2006-2007 REPORT

ICL UT

INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY of TENNESSEE | UT

Computer Science Department

INNOVATIVE COMPUTING LABORATORY 2006-2007 REPORT

EDITED BY SCOTT WELLS DESIGNED BY DAVID ROGERS

2006-2007 REPORT

CONTENTS

ICL INFOBYTES

ESTABLISHED	1989
STAFF	28
STUDENTS	11
ACTIVE PROJECTS	17
4	R&D 100 AWARDS
21	INTERNATIONAL COLLABORATORS
27	DOMESTIC RESEARCH PARTNERS
AVG. ANNUAL PAPERS/PUBLICATIONS	33
AVG. ANNUAL EVENTS ATTENDED (CONFERENCES, SEMINARS, WORKSHOPS)	28

FROM THE DIRECTOR

JACK DONGARRA
DIRECTOR, ICL

In 2006, the Innovative Computing Laboratory is celebrating 17 years of leadership in enabling technologies for high performance computing. Looking back over the 17-year period, the evolution and growth of the technology for computing has been truly astonishing. In an environment where technology changes every 18 months or less, ICL cannot afford to stand still. In 1989 the speed of a supercomputer was measured in gigaflops and in gigabytes. Today our measures are 100's of teraflops for speed and 100's of terabytes for memory, and petaflop computing is becoming a reality. The research that ICL has undertaken in the past 17 years has followed a natural progression and growth from our original tread of numerical linear algebra to performance evaluation and analysis, to software asset management, and to distributed computing.

Today the Innovative Computing Laboratory is addressing 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. One of the central challenges facing us today will be developing a methodology for programming multicore systems and investigating how the current and future generations of multicore processing chips can be used to provide more performance and processing power to important areas of scientific computing.

This is an extraordinary time to be involved in high performance computing. During these exciting times, I am grateful to our sponsors for their continued endorsement of our efforts. My special thanks and congratulations go to the ICL staff and students for their skill, dedication, and tireless efforts in making the ICL one of the best centers for enabling technology in the world.

OVERVIEW

Our mission at the Innovative Computing Laboratory (ICL) is to be a world leader in enabling technologies and software for scientific computing. Our vision is 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.

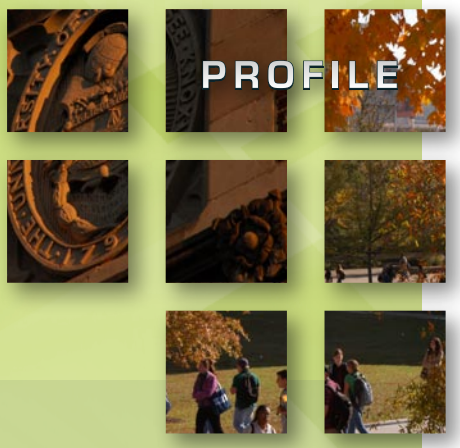

Located at the heart of the University of Tennessee campus in Knoxville, ICL continues to lead the way as one of the most respected academic, enabling technology research laboratories in the world. Our many contributions to technological discovery in the HPC community, as well as at UT, underscore our commitment to remain at the forefront of enabling technology research.

ICL's prominent role at the University of Tennessee is reflected by the Chancellor of the Knoxville campus, Dr. Loren Crabtree:

"Bolstered by a combination of talent and hard work, the staff and students of the Innovative Computing Laboratory represent the dedication and effort that have come to symbolize the contributions and achievements of the University of Tennessee. Led by Dr. Jack Dongarra, this research group and its outstanding accomplishments continue to embody our university's research mission and have helped solidify the stature of our university as a top publicly funded academic research institution. ICL has been, and will continue to be, instrumental in helping our university champion strategic partnerships around the globe as we pursue new alliances and collaborations with national and international research communities, including the nearby Oak Ridge National Laboratory. The University of Tennessee is proud to have the leadership of ICL in fostering our reputation for world-class research."

BACKGROUND

ICL was founded in 1989 by Dr. Jack Dongarra who came to the University of Tennessee from Argonne National Laboratory upon receiving a dual appointment as Distinguished Professor in the Computer Science Department and as Distinguished Scientist at nearby Oak Ridge National Laboratory (ORNL), two positions he holds today. Dr. Dongarra's arrival at UT ushered in a new era of enabling technology research as he built a strong foundation for collaboration and growth that attracted many post-doctoral researchers and professors from multi-disciplines such as mathematics, geology, chemistry, etc. Many of these scientists came to UT to work with Dr. Dongarra and remained as post-doctoral researchers, which began a long list of top research talent to pass through ICL. Below is a list of some of the researchers who have helped make ICL the respected organization it has become.

ZHAOJUN BAI
THE UNIVERSITY OF CALIFORNIA, DAVIS

RICHARD BARRETT
OAK RIDGE NATIONAL LABORATORY

ADAM BEGUELIN
AOL

SUSAN BLACKFORD
MYRICOM

HENRI CASANOVA
UNIVERSITY OF HAWAII

JAEYOUNG CHOI
SOONGSIL UNIVERSITY, KOREA

ANDY CLEARY
LAWRENCE LIVERMORE NATIONAL LABORATORY

FREDERIC DESPREZ
ÉNS-LYON

VICTOR EIJKHOUT
UNIVERSITY OF TEXAS, AUSTIN

GRAHAM FAGG
MICROSOFT

EDGAR GABRIEL
UNIVERSITY OF HOUSTON

ROBERT VAN DE GEIJN
THE UNIVERSITY OF TEXAS, AUSTIN

JULIEN LANGOU
UNIVERSITY OF 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
FORSCHUNGSZENTRUM JÜLICH

Over the past 17 years, ICL has produced numerous high value tools and applications that now compose the basic fabric of high performance, scientific computing. Some of the technologies that our research has produced include:

Active Netlib	**ATLAS**	**BLAS**	**FT-MPI**	**HARNESS**
LAPACK	**LINPACK Benchmark**	**MPI**	**Netlib**	**NetSolve**
PAPI	**PVM**	**RIB**	**ScaLAPACK**	**Top500**

Our successes continue along with current ICL efforts such as Fault Tolerant Linear Algebra, HPC Challenge benchmark suite (HPCC), LAPACK for Clusters (LFC), KOJAK, NetBuild, NetSolve/GridSolve, Open MPI, PAPI, SANS-Effort including SALSA, and vGrADS. Many of our efforts have been recognized nationally and internationally, which includes four R&D 100 awards; PVM in 1994, ATLAS and NetSolve in 1999, and then PAPI in 2001.

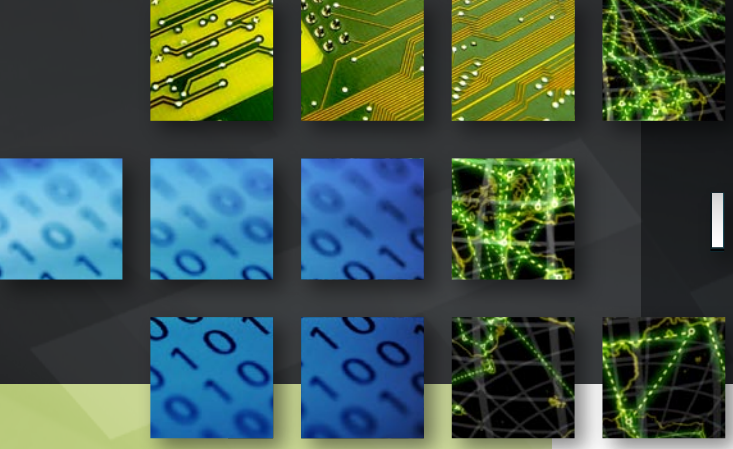

ICL RESEARCH

Our successful research efforts of the past have provided the foundation for addressing the challenges of the future. Recognizing that enabling technologies serve as catalysts for computational innovation, we continue to adapt to the ever increasing computational demands of the scientific community through our vision of what the next generation of enabling technologies can accomplish.

