a need for developing efficient recovery patterns for some specific kernel applications. The Fault Tolerant Linear Algebra (FT-LA) effort involves research to create new linear algebra algorithms that will efficiently handle several process failures. For iterative methods, our research has been focused on diskless checkpoint techniques. For dense methods, our research has been focused on ABFT techniques with a successful implementation of an extremely scalable, fault tolerant (dense) matrix-matrix multiply subroutine. Our future work in this area involves the development of a scalable fault-tolerant (dense) Cholesky, LU and QR factorizations in the context of PLASMA (see below).

LAPACK and ScaLAPACK

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 3.1.1 was released in February 2007. Recent work on LAPACK has revolved around variance of factorization and iterative refinement for symmetric positive definite systems. 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 LAPACK and ScaLAPACK under Windows, natively. LAPACK 3.2 is expected to be released in November 2008, which will include additional iterative refinement routines, new algorithm variants of factorization, and much more. ScaLAPACK 1.9.0 was released in late 2008, which includes the addition of the MRRR algorithm

PLASMA and Multi-core/Cell

As part of our multicore effort within the PLASMA (Parallel Linear Algebra for Scalable Multicore Architectures) project, we implemented a class of tiled algorithms for one-sided transformations in the context of multicore processors, including LU, Cholesky and QR factorizations, which outperform vendor implementations, in particular Intel's MKL. As of late 2008, we are extending it to two-sided transformations in order to solve un/symmetric eigenvalue problems (hessenberg reduction) and singular value problems (bidagonal reduction). These new algorithms utilize block data layout and process the input matrices by small, fixed-size tiles, which facilitates dynamic out-of-order scheduling of tasks and greatly reduces any idle time due to unnecessary stalls. By introducing inner-blocking of tiles, we have greatly minimized superfluous floating point operations that result from tiling.

Our work on the IBM Cell processor includes successful implementation of mixed-precision algorithms, which exploit the single precision speed of the chip, while still delivering full double-precision accuracy at the end. We have successfully utilized the technique for solution of linear systems of equations using LU and Cholesky factorizations. As part of this effort, we have also developed two types of single precision matrix multiplication (SGEMM) kernels for LU factorization and for Cholesky factorization, respectively, using Cell (SPE) assembly language.

As of late 2008, we are working on utilization of both Cell processors in the QS20 blade by efficiently exploiting the NUMA architecture of the blade using the libnuma library and ScaLAPACK-like block cyclic (1D) data partitioning. We have also successfully implemented a tile QR factorization with internal blocking, which shows performance on the Cell processor close to that of matrix multiplication.

Additional work involves exploring the capabilities of the Cell/SMP-SuperScalar environment to perform automatic thread-level (data dependent) parallelization of code and dynamic task scheduling. To date, the results for SMP systems are very encouraging. Also, the results for the Cell processor identify the scheduling mechanism as being a bottleneck when running on the PPE.

Nanotechnology

We are collaborating with several leading Physics groups on nanotechnology related problems. Nanotechnology, recognized by many as the driving force for the next "industrial revolution," is a multidisciplinary field for the study, manipulation and control of individual atoms and molecules on the nanometer scale. Numerical simulations, in order to correctly capture the physics on this scale, lead to very large computations, and thus easily require petascale (and beyond) computing resources and efficient software tools. We are involved in the development and tuning of new algorithms for the petascale regime, as well as profiling, performance modeling and automatic optimization of kernels. In particular, this includes iterative eigensolvers for certain interior eigenvalue problems, efficient preconditioning techniques for iterative eigensolvers of interest, dense LA kernels, as well as multisclale solvers, eigensolvers and preconditioners.

PERFORMANCE ANALYSIS AND BENCHMARKING

Generic Code Optimization (GCO)

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 to improve the process of finding the best code implementation.

HPCC

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 whose 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.

High Performance Linpack (HPL)

HPL is a software package that solves a (random) dense linear system in double precision (64 bits) arithmetic on distributed-memory computers. After nearly four years since the last stable release, HPL 2.0 was released in September 2008. Besides containing 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.

PERFORMANCE ANALYSIS AND BENCHMARKING CONTINUED

Performance API (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 performance 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 is available as a technology pre-release with example components provided for limited network counter and thermal monitoring support.

TOP500

For 15 years now, a ranking of the top 500 fastest computers in the world has been compiled biannually with published results occurring 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 ERSC/Lawrence Berkeley National Laboratory and the University of Mannheim, Germany to produce the rankings.

ASSET MANAGEMENT

RIB

The Repository in a Box (RIB) is an open-source software package for creating WWW metadata repositories. Metadata, from RIB's perspective, is information that describes reusable objects, such as software. RIB allows the user to enter metadata into a user friendly java applet which then sends the information to a RIB server via HTTP. The information is then stored in an SQL database where it is made available in a fully functional web site (catalog, search page, etc). Repositories which use similar data models can use the XML processing capabilities of RIB to share information via the Internet.

Netlib

Created in 1985, the Netlib repository contains freely available software, documents, and databases of interest to the numerical, computational science, scientific computing, and other communities. The repository is maintained by AT&T Bell Laboratories, the University of Tennessee and Oak Ridge National Laboratory, and by colleagues world-wide. The multi-collection is replicated at several sites around the world, automatically synchronized, to provide reliable and network efficient service to the global community. There have been more than half a billion accesses to Netlib as of 2008.

DISTRIBUTED AND GRID COMPUTING

FT-MPI and OpenMPI

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).

