INNOVATIVE COMPUTING LABORATORY

2007-2008 REPORT

THE UNIVERSITY of TENNESSEE
Department of Electrical Engineering and Computer Science

INNOVATIVE COMPUTING LABORATORY 2007-2008 REPORT

EDITED BY SCOTT WELLS DESIGNED BY DAVID ROGERS

INNOVATIVE COMPUTING LABORATORY 2007-2008 REPORT

Contents

JACK DONGARRA

FROM THE DIRECTOR

The Innovative Computing Laboratory now has 18 years under its wings and has seen many changes here at UTK and ORNL. This year (2007) 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
Director of the Innovative Computing Laboratory

OVERVIEW

At the Innovative Computing Laboratory (ICL) our mission is simple. We intend 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.

Since its inception in 1989, ICL has long been a part of the University of Tennessee's Computer Science Department. In July 2007, the Computer Science Department, originally part of the College of Arts and Sciences, and the Electrical and Computer Engineering Department, part of the College of Engineering, merged to form the Department of Electrical Engineering and Computer Science (EECS). With this merger, the new department is now solely part of the College of Engineering.

As a research center now affiliated with the College of Engineering, ICL is located at the heart of the Knoxville campus. Sustaining our history of success, we continue to lead the way as one of the most respected academic, enabling technology research centers in the world. Our commitment to excellence has been one of the keys to our many contributions to technological discovery in the HPC community.

Our importance to not only the University of Tennessee system and the state, but also our nation, is reflected by the Chancellor of the Knoxville campus, Dr. Loren Crabtree:

> "On behalf of the entire university, it is a privilege to recognize the importance of the Innovative Computing Laboratory to the university's research mission. Led by Distinguished Professor Jack Dongarra, ICL continues to set the standard for academic research centers in the 21st century. As one of the university's most respected centers, the students and staff of ICL continue to demonstrate the dedication, leadership, and accomplishments that embody the university's ongoing efforts to remain one of the top publicly funded academic research institutions in the United States. Going forward, I also expect ICL to continue to play a major role in helping the university establish and foster national and international collaborations, including our ongoing partnerships with Oak Ridge National Laboratory and the construction in Tennessee of the NSF's new petascale supercomputing center. The future of research demands that academic institutions raise the bar for instruction and exploration. The University of Tennessee is proud to be the home of world-class centers such as ICL and we look forward to its continued contributions to our nation's research agenda."

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. What began with Dr. Dongarra and a single graduate assistant has evolved into a fully functional center, with a staff of more than 40 researchers, students, and administrators. Througout the past 18 years, ICL has attracted many post-doctoral researchers and professors from multi-disciplines such as mathematics, chemistry, etc. Many of these scientists came to UT specifically to work with Dr. Dongarra, which began a long list of top research talent to pass through ICL and move on to make exciting contributions at other institutions and organizations. Below we recognize just a few who have helped make ICL the respected center it has become.

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	Piotr Luszczek	The MathWorks
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 Julich, Germany

Since its inception, ICL has produced many 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	LAPACK for Clusters (LFC)	LINPACK Benchmark
MPI	NetBuild	Netlib	NetSolve
PAPI	PVM	RIB	ScaLAPACK
Top500			

Our successes continue along with current ICL efforts such as Fault Tolerant Linear Algebra, Generic Code Optimization (GCO), HPC Challenge benchmark suite (HPCC), KOJAK, Multi-core and Cell effort, NetSolve/GridSolve, Open MPI, PAPI, SALSA, SCALASCA, and vGrADS. 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

The ever changing landscape of enabling technology research places incredible demands on innovative thinking. In addition, the rapid evolution in computational technology presents enormous challenges to remaining at the forefront of discovery. Our successful 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, and our determination continue to position ICL as a leader in academic research.

What originally began as in depth investigation of the numerical libraries that encode the use of linear algebra in software, our research portfolio has grown extensively. To address the changing shape of the computational science landscape and the challenging demands for enabling technology in high performance computing, we have evolved and expanded our research agenda. 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 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.

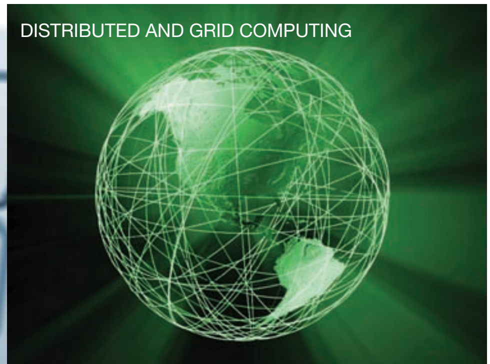

Demonstrating the range and diversity of our research, we will be engaged in more than 15 significant research projects during 2007-2008 across our four main areas of focus: numerical linear algebra, high performance distributed computing, performance analysis and benchmarking, and asset management. On the following pages, we provide brief summaries of our efforts in each of these areas.