Numerical linear algebra, specifically the numerical libraries that encode its use in software, still serves as the foundation of our research. However, as our research efforts progressed we recognized the changing shape of the computational science landscape and the changing demands for enabling technology in high performance computing. Thus we expanded our research to include work in high performance and distributed computing. As we gained a solid understanding of the challenges presented in these domains, we expanded our research into performance analysis and benchmarking for high-end computers. Finally, as a by-product of a long tradition of delivering high quality software produced from our research coupled with our experience with the development of the Netlib repository, we have embraced new challenges in building robust, comprehensive, and well-organized asset management tools.

Recognizing the future challenges of HPC and taking initiatives to address them have not only allowed ICL to grow but have also allowed us to demonstrate the range and diversity of our research. In 2006-2007 we will engage in more than 15 significant research projects. On the following pages, we provide brief summaries of the research in each of our four main areas of focus; numerical libraries, high performance distributed computing, performance analysis and benchmarking, and asset management.

ACKNOWLEDGMENT

We continue to be entrusted by the government to participate in the advancement of the nation's research agenda. Without the support of the many agencies and organizations that have funded, and continue to fund, our efforts we simply would not be able to conduct cutting edge research. The main source of support has been federal agencies that are charged with investing the nation's computational research funding. Therefore we acknowledge the following for their support of our efforts past and present:

NSF DOE DoD DARPA ONR NIH

Strong support from private industry has also played a significant role. Some organizations have targeted specific ICL projects. But 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:

CACI AMD hp Dolphin IBM intel The MathWorks Mellanox Technologies Microsoft Myricom NORTHROP GRUMMAN sgi Quadrics

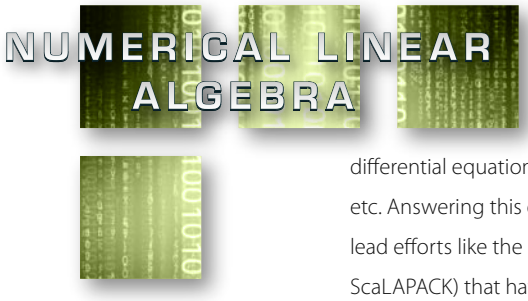

ICL has long been a leader in producing standards, algorithms, and software for numerical linear algebra, a quintessential ingredient of computational science. Sparse linear systems and eigenvalue calculations come from, among others, applications that involve partial differential equations, and dense operations arise from boundary element methods, quantum scattering, etc. Answering this demand, we have teamed up in the past with other researchers and industry to lead efforts like the BLAS Technical Forum and various linear algebra packages (e.g., LINPACK, LAPACK, ScaLAPACK) that have standardized programming interfaces and made performance portable across the plethora of modern hardware.

Our contributions to this community have drawn heavily on our expertise in high performance numerical linear algebra. Most of our on-going projects relate in one way or another to the concept of Self Adapting Numerical Software (SANS). The ATLAS project may be considered the bridge between the old style kernel computation optimization and SANS-style poly-algorithm approach. The latter is applied extensively in ICL's current projects such as Accels, LFC, and SALSA.

Members of the ICL team continue to address the issue of fault-tolerance in linear algebra codes, a growing concern for systems with ever increasing numbers of processors and hardware heterogeneity. We believe successful approaches to fault tolerance in next generation high end computing environments, where thousands of processors will be the norm, must leverage intimate knowledge of the application and its underlying numeric algorithms in order to achieve the efficiencies necessary in terms of adequate failure recovery times and system resource usage such as for additional disk and memory storage that is normally required. But the demands of this approach can be steep when compared to simpler, slower, resource hungry methods such as system wide, coordinated disk based checkpointing. As Supercomputer systems grow larger the possibilities of system and software failures increase, thus we are investing considerable effort into designing and implementing new custom fault tolerant numerical kernels through the FT-LA project. We have also recently begun to focus on algorithm based disk-free fault tolerance that uses various in-memory encodings to allow for scalable recovery of parallel codes.

Considerable progress has been made in the development of the SALSA system for heuristic decision making in the context of linear and nonlinear system solving. The software functions as an increasingly powerful testbed for iterative linear system solvers, using the available methods from the PETSc library, and attached packages such as Hypre. For internal use in the system, as well as for external use in matrix libraries or generally for communication between numerical software components, we have extended our metadata standard for matrix data that formalizes the matrix characteristics we analyze. We have

released a library that defines this standard through an API and an XML file format. A collection of analysis modules for generating characteristics of user input data was released in 2005. We also are continuing research on identifying statistical techniques and tools for building heuristics.

The LFC project draws on the conceptual underpinnings of the SANS software by choosing the best algorithm in a given context and tuning based on input data characteristics. The publicly available code includes all of ScaLAPACK's decompositional linear solvers. The in-development code adds computational server capabilities that add data persistence and all of ScaLAPACK's solvers including singular value and eigenvalue. LFC's ease of deployment has been proven time and time again while delivering the best possible performance levels across multiple architectures.

Recently, we have begun a major overhaul of the LAPACK and ScaLAPACK software collection with a focus on multicore. With renewed funding, new development is supplanting prior work done only by volunteers, which had fallen behind the theoretical and algorithmic advances that have been published or implemented elsewhere. Software and hardware landscapes continue to evolve requiring different techniques to make a successful programming library.

For the nano-physics community, we are studying certain eigenvalue problems where interior eigenvalues are needed, which are degenerate and close to a gap in the spectrum with known location (the gap separates energy levels of interest). This problem poses a considerable challenge to traditional methods, and we are investigating different solution approaches.

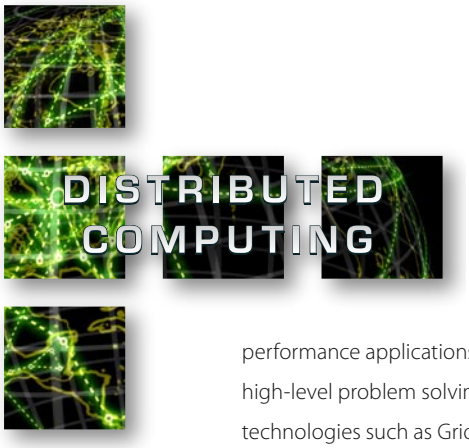

ICL's commitment to distributed computing spans more than fifteen years and has included involvement in a wide range of successful projects such as Parallel Virtual Machine (PVM) and the Message Passing Interface (MPI) standard. Distributed computing provides a fundamental platform for building efficient, modern, high performance applications. Currently, we are involved in multiple levels of distributed computing from high-level problem solving environments such as GridSolve/NetSolve and VGrADS, through middleware technologies such as GridRPC and various MPI implementations such as FT-MPI (Fault Tolerant MPI) and the open source community Open MPI project.

ICL RESEARCH

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 SuperComputing 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. Many features from FT-MPI such as runtime design, point-to-point RDMA messaging, buffer management and tuned collective communication algorithms are currently being applied to a new open source MPI implementation known as Open MPI. In a later stage of integration, the fault tolerant mechanisms of FT-MPI will also be added to Open MPI as a runtime user selectable module.

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. A new release of GridSolve addresses NAT and firewall traversal, improves service descriptions, and includes speed enhancements.

In addition, we are collaborating on the GridRPC initiative. The GridRPC API represents ongoing work to standardize and implement a portable and simple remote procedure call (RPC) mechanism for grid computing. This standardization effort is being pursued through the Global Grid Forum Research Group on Programming Models. The initial work on GridRPC shows that client access to existing grid computing systems such as NetSolve and Ninf can be unified via a common API, a task that has proven to be problematic in the past.

The Virtual Grid Application Development Software (VGrADS) is a multi-institution project lead by Rice University that addresses the fundamental problem of how to more effectively program and use highly complex and dynamic Grid systems. The VGrADS Execution System (vgES) has been developed to provide fast, scalable resource selection that allows end-users to manage their applications on a virtual grid that is dynamically instantiated at run time. We are also experimenting with a concept of statistical pseudo-reservations for resources, which allows large time-constrained work flow applications to ef-

fectively use multiple geographically distributed batch systems concurrently. ICL is also working on ways to schedule work flow applications that would be robust to the resource variations in such a dynamic Grid environment.