GridSolve

The purpose of GridSolve is to create the middleware necessary to provide a seamless bridge between the simple, standard programming interfaces and desktop Scientific Computing Environments (SCEs) that dominate the work of computational scientists and the rich supply of services supported by the emerging Grid architecture, so that the users of the former can easily access and reap the benefits (shared processing, storage, software, data resources, etc.) of using the latter. The GridSolve system is an RPC-based, client-agent-server system that includes service registration, service discovery, load balancing and service level fault tolerance. Recent work in GridSolve includes improved automatic workflow inference and execution, improved scheduling capabilities, improved client interfaces (Matlab, Octave, IDL, C and Fortran), MySQL compatibility, and various performance enhancements.

vGrADS

We are involved in the Virtual Grid Application Development Software (VGrADS) project, a multi-institution effort that addresses fundamental problems in effectively programming and using highly dynamic Grid systems. In order to take advantage of dynamic resources, an application describes its resource requirements (processors, connectivity, etc) as an abstract virtual grid. VGrADS can then schedule the application on the available resources. In order to accurately schedule the tasks in workflow applications, VGrADS implements a pseudo-reservation system by building on a statistical batch queue prediction technique. A fault-tolerant scheduler enables deadline-based scheduling in the face of dynamic and unreliable resources and services. In recent work, transparent, dynamic, cloud computing resource management is being added to VGrADS via the Eucalyptus project, which is an Amazon EC2 interface compatible, cloud computing infrastructure.

OTHER RESEARCH EFFORTS

DoD HPCMP

Since 1997, ICL has played a significant role in the Department of Defense's High Performance Computing Modernization Program (HPCMP). Specifically, ICL has been a lead center in the user Productivity Enhancement and Technology transfer (PET) program as part of the HPCMP. The PET program brings expertise from industry and academia to the DoD HPC user community, which is used to improve both performance and productivity within the DoD HPC centers. This improvement comes from efforts such as tool development, training, collaboration, and technology transfer. Technology transfer refers to efforts to ensure the latest technologies are available to the DoD HPC user community.

ICL currently leads the PET functional area known as Computational Environment (CE). Our CE team is responsible for ensuring there is a consistent, easy to use computational environment across DoD centers. This computational environment includes performance tools, debuggers, math libraries, data management tools, and Higher Level Languages (HLL). The CE team's duties revolve around enhancing and providing parallel programming tools and libraries as well as assisting DoD users with the use of these tools and libraries. On that front, the team spends much of its time working with users and user codes on activities including performance analysis, debugging, message passing issues, and parallel I/O. In addition, our CE team has provided training in areas such as MPI, parallel I/O, performance analysis, debugging, IPv6, compiler optimizations, and other topics.

DOE SciDAC

The Coordinated Infrastructure for Fault Tolerance Systems (CIFTS) is providing a coordinated infrastructure that will enable Fault Tolerance 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).

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, and University of Tennessee (ICL).

The Center for Scalable Application Development Software (CScADS) for Advanced Architectures was created 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 SDCI - 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 IGERT - SCALE-IT

In mid-2008, UT won two $3 million grants from NSF's Integrative Graduate Education and Research Traineeships (or IGERT) program. One of these grants, called Scalable Computing and Leading Edge Innovative Technologies, or SCALE-IT, was developed with the help of UT's Center for Information Technology Research (CITR - see page 26). As part of the SCALE-IT initiative, ICL will play a pivotal role by contributing to the scientific software toolkit that students of this program will use. Working with UT faculty and ORNL staff, SCALE-IT will unite students from several UT colleges to explore novel ways of applying high performance computing to solve many challenging biological problems. Every SCALE-IT student will participate in the recently formed computational science graduate program, Interdisciplinary Graduate Minor in Computational Science (IGMCS - see page 26), on the UT campus, and SCALE-IT and the IGMCS will work together to attract more faculty collaborators over the five year grant period.

ICL RESEARCH
WEB DIRECTORY

NUMERICAL LINEAR ALGEBRA

LAPACK	http://www.netlib.org/lapack/
ScaLAPACK	http://www.netlib.org/scalapack/
PLASMA	http://icl.cs.utk.edu/plasma/

PERFORMANCE ANALYSIS AND BENCHMARKING

HPCC	http://icl.cs.utk.edu/hpcc/
High Performance Linpack (HPL)	http://icl.cs.utk.edu/hpl/
Performance API (PAPI)	http://icl.cs.utk.edu/papi/
TOP500	http://www.top500.org/

ASSET MANAGEMENT

RIB	http://icl.cs.utk.edu/rib/
Netlib	http://www.netlib.org/

DISTRIBUTED AND GRID COMPUTING

FT-MPI	http://icl.cs.utk.edu/ft-mpi/
OpenMPI	http://www.open-mpi.org/
GridSolve	http://icl.cs.utk.edu/gridsolve/
vGrADS	http://vgrads.rice.edu/

OTHER RESEARCH EFFORTS

DoD HPCMP	http://www.hpcmo.hpc.mil/
DOE SciDAC	http://www.scidac.gov/
NSF SDCI - POINT	http://www.nic.uoregon.edu/point/

More information on ICL Research efforts can be found at http://icl.cs.utk.edu/

PEOPLE

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.