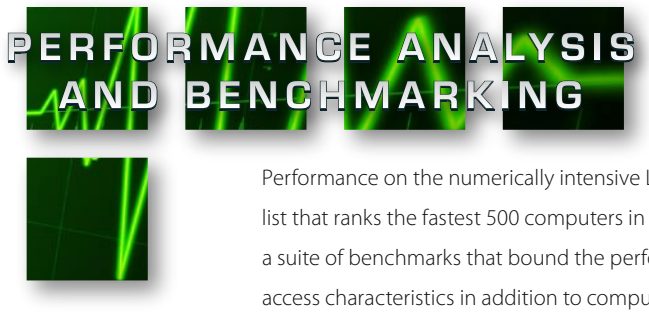

ICL continues to play a leadership role in benchmarking and performance evaluation efforts that measure and report performance on high performance computing (HPC) machines. Performance on the numerically intensive LINPACK benchmark is the basis of the semi-annual TOP500 list that ranks the fastest 500 computers in the world. The HPC Challenge (HPCC) project has developed a suite of benchmarks that bound the performance of many real applications as a function of memory access characteristics in addition to computational performance. HPCC includes several benchmarks for measuring sustainable memory bandwidth, the rate of random memory updates, and the latency and bandwidth of a number of communication patterns.

In addition to developing benchmarks, our research staff are actively involved in the development of performance evaluation tools and methodologies. As a basis for collection of accurate and relevant performance data, we have developed a portable library interface for access to hardware performance counters on most modern microprocessors. The interface, called the Performance API, or PAPI, is widely used by a number of research and commercial end-user performance analysis tools. KOJAK, a joint project with the Central Institute for Applied Mathematics at the Research Centre Jülich, is an automatic end-user performance analysis tool aimed at providing high-level feedback on performance behavior of parallel applications. KOJAK collects event traces at runtime and uses pattern recognition to convert the traces into information about performance bottlenecks relevant to developers.

Recent work on PAPI has extended the API to allow data and address range restrictions where this capability is supported by the underlying system. This extension allows hardware counter data to be collected in the context of specific data structures or instruction types. In order to enable access to both on-processor and off-processor counters and sensors simultaneously, another extension allows multiple platform-dependent substrates to be used together during the same application execution. PAPI has recently been ported to the new Montecito processor, and work is underway to provide a port

to the Cell processor. Multi-substrate support will be used to simultaneously measure counters for the heterogeneous components of hybrid multicore systems. The PAPI multi-substrate architecture is shown in the Figure below.

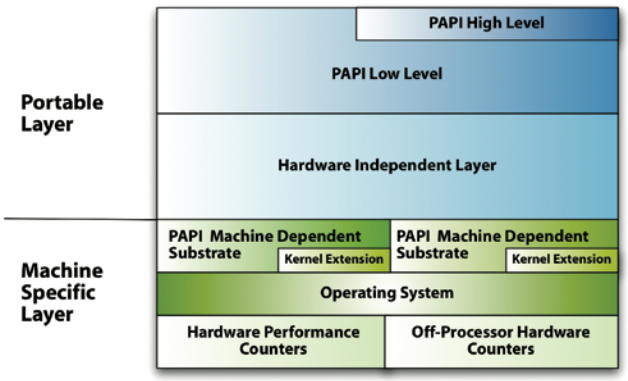

Extension to PAPI to support Multiple Substrates

Recent work on the KOJAK project is integrating trace data collection and feedback analysis into the OpenUH compiler framework being developed at the University of Houston. This integration will enable automated compiler optimization of OpenMP parallel programs based on performance feedback. Work is also underway to integrate PAPI and KOJAK with the Eclipse Parallel Tools Platform (PTP). Eclipse PTP provides an integrated development environment (IDE) for parallel programs which is extensible by means of plugins. A PAPI plugin will facilitate easy selection and invocation of hardware counter measurements for other performance tool plugins. A KOJAK plugin will automate the processes of selective trace instrumentation and distributed trace file analysis.

Other recent work in the performance optimization area is extending our previous work on automated empirical optimization of library software by applying the transformation, code generation, and search process to more general code segments and investigating different search algorithms, including parallel approaches.

ASSET MANAGEMENT

Our foundation for work in repositories and software asset management originates from the creation and development in the 1980s of the Netlib repository for mathematical software and other related resources.

Over the past two decades, Netlib has solidified its place as a forerunner of frameworks that store and distribute HPC software, tools, and other resources. Resulting from the proliferation of scientific computing and simulation in the US, the National High-performance Software Exchange (NHSE) was formed in the mid 1990s by several academic institutions and government agencies

with the primary goal of establishing discipline-oriented software repositories that could be contributed to and maintained by experts in their respective fields. ICL was one of the academic partners called upon to participate in this national effort.

One result of the NHSE effort was the development of the Repository in a Box (RIB) toolkit, which was produced to enable the creation and interoperability of discipline-oriented, web-based software repositories, specifically for the tools and applications generated by the HPC community. RIB has evolved to support the creation of repositories to store and share any type of digital object. RIB has been completely rewritten in Java to become more streamlined and flexible. Many new features have also been added, such as the ability to navigate through all object components in a tree-like structure from the catalog view and new intuitive interfaces to database operations that allow selective retrieving and displaying of information from multiple classes in an easy-to-read table format.

Building on our asset management experience, NetBuild was developed. This is a project to make it easier for authors and installers of application software to utilize standard computational software libraries, intercept calls to compilers and/or linkers, identify which libraries are needed for an application, locate those libraries, download them, install them, and link them into the executable. A new NetBuild client has been released with the ability to handle many more libraries, including the netBSD package collection. NetBuild can now support packages that use threads and multiple CPUs. Metadata support for representing compiler dependencies and data-format dependencies has also been added. Other recent changes include pre-loading support, include file support and cross-compiling support.

Another successful and ongoing project, The Remote Software Toolkit (ReST), develops a graphical framework for software management that allows software to be distributed, installed, monitored and tested in a distributed, heterogeneous environment. There are wizard tools that assist a software provider in developing packages for their software. The packages can then be deployed by simply placing them on a web page. When the packages are accessed, the ReST installer is automatically downloaded using Java Web Start, and the user is guided through installing the software on a (possibly remote) set of machines using a SSH based push mechanism. Additional tools for monitoring and testing software are being developed. All the tools are designed to be modular and easily extended for use in a wide range of projects.

PEOPLE

ICL Group, Fall 2006

Our most important asset, as with most organizations, is our staff. Our success hinges on our ability to skillfully apply our expertise to the computing challenges that confront the ever changing HPC community. The diversity of our full and part-time staff, which comprises individuals from all over the world including Italy, Poland, France, China, Thailand, and Korea, as well as the U.S., allows us to approach research problems from many directions.

Being a part of a Computer Science (CS) department at a large research university, we have a responsibility to combine teaching and research. With a CS program consisting of nearly 200 students we have been very proactive in securing graduate and undergraduate internships and assistantships for those students who are motivated, hard working, and willing to learn. Currently, we support more than a dozen students.

STAFF AND STUDENTS

THARA ANGSKUN
GRADUATE RESEARCH ASSISTANT

GEORGE BOSILCA
RESEARCH SCIENTIST

ALFREDO BUTTARI
SR RESEARCH ASSOCIATE

TOM CORTESE
RESEARCH SCIENTIST - CE LEAD

DAVID CRONK
RESEARCH DIRECTOR

JACK DONGARRA
UNIVERSITY DISTINGUISHED PROFESSOR

PENG DU
GRADUATE RESEARCH ASSISTANT

DON FIKE
RESEARCH ASSISTANT

TERESA FINCHUM
ADMINISTRATIVE SERVICES ASSISTANT

ERIKA FUENTES
GRADUATE RESEARCH ASSISTANT

JAN JONES
PUBLICATIONS COORDINATOR

JAKUB KURZAK
SR RESEARCH ASSOCIATE

JULIE LANGOU
RESEARCH ASSOCIATE

TRACY LEE
ACCOUNTING SPECIALIST II

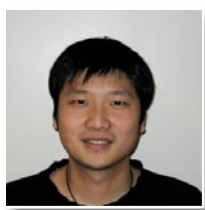

YINAN LI
GRADUATE RESEARCH ASSISTANT

DANIEL LUCIO
GRADUATE RESEARCH ASSISTANT

PIOTR LUSZCZEK
RESEARCH SCIENTIST

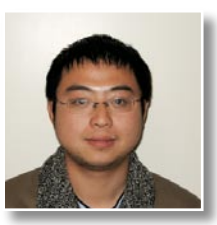

TENG MA
GRADUATE RESEARCH ASSISTANT

LEIGHANNE MILLS
ADMINISTRATIVE SUPPORT ASSISTANT

SHIRLEY MOORE
ASSOCIATE DIRECTOR

PEOPLE

STAFF AND STUDENTS CONTINUED

KEITH MOORE
SR. RESEARCH ASSOCIATE

TERRY MOORE
ASSOCIATE DIRECTOR

PHIL MUCCI
RESEARCH CONSULTANT

PAUL PELTZ
IT ADMINISTRATOR II

JELENA PJESIVAC-GRBOVIC
GRADUATE RESEARCH ASSISTANT

FARZONA PULATOVA
GRADUATE RESEARCH ASSISTANT

TRACY RAFFERTY
BUSINESS MANAGER

DAVID ROGERS
GRAPHIC DESIGNER

KEITH SEYMOUR
SR. RESEARCH ASSOCIATE

ZHIAO SHI
GRADUATE RESEARCH ASSISTANT

SERGEI SHINKAREV
IT ADMINISTRATOR

FENGGUANG SONG
GRADUATE RESEARCH ASSISTANT

DAN TERPSTRA
RESEARCH LEADER

JOE THOMAS
IT ANALYST III

STANIMIRE TOMOV
SR. RESEARCH ASSOCIATE

SCOTT WELLS
ASSISTANT DIRECTOR

JIAYI WU
GRADUATE RESEARCH ASSISTANT

ASIM YARKHAN
SR. RESEARCH ASSOCIATE

HAIHANG YOU
RESEARCH ASSOCIATE

As our list of collaborators and research partners continues to grow, so do the opportunities to host visitors from around the globe. Since our group was founded, we have routinely hosted numerous researchers, some of whom stay briefly to give seminars or presentations while many 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.

RECENT VISITORS TO ICL

DORIAN ARNOLD
UNIVERSITY OF WISCONSIN, MADISON

SCOTT ATCHLEY
MYRICOM

MARC BABOULIN
CENTER FOR RESEARCH AND ADVANCED TRAINING IN SCIENTIFIC COMPUTATION (CERFACS), FRANCE

RICHARD BARRETT
OAK RIDGE NATIONAL LABORATORY

SHAWN BROWN
PITTSBURGH SUPERCOMPUTER CENTER

LOUIS-CLAUDE CANON
ÉCOLE SUPÉRIEURE D'ÉLECTRONIQUE DE L'OUEST, FRANCE

MANOEL CUNHA
FEDERAL UNIVERSITY OF RIO DE JANEIRO

REMI DELMAS
INSTITUT NATIONAL DES SCIENCES APPLIQUÉES DE TOULOUSE (INSA), FRANCE

VICTOR EIJKHOUT
TEXAS ADVANCED COMPUTING CENTER

PATRICK GEOFFRAY
MYRICOM

EMMANUEL JEANNOT
LABORATOIRE LORRAIN DE RECHERCHE EN INFORMATIQUE ET SES APPLICATIONS (LORIA), FRANCE

MYUNG HO KIM
SOONGSIL UNIVERSITY - SEOUL, SOUTH KOREA

KEVIN LONDON
MICROSOFT

JEAN-FRANÇOIS PINEAU
LABORATOIRE DE L'INFORMATIQUE DU PARALLÉLISME
ÉCOLE NORMALE SUPÉRIEURE DE LYON, FRANCE

CRAIG RASMUSSEN
LOS ALAMOS NATIONAL LABORATORY

YVES ROBERT
LABORATOIRE DE L'INFORMATIQUE DU PARALLÉLISME
ÉCOLE NORMALE SUPÉRIEURE DE LYON, FRANCE

KEN ROCHE
OAK RIDGE NATIONAL LABORATORY

BERNARD TOURANCHEAU
UNIVERSITY OF LYON, FRANCE

FRÉDÉRIC VIVIEN
LABORATOIRE DE L'INFORMATIQUE DU PARALLÉLISME
ÉCOLE NORMALE SUPÉRIEURE DE LYON, FRANCE

BRIAN WYLIE
FORSCHUNGSZENTRUM JÜLICH, GERMANY

ICL ALUMNI

Proudly, we can boast that we have provided numerous students and Post-docs with invaluable experience. Many of our students have gone on to apply the experience and knowledge gained while working at ICL to careers with many of the largest companies in the computing industry including AOL, Cray, Hewlett Packard, Hitachi, IBM, Inktomi, Intel, Microsoft, Myricom, NEC, SGI, and many others. Furthermore, many of our former staff and students are now faculty at academic institutions around the world.

CAROLYN AEBISCHER 1990-1993
SUDESH AGRAWAL 2001-2006
BIVEK AGRAWAL 2004-2006
ED ANDERSON 1989-1991
PAPA ARKHURST 2003
DORIAN ARNOLD 1999-2001
ZHAOJUN BAI 1990-1992
ASHWIN BALAKRISHNAN 2001-2002
RICHARD BARRETT 1992-1994
ALEX BASSI 2000-2001
MICAH BECK 2000-2001
ADAM BEGUELIN 1991
ANNAMARIA BENZONI 1991
SCOTT BETTS 1997-1998
NIKHIL BHATIA 2003-2005
NOEL BLACK 2002-2003
SUSAN BLACKFORD 1989-2001
FERNANDO BOND 1999-2000
RANDY BROWN 1997-1999
MURRAY BROWNE 1998-1999
CYNTHIA BROWNE 2005
ANTONIN BUKOVSKY 1998-2003
GREG BUNCH 1995
HENRI CASANOVA 1995-1998
RAMKRISHNA CHAKRABARTY 2005
SHARON CHAMBERS 1998-2000
ZIZHONG CHEN 2001-2006
JAEYOUNG CHOI 1994-1995
ERIC CLARKSON 1998
ANDY CLEARY 1995-1997
JASON COX 1993-1997
CRICKET DEANE 1998-1999
FREDERIC DESPREZ 1994-1995
JUN DING 2001-2003
JIN DING 2003
MARTIN DO 1993-1994
LEON DONG 2000-2001
DAVID DOOLIN 1997
ANDREW DOWNEY 1998-2003
MARY DRAKE 1989-1992
JULIO DRIGGS 2002-2004
BRIAN DRUM 2001-2004
VICTOR EIJKHOUT 1992-2005
BRETT ELLIS 1995-2005
SHAWN ERICSON 2004
ZACHARY EYLER-WALKER 1997-1998
LISA EZZELL 2003-2004
GRAHAM FAGG 1996-2006
MARKUS FISCHER 1997-1998
XIAOQUAN FU 2003-2004
MEGAN FULLER 2006
EDGAR GABRIEL 2003-2004
LYNN GANGWER 2000-2001
TRACY GANGWER 1992-1993
NATHAN GARNER 2001-2006
KELLEY GARNER 1998
JONATHAN GETTLER 1996
ERIC GREASER 1993
STAN GREEN 1992-1996
ALICE GREGORY 2004-2006