ICL GROUP, FALL 2008

ICL STAFF AND STUDENTS

- **Wes Alvaro** — Graduate Research Assistant
- **Wesley Bland** — Graduate Research Assistant
- **Kartheek Bodanki** — Graduate Research Assistant
- **George Bosilca** — Research Asst. Professor
- **Aurelien Bouteiller** — Senior Research Associate
- **Barry Britt** — IT Administrator I
- **Tom Cortese** — Research Scientist - CE Lead
- **David Cronk** — Research Director
- **Yuanshun (Shaun) Dai** — Assistant Professor
- **Jack Dongarra** — University Distinguished Professor
- **Peng Du** — Graduate Research Assistant
- **Don Fike** — Research Assistant
- **Teresa Finchum** — Administrative Services Assistant
- **Bilel Hadri** — Senior Research Associate
- **Josh Hoffman** — Undergrad Student Asstistant
- **Heike Jagode** — Senior Research Associate
- **Jakub Kurzak** — Senior Research Associate
- **Julie Langou** — Research Associate
- **Tracy Lee** — Accounting Specialist III
- **Pierre Lemarinier** — Senior Research Associate
- **Yinan Li** — Graduate Research Assistant
- **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
- **Keith Seymour** — Senior Research Associate
- **Matt Skinner** — Graduate Research Assistant
- **Gwang Son** — Graduate Research Assistant
- **Fengguang Song** — Graduate Research Assistant
- **Dan Terpstra** — Research Leader
- **Joe Thomas** — IT Analyst III
- **Stanimire Tomov** — Research Scientist
- **Scott Wells** — Assistant Director
- **Asim YarKhan** — Senior Research Associate
- **Haihang You** — Research Associate

VISITORS

Twenty-first century research agenda's 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 are the enormous opportunities to invite and work with researchers from all around the world. Since our group 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 extremely beneficial to us and we will continue providing opportunities for visits from our national and international colleagues in research.

ICL ALUMNI

Since its formation, ICL has employed many students and staff from multiple background and academic disciplines. From undergraduate students to experienced research scientists, more than 170 individuals have passed through our doors, many of whom have gone on to both achieve and contribute great things in academia, government, and the private sector. Many of our former staff and students are now faculty members at some of the most prominent 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.com, Cray, Google, Hewlett Packard, IBM, Intel, Microsoft, Myricom, NEC, and SiCortex just to name a few. And many have even gone on to form companies of their own.

1990-1993 **Carolyn Aebischer**	2001-2006 **Sudesh Agrawal**	2004-2006 **Bivek Agrawal**	1989-1991 **Ed Anderson**	2007 **Daniel Andrzejewski**
2003-2007 **Thara Angskun**	2003 **Papa Arkhurst**	1999-2001 **Dorian Arnold**	2008 **Marc Baboulin**	1990-1992 **Zhaojun Bai**
2001-2002 **Ashwin Balakrishnan**	1992-1994 **Richard Barrett**	2000-2001 **Alex Bassi**	2000-2001 **Micah Beck**	2007 **Daniel Becker**
1991 **Adam Beguelin**	1991 **Annamaria Benzoni**	1997-1998 **Scott Betts**	2003-2005 **Nikhil Bhatia**	2002-2003 **Noel Black**
1989-2001 **Susan Blackford**	1999-2000 **Fernando Bond**	1997-1999 **Randy Brown**	1998-1999 **Murray Browne**	2005 **Cynthia Browne**
1998-2003 **Antonin Bukovsky**	1995 **Greg Bunch**	2006-2007 **Alfredo Buttari**	1995-1998 **Henri Casanova**	2005 **Ramkrishna Chakrabarty**
1998-2000 **Sharon Chambers**	2001-2006 **Zizhong Chen**	1994-1995 **Jaeyoung Choi**	1998 **Eric Clarkson**	1995-1997 **Andy Cleary**
2007 **Camille Coti**	1993-1997 **Jason Cox**	1998-1999 **Cricket Deane**	1994-1995 **Frederic Desprez**	2001-2003 **Jun Ding**
2003 **Jin Ding**	1993-1994 **Martin Do**	2000-2001 **Leon Dong**	1997 **David Doolin**	1998-2003 **Andrew Downey**
1989-1992 **Mary Drake**	2002-2004 **Julio Driggs**	2001-2004 **Brian Drum**	1992-2005 **Victor Eijkhout**	1995-2005 **Brett Ellis**
2004 **Shawn Ericson**	1997-1998 **Zachary Eyler-Walker**	2003-2004 **Lisa Ezzell**	1996-2006 **Graham Fagg**	1997-1998 **Markus Fischer**

RECENT VISITORS TO ICL

Phil Andrews National Institute for Computational Sciences	Marc Baboulin Universidade de Coimbra, Portugal	Petter Bjorstad University of Bergen, Norway
Elisabeth Brunet Université Bordeaux I, France	Alfredo Buttari ENS Lyon, France	Franck Cappello INRIA, France
Anthony Danalis University of Delaware	Karl Fuerlinger University of California at Berkeley	Christoph Geile Juelich Supercomputing Centre, Germany
Sven Hammarling Numerical Algorithms Group (NAG), UK	Thomas Herault University of Paris-Sud, France	Emmanuel Jeannot INRIA, France
Bo Kagstrom University of Umea, Sweden	Julien Langou University of Colorado at Denver	Jeff Larkin Cray, Inc.
James McCombs College of William and Mary	Cleve Moler The MathWorks	Sushil Prasad Georgia State University
Narapat Saengpatsa Suranaree University of Technology, Thailand	Martin Swany University of Delaware	Bernard Tourancheua University of Lyon, France
Felix Wolf Juelich Supercomputing Centre, Germany	Ailing Zhao University of Illinois, Chicago	

2003-2004 **Xiaoquan Fu**	2003-2007 **Erika Fuentes**	2006-2008 **Karl Fuerlinger**	2006-2008 **Megan Fuller**	2003-2004 **Edgar Gabriel**
2000-2001 **Lynn Gangwer**	1992-1993 **Tracy Gangwer**	2001-2006 **Nathan Garner**	1998 **Kelley Garner**	1996 **Jonathan Gettler**
1993 **Eric Greaser**	1992-1996 **Stan Green**	2004-2006 **Alice Gregory**	2000-2001 **Hunter Hagewood**	1996-1997 **Christian Halloy**
1996-1997 **Sven Hammarling**	1995-1996 **Hidehiko Hasegawa**	1995-1996 **Satomi Hasegawa**	1996 **Chris Hastings**	1999-2001 **David Henderson**
1996 **Greg Henry**	1996-1998 **Sid Hill**	1998-2000 **George Ho**	1995-1999 **Jeff Horner**	2000-2001 **Yan Huang**
2002 **Chris Hurt**	1992-1995 **Paul Jacobs**	1999-2000 **Weizhong Ji**	1992-1995 **Weicheng Jiang**	1997-1998 **Song Jin**
1992-2008 **Jan Jones**	2007 **Venkata Kakani**	2001 **Balajee Kannan**	2008 **Madhuri Kasam**	2002 **David Katz**
2008 **Supriya Kilambi**	1992-1996 **Youngbae Kim**	2008 **Jenya Kirshtein**	1993-1996 **Michael Kolatis**	2005 **Coire Kyle**
2003-2004 **Amanda Laake**	2003-2005 **Jeff Larkin**	2000-2002 **DongWoo Lee**	1993-1994 **Todd Letsche**	1992-1995 **Sharon Lewis**
2001 **Xiang Li**	2002 **Weiran Li**	2000 **Chaoyang Liu**	1996-2005 **Kevin London**	1999 **Matt Longley**
2005-2008 **Daniel Lucio**	2000-2001 **Richard Luczak**	1990-1996 **Robert Manchek**	2004 **Tushti Marwah**	1994-2000 **Paul McMahan**
2003-2006 **Eric Meek**	1998-2002 **Jeremy Millar**	1999-2003 **Michelle Miller**	2001-2002 **Cindy Mitchell**	1987-2007 **Keith Moore**
1991-1993 **Steven Moulton**	2008 **Matthew Nabity**	2004-2005 **Shankar Narasimhaswami**	1994-1995 **Peter Newton**	1997-1998 **Caroline Papadopoulos**
2002 **Leelinda Parker**	2008 **Dilip Patlolla**	1993-2001 **Antoine Petitet**	2008 **Vlado Pjesivac**	2003-2007 **Jelena Pjesivac-Grbovic**
1991-1992 **James S. Plank**	1992-1994 **Roldan Pozo**	2005-2006 **Farzona Pulatova**	1999-2001 **Tammy Race**	1998-2000 **Ganapathy Raman**
2003 **Kamesh Ramani**	1999-2004 **Mei Ran**	1996-1997 **Yves Robert**	1999-2004 **Ken Roche**	1997-1998 **Tom Rothrock**
1993-1997 **Tom Rowan**	2008 **Narapat (Ice) Saengpatsa**	2001-2005 **Kiran Sagi**	1998-1999 **Evelyn Sams**	2001 **Farial Shahnaz**
2001-2007 **Zhiao Shi**	2005-2007 **Sergei Shinkarev**	1991-1992 **Majed Sidani**	1996-1998 **Shilpa Singhal**	1999-2001 **Thomas Spencer**
1995-2001 **Erich Strohmaier**	1996-1999 **Martin Swany**	2002 **Daisuke Takahashi**	1993-1999 **Judi Talley**	2005-2006 **Yuan Tang**
1998 **Keita Teranishi**	1998-1999 **John Thurman**	1997 **Francoise Tisseur**	1993-1994 **Bernard Tourancheau**	2004 **Lauren Vaca**
1999-2003 **Sathish Vadhiyar**	1990-1991 **Robert van de Geijn**	2008 **Eugene Vecharynski**	1993-1995 **Scott Venckus**	1990-1996 **Reed Wade**
2001-2005 **Michael Walters**	1991-2001 **R. Clint Whaley**	2000-2001 **Susan Wo**	2003-2005 **Felix Wolf**	2004-2007 **Jiayi Wu**
2004-2005 **Qiu Xia**	1998-2000 **Tinghua Xu**	1999 **Tao Yang**	2001-2005 **Yuanlei Zhang**	2001 **Yong Zheng**
2000-2001 **Luke Zhou**	2002-2004 **Min Zhou**			

PARTNERSHIPS

Since 1989, ICL has continuously 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.

DOMESTIC COLLABORATORS

The Aerospace Corporation		ANL	Argonne National Laboratory	BLAST	Basic Linear Algebra Subprograms Technical Forum
CACR	Caltech Center for Advanced Computing Research	DARPA	Defense Advanced Research Projects Agency	DoD	U.S. Department of Defense
DoD HPMCP	High Performance Computing Modernization Program	DOE	U.S. Department of Energy	Emory University	
Google		HiPerSoft	Center for High Performance Software Research	IBM	
Intel Corporation		Internet2		ISI	Information Sciences Institute
I2-DSI	The Internet2 Distributed Storage Infrastructure	JICS	Joint Institute for Computational Science	LANL	Los Alamos National Laboratory
LBL	Lawrence Berkely National Laboratory	LLNL	Lawrence Livermore National Laboratory	MathWorks	
Microsoft Research		MRA	MetaCenter Regional Alliance	NASA	National Aeronautics and Space Administration
NCSA	The National Computational Science Alliance	NHSE	The National HPCC Software Exchange	NICS	National Institute of Standards and Technology
NIST	National Institute of Standards and Technology	NPACI	National Partnership for Advanced Computational Infrastructure	NSF	National Science Foundation
ORNL CSMD	ORNL Computer Science and Mathematics Division	Rice University		SDSC	San Diego Supercomputing Center
SGI		Sun Microsystems		TACC	Texas Advanced Computing Center
The University of California, Berkeley		The University of California, San Diego		The University of California, Santa Barbara	
The University of Illinois, Urbana-Champaign		The University of North Carolina			