HUNTER HAGEWOOD 2000-2001
CHRISTIAN HALLOY 1996-1997
SVEN HAMMARLING 1996-1997
HIDEHIKO HASEGAWA 1995-1996
SATOMI HASEGAWA 1995-1996
CHRIS HASTINGS 1996
DAVID HENDERSON 1999-2001
GREG HENRY 1996
SID HILL 1996-1998
GEORGE HO 1998-2000
JEFF HORNER 1995-1999
YAN HUANG 2000-2001
CHRIS HURT 2002
PAUL JACOBS 1992-1995
WEIZHONG JI 1999-2000
WEICHENG JIANG 1992-1995
SONG JIN 1997-1998
BALAJEE KANNAN 2001
DAVID KATZ 2002
YOUNGBAE KIM 1992-1996
MICHAEL KOLATIS 1993-1996
COIRE KYLE 2005
AMANDA LAAKE 2003-2004
JULIEN LANGOU 2003-2006
JEFF LARKIN 2003-2005
DONGWOO LEE 2000-2002
TODD LETSCHE 1993-1994
SHARON LEWIS 1992-1995
XIANG LI 2001
WEIRAN LI 2002
CHAOYANG LIU 2000
KEVIN LONDON 1996-2005
MATT LONGLEY 1999
RICHARD LUCZAK 2000-2001
ROBERT MANCHEK 1990-1996
TUSHTI MARWAH 2004
PAUL MCMAHAN 1994-2000
ERIC MEEK 2003-2006
JEREMY MILLAR 1998-2002
MICHELLE MILLER 1999-2003
CINDY MITCHELL 2001-2002
STEVEN MOULTON 1991-1993
SHANKAR NARASIMHASWAMI 2004-2005
PETER NEWTON 1994-1995
CAROLINE PAPADOPOULOS 1997-1998
LEELINDA PARKER 2002
ANTOINE PETITET 1993-2001
JAMES S. PLANK 1991-1992
ROLDAN POZO 1992-1994
TAMMY RACE 1999-2001
GANAPATHY RAMAN 1998-2000
KAMESH RAMANI 2003
MEI RAN 1999-2004
YVES ROBERT 1996-1997
KEN ROCHE 1999-2004
TOM ROTHROCK 1997-1998
TOM ROWAN 1993-1997
KIRAN SAGI 2001-2003
EVELYN SAMS 1998-1999
FARIAL SHAHNAZ 2001
MAJED SIDANI 1991-1992
SHILPA SINGHAL 1996-1998
THOMAS SPENCER 1999-2001
ERICH STROHMAIER 1995-2001
MARTIN SWANY 1996-1999
DAISUKE TAKAHASHI 2002
JUDI TALLEY 1993-1999
YUAN TANG 2005-2006
KEITA TERANISHI 1998
JOHN THURMAN 1998-1999
FRANCOISE TISSEUR 1997
BERNARD TOURANCHEAU 1993-1994
LAUREN VACA 2004
SATHISH VADHIYAR 1999-2003
ROBERT VAN DE GEIJN 1990-1991
SCOTT VENCKUS 1993-1995
REED WADE 1990-1996
MICHAEL WALTERS 2001-2005
R. CLINT WHALEY 1991-2001
SUSAN WO 2000-2001
FELIX WOLF 2003-2005
QIU XIA 2004-2005
TINGHUA XU 1998-2000
TAO YANG 1999
YUANLEI ZHANG 2001-2005
YONG ZHENG 2001
LUKE ZHOU 2000-2001
MIN ZHOU 2002-2004

PARTNERSHIPS

Over the past 17 years, we have aggressively sought to build lasting, collaborative relationships with both domestic and international research institutions. These relationships within the high performance computing (HPC) community have played an important role in our success. Leveraging the incredible growth of computational science, 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.

As part of our collaborative initiatives, we have built a strong portfolio of shared resources, both material and intellectual. We have also forged many lasting relationships with application and tool vendors, many of whom 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 are based on the BLAS, LAPACK, and ScaLAPACK specifications and software developed by ICL. The following lists highlight many of our domestic partners and collaborators. As our list of government and academic partners continues to grow, we also continue to search for opportunities to establish partnerships with HPC vendors.

The world map on page 21 shows the location of many of the domestic and international partners and collaborators in research with whom we continue to work.

DOMESTIC COLLABORATORS INCLUDE:

ANL
ARGONNE NATIONAL LABORATORY

BLAST
THE BASIC LINEAR ALGEBRA SUBPROGRAMS TECHNICAL FORUM

CACR
CALIFORNIA INSTITUE OF TECHNOLOGY CENTER FOR ADVANCED COMPUTING RESEARCH

DARPA
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

DoD
THE UNITED STATES DEPARTMENT OF DEFENSE

DoD HPCMP
THE DoD HIGH PERFORMANCE COMPUTING MODERNIZATION PROGRAM

DOE
THE UNITED STATES DEPARTMENT OF ENERGY

DOE2000

EMORY UNIVERSITY

HiPerSoft
CENTER FOR HIGH PERFORMANCE SOFTWARE RESEARCH

IBM
INTERNATIONAL BUSINESS MACHINES

INTEL CORPORATION

INTERNET2

ISI
INFORMATION SCIENCES INSTITUTE

I2-DSI
THE INTERNET2 DISTRIBUTED STORAGE INFRASTRUCTURE

JICS
THE JOINT INSTITUTE FOR COMPUTATIONAL SCIENCE

LANL
LOS ALAMOS NATIONAL LABORATORY

LLNL
LAWRENCE LIVERMORE NATIONAL LABORATORY

MICROSOFT RESEARCH

MRA
METACENTER REGIONAL ALLIANCE

NASA
NATIONAL AERONAUTICS AND SPACE ADMINISTRATION

NCSA
THE NATIONAL COMPUTATIONAL SCIENCE ALLIANCE

NHSE
THE NATIONAL HPCC SOFTWARE EXCHANGE

NIST
NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY

NPACI
THE NATIONAL PARTNERSHIP FOR ADVANCED COMPUTATIONAL INFRASTRUCTURE

NSF
THE NATIONAL SCIENCE FOUNDATION

ORNL CSMD
THE COMPUTER SCIENCE AND MATHEMATICS DIVISION OF OAK RIDGE NATIONAL LABORATORY

RICE UNIVERSITY

SDSC
SAN DIEGO SUPERCOMPUTING CENTER

SGI
SILICON GRAPHICS INCORPORATED

SUN MICROSYSTEMS

UNIVERSITY OF CALIFORNIA, BERKELEY

UNIVERSITY OF CALIFORNIA, SAN DIEGO

UTK-CS
THE COMPUTER SCIENCE DEPARTMENT OF THE UNIVERSITY OF TENNESSEE

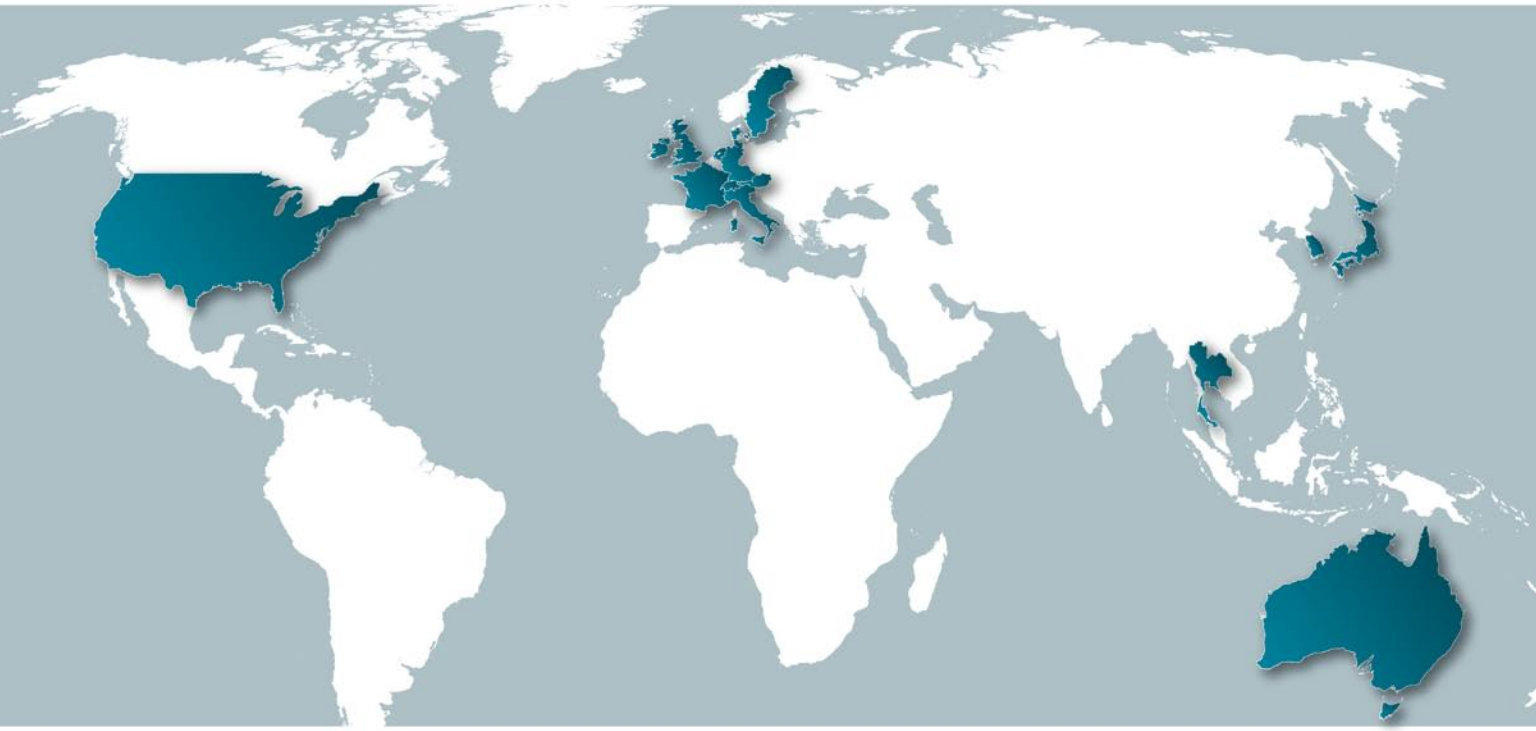

INTERNATIONAL COLLABORATORS INCLUDE:

DANISH COMPUTING CENTER FOR RESEARCH AND EDUCATION
LYNGBY, DENMARK

DEPARTMENT OF MATHEMATICAL AND COMPUTING SCIENCES, TOKYO INSTITUTE OF TECHNOLOGY TOKYO, JAPAN

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MANCHESTER
MANCHESTER, ENGLAND

EUROPEAN CENTRE FOR RESEARCH AND ADVANCED TRAINING IN SCIENTIFIC COMPUTING (CERFACS) TOULOUSE, FRANCE

UNIVERSITY COLLEGE DUBLIN DUBLIN, IRELAND

ECOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE LAUSANNE, SWITZERLAND

FAKULTÄT FÜR MATHEMATIK UND INFORMATIK 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 ICONE" DEL C.N.R. ROME, ITALY

INTELLIGENT SYSTEMS DESIGN LABORATORY, DOSHISHA UNIVERSITY KYOTO, JAPAN

KASETSART UNIVERSITY BANGKOK, THAILAND

LABORATOIRE DE L'INFORMATIQUE DU PARALLELISME, ÉCOLE NORMAL SUPERIEURE DE LYON LYON, FRANCE

MATHEMATICAL INSITUTE, UTRECHT UNIVERSITY NETHERLANDS

NUMERICAL ALGORITHMS GROUP LTD OXFORD, ENGLAND

PARALLEL AND HPC APPLICATION SOFTWARE EXCHANGE (PHASE)
TSUKUBA, JAPAN

LABORATOIRE RÉSEAUX HAUT DÉBITS ET SUPPORT D'APPLICATOINS MULTIMEDIA JEUNE EQUIPE DE L'UNIVERSITÉ CLAUDE BERNARD DE LYON (RESAM) LYON, FRANCE

RUTHERFORD APPLETON LABORATORY OXFORD, ENGLAND

SOONGSIL UNIVERSITY SEOUL, SOUTH KOREA

TECHNISCHE UNIVERSITAET WIEN VIENNA, AUSTRIA

UNIVERSITY OF UMEÅ UMEÅ, SWEDEN

HARDWARE RESOURCES

Understanding the nuances associated with hardware/software combinations requires testing and development on multiple platforms and architectures. Such heterogeneity is exhibited by the hardware resources on which we conduct our research. Not only do we have multiple, heterogeneous systems in-house, 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:

16 NODE INTEL P4 CLUSTER CONNECTED WITH MELLANOX AND MYRINET 10G

64 NODE INTEL EM64T CLUSTER CONNECTED WITH MYRINET 2000

INTEL ITANIUM CLUSTER

24 NODE AMD OPTERON CLUSTER CONNECTED WITH SILVERSTORM

64 NODE AMD OPTERON CLUSTER CONNECTED WITH MYRINET 2000

In addition to our own resources, we have access to multiple systems belonging to the Computer Science department, which include several server class machines and several HPC clusters. These clusters consist of multiple architectures including Itaniums, Itanium2s, Pentium 4s, and AMD processors that 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.

In addition, 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. The following are some of the systems that we currently utilize around the country:

CRAY X1E, XT3, AND XD1
IBM POWER 4, 4+, 5, CLUSTER 1600, AND BLUEGENE/L
HP XC SYSTEM
SEVERAL LARGE (256+ PROC) LINUX CLUSTERS
SGI ORIGIN 3000, 3800, 3900, AND ALTIX

For our Grid computing research, we utilize and help maintain a campuswide grid architecture known as the Scalable Intracampus Research Grid (SInRG), an NSF funded research infrastructure established by the Computer Science department under ICL leadership. This infrastructure provides hardware computing resources within the boundaries of the Knoxville campus for interdisciplinary research collaborations that are indicative of the national and international technology Grid. SInRG allows students, faculty, and other researchers at UTK, including ICL, to address important challenges of grid-based computing utilizing the advantages of local communication and central administration.

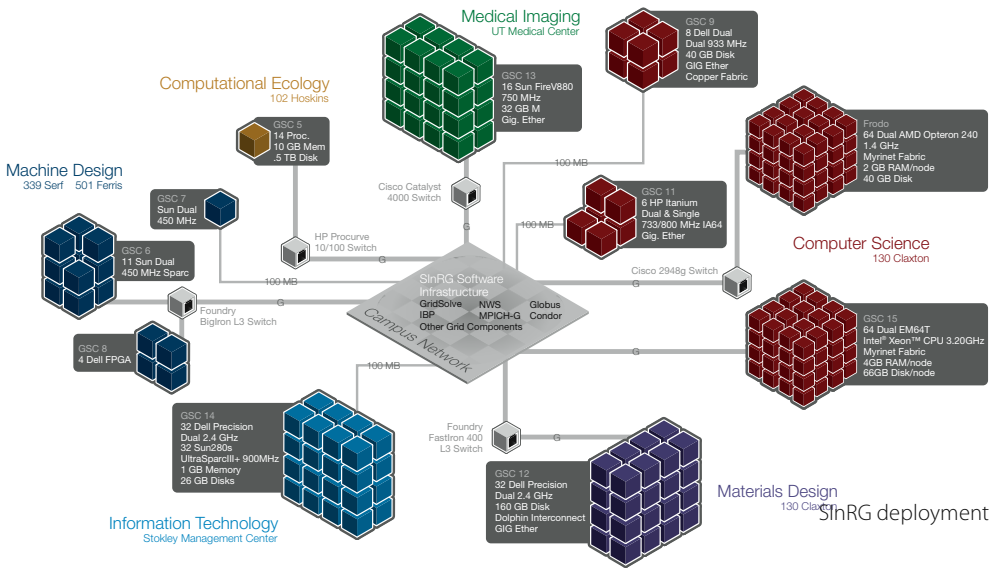

SInRG deployment

In addition, we possess an Access Grid (AG) node, which consists of various interfaces and environments on the Grid that support distributed meetings, lectures, tutorials, and other collaborative efforts. The AG comprises multiple video cameras, speakers, projectors, and PCs to form a seamless resource for conducting timely, online collaborative activities. The AG has become an invaluable tool and resource for collaborating with the many organizations and institutions with which we conduct joint research.

PUBLICATIONS

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

2006

Aguilera, G., Teller, P. J., Taufer, M., Wolf, F. **"A Systematic Multi-step Methodology for Performance Analysis of Communication Traces of Distributed Applications based on Hierarchical Clustering,"** *Proc. of the 5th International Workshop on Performance Modeling, Evaluation, and Organization of Parallel and Distributed Systems (PMEO-PDS 2006)* (to appear), Rhodes Island, Greece, IEEE Computer Society, ICL-UT-05-06, April 25-29, 2006.

Angskun, T., Fagg, G., Bosilca, G., Pjesivac-Grbovic, J., Dongarra, J. **"Scalable Fault Tolerant Protocol for Parallel Runtime Environments,"** *2006 Euro PVM/MPI,* Bonn, Germany, ICL-UT-06-12, 2006.