INTERNATIONAL COLLABORATORS

Danish Computing Center for Research and Education
Lyngby, Denmark

Department of Mathematical and Computing Sciences, Tokyo Institute of Technology Tokyo, Japan

University of Manchester
Manchester, England

European Centre for Research and Advanced Training in Scientific Computing
Toulouse, France

University College Dublin
Dublin, Ireland

École Polytechnique Fédérale de Lausanne
Lausanne, Switzerland

Institut für Mathematik der Universtität Mannheim
Mannheim, Germany

Forschungszentrum Jülich Central Institute for Applied Mathematics
Jülich, Germany

High Performance Computing Center (HLRS) Stuttgart
Stuttgart, Germany

Institut ETH Zentrum
Zurich, Switzerland

Istituto per le Applicazioni del Calcolo "Mauro Picone"del C.N.R.
Rome, Italy

Intelligent Systems Design Laboratory, Doshisha University
Kyoto, Japan

Kasetsart University
Bangkok, Thailand

Laboratoire de l'Informatique du Parallélisme, École Normale Supérieure de Lyon Lyon, France

Laboratoire de Recherche en Informatique et en Automatique, Université de Paris-Sud Cedex, France

Mathematical Insitute, Utrecht University
Netherlands

Numerical Algorithms Group Ltd
Oxford, England

Parallel and HPC Application Software Exchange (PHASE)
Tsukuba, Japan

Laboratoire Réseaux Haut Débits et Support d'Applicatoins Multimedia (RESAM) Lyon, France

Rutherford Appleton Laboratory
Oxford, England

Soongsil University
Seoul, South Korea

Technische Universität Wien
Vienna, Austria

Umeå University
Umeå, Sweden

Virtual Institute High-Productivity Supercomputing (VI-HPS)
Jülich, Germany

HARDWARE RESOURCES

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 presents 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 Silver-storm 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

ICL LOCAL RESOURCES

JUGENE SUPERCOMPUTER
AT JÜLICH SUPERCOMPUTING CENTRE

DAVINCI, IBM SP6
AT NAVAL OCEANOGRAPHIC OFFICE MAJOR SHARED RESOURCE CENTER AT STENNIS SPACE CENTER, MISS.

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 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, exclusive 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 the Cray XT4 system (Kraken), a UT resource that is being upgraded to a 10,000+ compute socket XT5, as well as an existing Cray XT4 also being upgraded to an XT5 (Jaguar). 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, and the Cell
Several large (512+ proc) Linux Clusters
SGI Altix
SiCortex

JAGUAR, CRAY XT5
AT OAK RIDGE NATIONAL LABORATORY

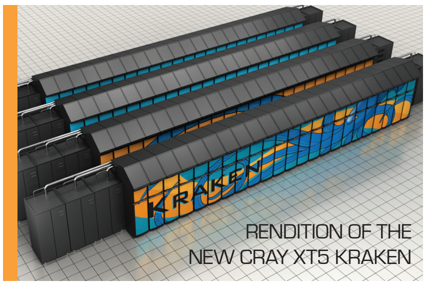

RENDITION OF THE
NEW CRAY XT5 KRAKEN

RECENT PUBLICATIONS

7th International parallel Processing and Applied Mathematics Conference, Lecture Notes in Comptuer Science, Wyrzykowski, R., Dongarra, J., Karczewski, K., Wasniewski, J. eds. Gdansk, Poland, Springer Berlin, Vol. 4967, September, 2008.

8th International Conference on Computational Science (ICCS), Proceedings Parts I, II, and III, Lecture Notes in Computer Science, Bubak, M., Albada, G., Dongarra, J., Sloot, P. eds. Krakow, Poland, Springer Berlin, Vol. 5101, June 23-25, 2008.

15th European PVM/MPI Users' Group Meeting, Recent Advances in Parallel Virtual Machine and Message Passing Interface, Lecture Notes in Computer Science, Lastovetsky, A., Kechadi, T., Dongarra, J. eds. Dublin Ireland, Springer Berlin, Vol. 5205, 2008.

Alvaro, W., Kurzak, J., Dongarra, J. "**Fast and Small Short Vector SIMD Matrix Multiplication Kernels for the CELL Processor**," University of Tennessee Computer Science Technical Report, UT-CS-08-609, (also LAPACK Working Note 189), January, 2008.

Baboulin, M., Dongarra, J., Gratton, S., Langou, J. "**Computing the Conditioning of the Components of a Linear Least Squares Solution**," VECPAR '08, High Performance Computing for Computational Science, Toulouse, France, June 24-27, 2008.

Baboulin, M., Dongarra, J., Tomov, S. "**Some Issues in Dense Linear Algebra for Multicore and Special Purpose Architectures**," University of Tennessee Computer Science Technical Report, UT-CS-08-615 (also LAPACK Working Note 200), May 6, 2008.

Baboulin, M., Gratton, S., "**Using dual techniques to derive componentwise and mixed condition numbers for a linear functional of a linear least squares solution**," University of Tennessee Computer Science Technical Report, UT-CS-08-622 (also LAPACK Working Note 207), August 4, 2008.

Bosilca, G., Delmas, R., Dongarra, J., Langou, J. "**Algorithmic Based Fault Tolerance Applied to High Performance Computing**," University of Tennessee Computer Science Technical Report, UT-CS-08-620 (also LAPACK Working Note 205), June 19, 2008.

Bouteiller, A., Bosilca, G., Dongarra, J. "**Redesigning the Message Logging Model for High Performance**," International Supercomputer Conference (ISC 2008), Dresden, Germany, June 17, 2008.

Bouteiller, A., Desprez, F. "**Fault Tolerance Management for a Hierarchical GridRPC Middleware**," 8th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2008), Lyon, France, May 19-22, 2008.

Buttari, A., Dongarra, J., Kurzak, J., Langou, J., Langou, J., Luszczek, P., Tomov, S. "**Exploiting Mixed Precision Floating Point Hardware in Scientific Computations**," in High Performance Computing and Grids in Action, Grandinetti, L. eds. Amsterdam, IOS Press, 2008.

Buttari, A., Dongarra, J., Kurzak, J., Luszczek, P., Tomov, S. "**Using Mixed Precision for Sparse Matrix Computations to Enhance the Performance while Achieving 64-bit Accuracy**," ACM Transactions on Mathematical Software, Vol 34, Number 4, 2008.

Buttari, A., Langou, J., Kurzak, J., Dongarra, J. "**Parallel Tiled QR Factorization for Multicore Architectures**," Concurrency and Computation: Practice and Experience, Vol. 20, pp. 1573-1590, 2008.

Caniou, Y., Caron, E., Desprez, F., Nakada, H., Tanaka, Y., Seymour, K. "**High Performance GridRPC Middleware**," Recent developments in Grid Technology and Applications, Gravvanis, G.A., Morrison, J.P., Arabnia, H.R. and Power, D.A. eds. Nova Science Publishers, 2008.

Canning, A., Dongarra, J., Langou, J., Marques, O., Tomov, S., Voemel, C., Wang, L.-W. "**Interior State Computation of Nano Structures**," PARA 2008, 9th International Workshop on State-of-the-Art in Scientific and Parallel Computing, Trondheim, Norway, May 13-16, 2008.
Chen, Z., Dongarra, J. "**Algorithm-Based Fault Tolerance for Fail-Stop Failures**," IEEE Transactions on Parallel and Distributed Systems, Volume. 19, No. 12, December, 2008.

Dongarra, J. "**A Tribute to Gene Golub**," Computing in Science and Engineering, IEEE, pp. 5, March/April, 2008.

Dongarra, J. "**Performance of Various Computers Using Standard Linear Equations Software (Linpack Benchmark Report)**," University of Tennessee Computer Science Technical Report, CS-89-85, September 30, 2008.

Dongarra, J., Demmel, J., Husbands, P., Luszczek, P. "**HPCS Library Study Effort**," University of Tennessee Computer Science Technical Report, UT-CS-08-617, 2008.

Dongarra, J., Golub, G. H., Grosse, E., Moler, C., Moore, K. "**Netlib and NA-Net: Building a Scientific Computing Community**," IEEE Annals of the History of Computing, Vol. 30, number 2, pp. 30-41, April - June, 2008.

Dongarra, J., Graybill, R., Harrod, W., Lucas, R., Lusk, E., Luszczek, P., McMahon, J., Snavely, A., Vetter, J., Yelick, K., Alam, S., Campbell, R., Carringon, L., Chen, T.-Y., Khalili, O., Meredith, J., Tikir, M. "**DARPA's HPCS Program: History, Models, Tools, Languages**," in Advances in Computers, Zelkowitz, M. eds. Elsevier, Vol. 72, 2008.

Dongarra, J., Langou, J. "**The Problem with the Linpack Benchmark Matrix Generator**," University of Tennessee Computer Science Technical Report, UT-CS-08-621 (also LAPACK Working Note 206), June 28, 2008.

Dongarra, J., Luszczek, P. "**How Elegant Code Evolves With Hardware: The Case Of Gaussian Elimination**," in Beautiful Code Leading Programmers Explain How They Think (Chapter 14), Oram, A., Wilson, G. eds. Japanese Edition, pp. 243-282, 2008.

Dongarra, J., Moore, S., Peterson, G., Tomov, S., Allred, J., Natoli, V., Richie, D. "**Exploring New Architectures in Accelerating CFD for Air Force Applications**," DoD HPCMP User Group Conference (submitted), Seattle, Washington, July 14-17, 2008.

Dongarra, J., Pineau, J.-F., Robert, Y., Vivien, F. "**Matrix Product on Heterogeneous Master Worker Platforms**," 2008 PPoPP Conference, Salt Lake City, Utah, February 20-23, 2008.

A complete bibliography of our publications and technical reports can be found at http://icl.cs.utk.edu/

Fuerlinger, K., Moore, S. "**Detection and Analysis of Iterative Behavior in Parallel Applications**," Proceedings of the 2008 International Conference on Computational Science (ICCS 2008), Krakow, Poland, Lecture Notes in Computer Science 5103, pp. 261-267, June, 2008.

Gustavson, F. G., Wasniewski, J., Dongarra, J. "**Rectangular Full Packed Format for Cholesky's Algorithm: Factorization, Solution and Inversion**," University of Tennessee Computer Science Technical Report, UT-CS-08-614 (also LAPACK Working Note 199), April 28, 2008.