Angskun, T., Fagg, G., Bosilca, G., Pjesivac-Grbovic, J., Dongarra, J. **"Self-Healing Network for Scalable Fault Tolerant Runtime Environments,"** *DAPSYS 2006, 6th Austrian-Hungarian Workshop on Distributed and Parallel Systems,* Innsbruck, Austria, September 21-23, 2006.

Bosilca, G., Chen, Z., Dongarra, J., Eijkhout, V., Fagg, G., Fuentes, E., Langou, J., Luszczek, P., Pjesivac-Grbovic, J., Seymour, K., You, H., Vadhiyar, S. **"Self Adapting Numerical Software SANS Effort,"** *IBM Journal of Research and Development,* Volume 50, number 2/3, pp. 223-238, 2006.

Canning, A., Dongarra, J., Langou, J., Marques, O., Tomov, S., Voemel, C., Wang, L-W. **"Towards bulk based preconditioning for quantum dot computations,"** *IEEE/ACM Proceedings of HPCNano SC06* (to appear), 2006.

Canning, A., Dongarra, J., Langou, J., Marques, O., Tomov, S., Voemel, C., Wang, L-W. "Performance evaluation of eigensolvers in nano-structure computations," IEEE/ACM Proceedings of HPCNano SC06 (to appear), 2006.

Chen, Z., Dongarra, J. **"Algorithm-Based Checkpoint-Free Fault Tolerance for Parallel Matrix Computations on Volatile Resources,"** *IPDPS 2006, 20th IEEE International Parallel and Distributed Processing Symposium*, Rhodes Island, Greece, April 25-29, 2006.

Dewolfs, D., Broeckhove, J., Sunderam, V., Fagg, G. **"FT-MPI, Fault-Tolerant Metacomputing and Generic Name Services: A Case Study,"** *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, ICL-UT-06-14, Vol. 4192, Number 2006, pp. 133-140, 2006.

Dongarra, J. **"Performance of Various Computers Using Standard Linear Equations Software (Linpack Benchmark Report),"** *University of Tennessee Computer Science Department Technical Report*, UT-CS-04-526, UT-CS-89-95, 2006.

Dongarra, J., Bosilca, G., Chen, Z., Eijkhout, V., Fagg, G., Fuentes, E., Langou, J., Luszczek, P., Pjesivac-Grbovic, J., Seymour, K., You, H., Vadhiyar, S. **"Self-adapting numerical software (SANS) effort,"** *IBM Journal of Research and Development*, Vol. 50, no. 2/3, pp. 223-238, March/May, 2006.

Dongarra, J., Golub, G., Grosse, E., Moler, C., Moore, K. **"Twenty-Plus Years of Netlib and NA-Net,"** *University of Tennessee Computer Science Department Technical Report*, UT-CS-04-526, 2006.

Fagg, G., Pjesivac-Grbovic, J., Bosilca, G., Angskun, T., Dongarra, J. **"Flexible collective communication tuning architecture applied to Open MPI,"** *2006 Euro PVM/MPI* (submitted), Bonn, Germany, September, 2006.

Hernandez, O., Song, F., Chapman, B., Dongarra, J., Mohr, B., Moore, S., Wolf, F **"Performance Instrumentation and Compiler Optimizations for MPI/OpenMP Applications,"** *Second International Workshop on OpenMP*, Reims, France, June, 2006.

Keller, R., Bosilca, G., Fagg, G., Resch, M., Dongarra, J. **"Implementation and Usage of the PERUSE-Interface in Open MPI,"** *Recent Advances in Parallel Virtual Machine and Message Passing Interface, 13th European PVM/MPI User's Group Meeting, Lecture Notes in Computer Science*, Springer Heidelberg. Mohr, B., Traff, J. L., Worringen, J., Dongarra, J. (Eds.), Vol. 4192, pp 347 - 355. ISBN:3-540-39110-X, ISSN 0302-9743.

Kurzak, J., Dongarra, J. **"Implementation of the Mixed-Precision High Performance LINPACK Benchmark on the CELL Processor,"** *University of Tennessee Computer Science Tech Report*, UT-CS-06-580, LAPACK Working Note #177, September, 2006.

Kurzak, J., Dongarra, J. **"Implementing Linear Algebra Routines on Multi-Core Processors with Pipelining and a Look Ahead,"** *University of Tennessee Computer Science Tech Report*, UT-CS-06-581, LAPACK Working Note #178, September, 2006.

Kurzak, J., Dongarra, J. **"Linear Algebra Routines on Multi-Core Processors with Pipelining and a Look Ahead,"** *PARA 2006*, Umea Sweden, June 2006. (Also *LAPACK Working Note* 178 and *University of Tennessee CS Department Tech Report* UT-CS-06-581).

Langou, J., Langou, J., Luszczek, P., Kurzak, J., Buttari, A., Dongarra, J. **"Exploiting the Performance of 32 bit Floating Point Arithmetic in Obtaining 64 bit Accuracy,"** *University of Tennessee Computer Science Tech Report*, UT-CS-06-574, LAPACK Working Note #175, April, 2006.

London, K., Moore, S., Terpstra, D., Dongarra, J. **"Statistical Estimation of Processor Thermal Properties from a Standard Hardware Counter Interface,"** *The Second Workshop on High-Performance, Power-Aware Computing*, April 25, 2006, Rhodes, Greece in conjunction with International Parallel and Distributed Processing Symposium - IPDPS'2006 April 25-29, 2006, Rhodes, Greece.

Luszczek, P. **"High Performance Development for High End Computing with Python Language Wrapper (PLW),"** *International Journal of High Performance Computing Applications* (to appear), 2006.

Luszczek, P., Dongarra, J. **"Overview of the HPCC Benchmark Suite,"** *SPEC 2006 Workshop*, University of Texas at Austin, January 23, 2006.

Mohr, B., Traff, J. L., Worringen, J., Dongarra, J. (Eds.) **"Recent Advances in Parallel Virtual Machine and Message Passing Interface,"** *13th European PVM/MPI User's Group Meeting, Lecture Notes in Computer Science*, Vol. 4192, Springer Heidelberg, 2006. ISBN: 3-540-39110-X, ISSN 0302-9743.

Ni, J., Dongarra, J., Zheng, Y., Gu, G., Gentzsch, W., Jin,H. *Proceedings of the First International Multi-symposiums on Computer and Computational Sciences (IMSCCS|06),* IEEE Computer Society Press, Los Alamitos, California, USA, Vol. 1 and 2, 2006. ISBN 0-7695-2581-4.

Pjesivac-Grbovic, J., Angskun, Bosilca, G., Fagg, G., Gabriel, E., Dongarra, J. **"Performance Analysis of MPI Collective Operations,"** *Cluster Computing Journal* (to appear), 2006.

Pjesivac-Grbovic, J., Fagg, G., Angskun, T., Bosilca, G., Dongarra, J. **"MPI Collective Algorithm Selection and Quadtree Encoding,"** *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, ICL-UT-06-13, Vol. 4192, Number 2006, pp. 40-48, September, 2006.

Pjesivac-Grbovic, J., Fagg, G., Angskun, T., Bosilca, G., Dongarra, J. **"MPI Collective Algorithm Selection and Quadtree Encoding,"** *ICL Technical Report*, ICL-UT-06-11, 2006.

Seymour, K., You, H., Dongarra, J. **"ATLAS on the BlueGene/L -- Preliminary Results,"** *ICL Technical Report*, ICL-UT-06-10, May, 2006.

Shi, Z., Dongarra, J. **"Scheduling Workflow Applications on Processors with Different Capabilities,"** *Future Generation Computing Systems*, Vol. 22, pp 665-675.

Song, F., Dongarra, J., Moore, S. **"Experiments with Strassen's Algorithm: From Sequential to Parallel,"** *18th IASTED International Conference on Parallel and Distributed Computing and Systems PDCS 2006* (submitted), Dallas, Texas, November 13-15, 2006.

Tang, Y. **"Technical Comparison between several representative checkpoint/rollback solutions for MPI programs,"** *ICL Technical Report*, ICL-UT-06-09, 2006.