Hardt, M., Seymour, K., Dongarra, J., Zapf, M., Ruiter, N.V. "**Interactive Grid-Access Using Gridsolve and Giggle**," Computing and Informatics, Vol. 27, No. 2, 233-248, ISSN 1335-9150, 2008.

Hernandez, O., Song, F., Chapman, B., Dongarra, J., Mohr, M., Moore, S., Wolf, F. "**Performance Instrumentation and Compiler Optimizations for MPI/OpenMP Applications**," Lecture Notes in Computer Science, OpenMP Shared Memory Parallel Programming, Springer Berlin / Heidelberg, Vol. 4315, 2008.

Jagode, H., Hein, J. "**Custom assignment of MPI ranks for parallel multi-dimensional FFTs: Evaluation of BG/P versus BG/L** (submitted)," The 2008 IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA-08), Sydney, Australia, Springer, InderScience, December 10-12, 2008.

Jagode, H., Hein, J., Trew, A. "**Task placement of parallel multi-dimensional FFTs on a mesh communication network**," University of Tennessee Computer Science Technical Report, UT-CS-08-613, April 22, 2008.

Kurzak, J., Buttari, A., Dongarra, J. "**Solving Systems of Linear Equations on the CELL Processor Using Cholesky Factorization**," IEEE Transactions on Parallel and Distributed Systems, Vol. 19, Number 9, pp. 1-11, September, 2008.

Kurzak, J., Buttari, A., Luszczek, P., Dongarra, J. "**The PlayStation 3 for High Performance Scientific Computing**," University of Tennessee Computer Science Technical Report, UT-CS-08-608, January 3, 2008.

Kurzak, J., Buttari, A., Luszczek, P., Dongarra, J. "**The PlayStation 3 for High Performance Scientific Computing**," Computing in Science and Engineering, pp. 80-83, May/June, 2008.

Kurzak, J., Dongarra, J. "**QR Factorization for the CELL Processor**," University of Tennessee Computer Science Technical Report, UT-CS-08-616 (also LAPACK Working Note 201), May, 2008.

Li, Y., Dongarra, J. "**Request Sequencing: Enabling Workflow for Efficient Parallel Problem Solving in GridSolve**," ICL Technical Report, ICL-UT-08-01, April, 2008.

Li, Y., Dongarra, J., Seymour, K., YarKhan, A. "**Request Sequencing: Enabling Workflow for Efficient Problem Solving in GridSolve**," International Conference on Grid and Cooperative Computing (GCC 2008) (submitted), Shenzhen, China, October 26-28, 2008.

Ltaief, H., Kurzak, J., Dongarra, J. "**Parallel Band Two-Sided Matrix Bidiagonalization for Multicore Architectures**," University of Tennessee Computer Science Technical Report, UT-CS-08-631 (also LAPACK Working Note 209), October 1, 2008.

Ltaief, H., Kurzak, J., Dongarra., J. "**Parallel Block Hessenberg Reduction using Algorithms-By-Tiles for Multicore Architectures Revisited**," University of Tennessee Computer Science Technical Report, UT-CS-08-624 (also LAPACK Working Note 208), August 7, 2008.

Seymour, K., You, H., Dongarra, J. "**A Comparison of Search Heuristics for Empirical Code Optimization**," The 3rd international Workshop on Automatic Performance Tuning, Tsukuba, Japan, October 1st, 2008.

Song, F., Moore, S., Dongarra, J., "**Analytical Modeling for Affinity-Based Thread Scheduling on Multicore Platforms**," University of Tennessee Computer Science Technical Report, UT-CS-08-626, August 12, 2008.

Tomov, S., Dongarra, J., Baboulin, M. "**Towards Dense Linear Algebra for Hybrid GPU Accelerated Manycore Systems**," University of Tennessee Computer Science Technical Report, UT-CS-08-632 (also LAPACK Working Note 210), October 17, 2008.

Vomel, C., Tomov, S., Marques, O., Canning, A., Wang, L.-W., Dongarra, J. "**State-of-the-Art Eigensolvers for Electronic Structure Calculations of Large Scale Nano-Systems**," Journal of Computational Physics, Vol. 227, Issue15, pp. 7113-7124, July, 2008.

Wolf, F., Wylie, B., Abraham, E. Becker, D., Frings, W., Fuerlinger, K., Geimer, M. Hermanns, M., Mohr, B., Moore, S., Pfeifer, M. Szebenyi, Z. "**Usage of the Scalasca Toolset for Scalable Performance Analysis of Large-scale Parallel Applications**," Proceedings of the 2nd International Workshop on Tools for High Performance Computing, Resch, M., Keller, R., Himmler, V., Krammer, B., Schulz, A. eds. Stuttgart, Germany, Springer, pp. 157-167, July, 2008.

Youseff, L., Seymour, K., You, H., Dongarra, J., Wolski, R. "**The Impact of Paravirtualized Memory Hierarchy on Linear Algebra Computational Kernels and Software**," ACM/IEEE International Symposium on High Performance Distributed Computing, Boston, MA., June 23-27, 2008.

Japanese Translation: **How Elegant Code Evolves With Hardware: The Case Of Gaussian Elimination**, Jack Dongarra and Piotr Luszczek, ビューティフルコード, Andy Oram編 Greg Wilson,編 Brian Kernighan（他）著 オライリー・ジャパン 版, 2008年04月 発行 ISBN 978-4-87311-363-0.