Tang, Y., Fagg, G., Dongarra, J. **"Proposal of MPI operation level Checkpoint/Rollback and one implementation,"** *Proceedings of IEEE CCGrid 2006*, IEEE Computer Society, 2006.

Voemel, C., Tomov, S., Wang, L-W., Marques, O., Dongarra, J. **"The use of bulk states to accelerate the band edge state calculation of a semiconductor quantum dot,"** *Journal of Computational Physics* (submitted), 2006.

Wolf, F., Freitag, F., Mohr, B., Moore, S., Wylie, B. **"Large Event Traces in Parallel Performance Analysis,"** *8th Workshop 'Parallel Systems and Algorithms' (PASA), Lecture Notes in Informatics*, Frankfurt/Main, Germany, Gesellschaft für Informatik, ICL-UT-06-08, March 13-16, 2006.

Wyrzykowski, R., Dongarra, J., Meyer, N., Wasniewski, J. (Eds.) *Proceedings of Parallel Processing and Applied Mathematics*, *6th International Conference PPAM 2005*, Springer Berlin/Heidelberg, 2006. ISBN: 3-540-34141-2, ISSN: 0302-9743.

YarKhan, A., Seymour, K., Sagi, K., Shi, Z., Dongarra, J. **"Recent Developments in GridSolve,"** *International Journal of High Performance Computing Applications* (Special Issue: Scheduling for Large-Scale Heterogeneous Platforms), Robert, Y eds. Sage Science Press, Vol. 20, Number 1, Spring, 2006.

Zunger, A., Franceschetti, A., Bester, G., Jones, W. B., Kim, K., Graf, P. A., Wang, L-W., Canning, A., Marques, O., Voemel, C., Dongarra, J., Langou, J., Tomov, S. **"Predicting the electronic properties of 3D, million-atom semiconductor nanostructure architectures,"** *J. Phys.*: Conf. Ser. 46, doi:10.1088/1742-6596/46/1/040, 292-298, 2006.

LATE 2005

Luszczek, P., Koester, D. **"HPC Challenge v1.x Benchmark Suite,"** *SC|05 Tutorial - S13*, Seattle, Washington, November 13, 2005.

Shimosaka, H., Hiroyasu, T., Mitsunori, M., Dongarra, J. **"Optimization Problem Solving System Using GridRPC,"** *IEEE Transactions on Parallel and Distributed Systems* (submitted), October, 2005.

Cyberinfrastructure Technology Watch

http://www.ctwatch.org/

In collaboration with the NSF funded Cyberinfrastructure Partnership (CIP), which includes the San Diego Supercomputing Center (SDSC) and the National Center for Supercomputing Applications (NCSA), ICL is leading a broad ranging publication effort called Cyberinfrastructure Technology Watch (CTWatch). CTWatch is intended to serve as a forum for ideas and opinion 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 today requires, CTWatch has been 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, we have created *CTWatch Quarterly*, an on-line serial publication modeled on a more traditional academic journal. Along a more experimental line, we have created CTWatch Blog, an on-line Weblog that provides commentary and informative links on the most recent developments and ideas occurring in the national and international cyberinfrastructure community.

CTWatch Quarterly is designed to be published on-line and is made available in both HTML and in a high quality format intended for printing on-demand. Each issue revolves around a particular area of interest for the cyberinfrastructure community and is organized by a guest editor who is a leader in that field. The focus topics (and corresponding guest editors) for 2006 included *International Cyberinfrastructure: Activities Around the Globe* (Radha Nandkumar and Thom Dunning), *Designing and Supporting Science-Driven Infrastructure* (Fran Berman and Thom Dunning), *Trends and Tools in Bioinformatics and Computational Biology* (Rick Stevens) and *High Productivity Computer Systems* (Jeremy Kepner).

In the near future we are planning new issues of the *Quarterly* on the approaching and highly disruptive watershed in high performance computing caused by new multicore and heterogeneous processor architectures, cyberinfrastructure and the social sciences, and the growing revolution in scientific communications.

CTWatch QUARTERLY

Available at http://www.ctwatch.org/quarterly/

2006 ISSUES

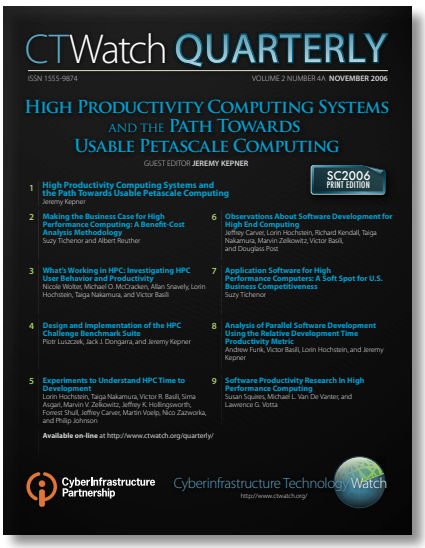

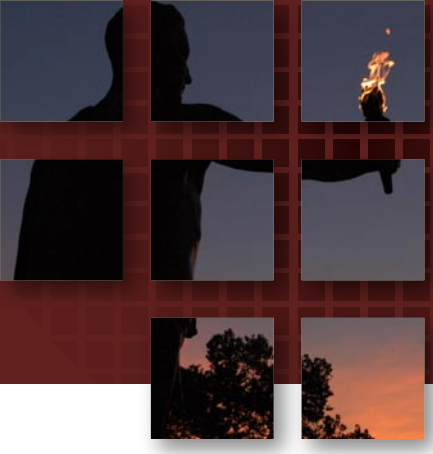

CITR

Center for Information Technology Research

From its inception in the spring of 2001, the Center for Information Technology Research (CITR), directed by Dr. Jack Dongarra and co-located with ICL, has fulfilled all the expectations that the University of Tennessee (UT) had when it established the Research Center program. CITR's mission has been to develop and enhance opportunities for multi-disciplinary and innovative Information Technology Research (ITR) at the University of Tennessee. In order to build up a thriving, well-funded community in basic and applied ITR, CITR's primary strategy has been to invest in a diverse group of ITR laboratories, each one led by an established researcher or an emerging leader in some significant area of ITR. Since first rate students and staff are indispensable to the success of such a strategy, CITR is also working to help develop the kind of educational environment that can recruit and train the people that a top flight IT research university requires.

This year CITR will work with faculty and administrators from several departments and colleges to help establish a new, University wide program in Computational Science that supports advanced degree concentrations in this critical new area across the curriculum. Under this program, 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. Although CITR has also made small investments in collateral activities - challenge grants for new IT researchers, contributions to start up packages for stellar new faculty, enhanced graduate stipends in ITR-related fields - it has concentrated on the ITR laboratories, and this concentration has produced the majority of its successes. Of the nine research centers of excellence - five in Knoxville and four at the Health Science Center in Memphis - CITR ranked second, bringing in $36.1 million in new research funding, just behind the Center of Genomics and Bioinformatics. Since UT's investment over that period was $2.7 million, CITR's rate of return on investment has been 13 to 1.

The full list of UT's Research Centers of Excellence (and their respective directors), includes the following:

CENTERS BASED IN KNOXVILLE

Center for Information Technology Research
Dr. Jack Dongarra

Tennessee Advanced Materials Laboratory
Dr. Ward Plummer

Center for Environmental Biotechnology
Dr. Gary Sayler

Food Safety Center
Dr. Stephen P. Oliver & Dr. Ann Draughon

Center for Structural Biology
Dr. Engin Serpesu

CENTERS BASED IN MEMPHIS

Center for Diseases of Connective Tissue
Dr. Andrew H. Kang

Center of Genomics and Bioinformatics
Dr. Dan Goldowitz & Dr. Robert W. Williams

Center for Neurobiology of Brain Diseases
Dr. William A. Pulsinelli

Vascular Biology Center
Dr. Lisa Jennings

ADDRESS	Innovative Computing Laboratory Suite 413 Computer Science Department 1122 Volunteer Blvd University of Tennessee Knoxville, TN 37996-3450
WEB SITE	http://icl.cs.utk.edu/
E-MAIL	icl@cs.utk.edu
PHONE	(865) 974-8295
FAX	(865) 974-8296

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