Chinese Translation:《并 行计算综 论》,（美）多加拉等编著, (Sourcebook of Parallel Computing, by J. Dongarra, I. Foster, G. Fox, W. Gropp, K. Kennedy, L. Torczon, and A. White), translated by 莫则尧等译, in 国外计算机科学教材系列, published by 电子工业出版社 (Publishing House of Electronics Industry), 出版日期：2005-05.

Cyberinfrastructure Technology Watch

For the past four years, ICL has collaborated with the NSF funded Cyberinfrastructure Partnership (CIP) to lead a broad ranging publication effort called Cyberinfrastructure Technology Watch (CTWatch). Although 2008 saw the final issue of its main publication, *CTWatch Quarterly*, CTWatch served both as a forum for ideas and opinions on issues of importance to the cyberinfrastructure community and as an ongoing source of information and analysis concerning the latest innovations in cyberinfrastructure technology.

To create the kind of productive mix of news, information, and dialogue that rapid progress in shared cyberinfrastructure requires, CTWatch was developed along two complementary paths, one based on a more traditional publishing paradigm and the other including new types of non-traditional, Internet-based communication and publishing. On the conventional front, *CTWatch Quarterly* offered an on-line serial publication modeled on a more traditional academic journal. Along a more experimental line, CTWatch Blog provided commentary and informative links on the most recent developments and ideas occurring in the national and international cyberinfrastructure community.

CTWatch Quarterly was published on-line and all of its issues will continue to be available indefinitely in both HTML and in high quality PDF format. Each issue was organized by a guest editor and revolved around an area of interest for the cyberinfrastructure community in which the guest editor had a leadership role. The focus topics (and corresponding guest editors) for 2007 included "The Promise and Perils of the Coming Multicore Revolution and Its Impact" (Jack Dongarra), "Socializing Cyberinfrastructure: Networking the Humanities, Arts, and Social Sciences" (David Theo Goldberg and Kevin D. Franklin), "The Coming Revolution in Scholarly Communications & Cyberinfrastructure" (Lee Dirks and Tony Hey) and "Software Enabling Technologies for Petascale Science" (Fred Johnson). The final issue – Spring 2008 – was guest edited by Pete Beckman and focused on "Urgent Computing," the application of high-performance and grid computing to emergency situation, such as epidemics and extreme weather.

Looking back over the course of its run, it is clear that CTWatch offered a remarkable record of the most advanced thinking about cyberinfrastructure during this period. An incredibly broad range of salient cyberinfrastructure topics have been highlighted via the Quarterly: from low-power, high performance computing to cyberinfrastructure in the social sciences, coverage has included both domestic and international efforts. The future of cyberinfrastructure could not be more exciting, and we are grateful for having been given the opportunity to lead the CTWatch effort. Many thanks go out to NSF, our CIP collaborators at San Diego Supercomputing Center (SDSC) and the National Center for Supercomputing Applications (NCSA), the guest editors, the article authors, the editorial board, and the entire production staff for successfully meeting the original goals of the Quarterly. Together this group played a truly significant role in increasing cyberinfrastructure visibility within academia, government, and industry.

http://www.ctwatch.org/

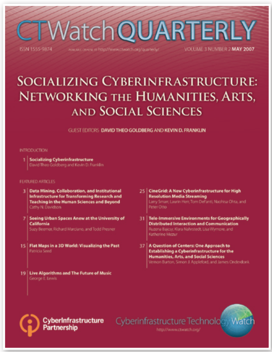

14 ISSUES OF **CTWATCH QUARTERLY** WERE PUBLISHED FROM FEBRUARY 2005 TO MARCH 2008

www.ctwatch.org/quarterly/

CITR

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

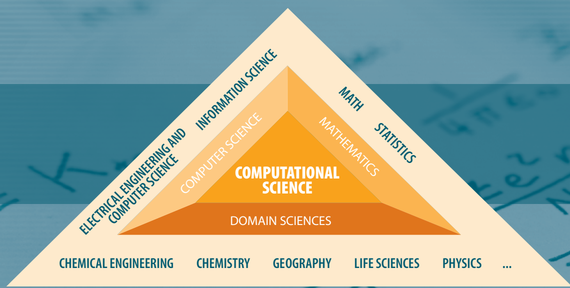

COMPUTATIONAL SCIENCE: INHERENTLY INTERDISCIPLINARY

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 2008, there were 15 different departments contributing nearly 100 courses to the program.

http://igmcs.cs.utk.edu/

TECHNISCHE UNIVERSITÄT DRESDEN, CENTER FOR INFORMATION SERVICES AND HIGH PERFORMANCE COMPUTING

FORSCHUNGSZENTRUM JÜLICH, JÜLICH SUPERCOMPUTING CENTRE

UNIVERSITY OF TENNESSEE INNOVATIVE COMPUTING LABORATORY

RWTH AACHEN UNIVERSITY, CENTER FOR COMPUTING AND COMMUNICATION

UNIVERSITY OF STUTTGART, HIGH PERFORMANCE COMPUTING CENTRE

VIRTUAL INSTITUTE HIGH PRODUCTIVITY SUPERCOMPUTING

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

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 three primary institutions in Germany plus ICL, joins 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.”

SPONSORS

Over the past 19 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 Institute 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:

ICL UT INNOVATIVE COMPUTING LABORATORY

ADDRESS 1122 Volunteer Blvd Suite 413
University of Tennessee
Knoxville, TN 37996-3450

WEBSITE http://icl.cs.utk.edu/

EMAIL icl@cs.utk.edu

PHONE (865) 974-8295

FAX (865) 974-8296

ICL UT