NUMERICAL LINEAR ALGEBRA

FPGAs

ICL is involved in the development and evaluation of innovative linear algebra (LA) algorithms for Field Programmable Gate Arrays (FPGAs). FPGAs are specialized semiconductor devices that can change their hardware function from moment to moment through a simple re-programmable mechanism. Their appeal is based on their power efficiency and on the fact that operations executed on them can be orders of magnitude faster than their software equivalents. We are interested in the use of FPGAs to speed up not only basic dense linear algebra algorithms, but also compute intensive kernels of complex scientific codes, especially as applied to electromagnetic wave propagation, computational fluid dynamics, and particle and radiation transport.

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. One of the accomplishments for 2007 was to develop an extremely scalable, fault tolerant (dense) matrix-matrix multiply subroutine. Our future work in this area involves the implementation of a scalable fault-tolerant (dense)LU factorization.

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 multi-core BLAS library whereas the ScaLAPACK libray is parallel distributed, relies on BLAS, LAPACK, MPI, and BLACS libraries. LAPACK 3.1.1 was released in February 2007. Recent work on LAPACK has evolved around variance of factorization, iterative refinement for symmetric positive definite systems, and porting to Windows. In 2007, a special effort was made to port LAPACK and ScaLAPACK under Windows natively, and thus LAPACK 3.1.1 for Windows and ScaLAPACK 1.8.0 for Windows have been released. ScaLAPACK 1.8.0 is expected to be released in April 2008, which will involve the externalization of the LAPACK routines; starting from 1.8.0, the LAPACK library will need to be installed on the machine in order to link/run a ScaLAPACK application.

Multi-core/Cell

As part of our multi-core effort, we have designed a class of tiled algorithms for multi-core processors, including LU, Cholesky and QR factorizations, which outperform vendor implementations, in particular Intel's MKL. The 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 2007, we are working on utilization of both Cell processors in the QS20 blade by efficiently exploiting the NUMA architecture of the blade using libnuma library and ScaLAPACK-like block cyclic (1D) data partitioning. We are also planning an implementation of our recently designed tilet QR algorithm with internal blocking, which promises 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. As of late 2007 the results for SMP systems are very encouraging. The results for the Cell processor identify the scheduling mechanism as being a bottleneck when running on the PPE.

Nanotechnology

As part of our Department of Energy Nanotechnology effort, 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.

SALSA

Considerable progress has been made in the development of the Self-Adapting Large-scale Solver Architecture (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 continued to research different statistical and machine learning techniques to develop strategies that can successfully help us determine whether a solver is reliable or not, and when it is the most efficient choice for certain problems. Based on this research and various experimental results, we have developed a basic recommendation approach for selecting the most suitable methods for new problems. We are continuing to research different parameters and options to better tune our recommendation heuristics, as well as experimentation with new data.

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. The ROSE project from Lawrence Livermore National Laboratory provides, among other things, a source-to-source code transformation tool that can produce blocked and unrolled versions of the input code. Combined with our search heuristic, we can use the ROSE LoopProcessor to perform empirical code optimization.

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 recent version of the code was released earlier 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.

KOJAK

KOJAK is a package for automated performance analysis of message passing and shared memory programming paradigms. KOJAK's tracing system has already been ported to the Windows Compute Cluster platform to enable analysis of applications using Microsoft's MPI infrastructure. Recent work has been devoted to investigating techniques for coupling automated inefficiency pattern detection with manual trace visualization, highlighting inefficiency patterns in the visualized trace file. Work continues on the integration of KOJAK with the OpenUH compiler infrastructure and on developing Eclipse plug-ins, integrating performance tools into the popular development IDE.

Performance API (PAPI)

PAPI, the Performance API, has become the de facto standard within the High Performance Computing 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 architectures, including the Cray X2 (Black Widow), IBM Cell and POWER6, and multicore offerings from Intel and AMD. Architecturally, PAPI has been restructured into a framework and multiple components. This allows access to both on-processor and off-processor counters and sensors simultaneously. Component PAPI is available as a technology pre-release with example components provided for limited network counter and thermal monitoring support.

PERI

The Performance Engineering Research Institute (PERI) is a Department of Energy sponsored Scientific Discovery through Advanced Computing (SciDAC) project focusing on the performance modeling and automated tuning of parallel scientific applications. ICL's involvement in the automated tuning part of the effort builds on previous success of the ATLAS project, with the idea being to extend ATLAS's empirical tuning approach to additional library routines and key application kernels used by SciDAC codes. ICL also leads the application engagement "Tiger Team" for PERI's interactions with the Gyrokinetic Toroidal Code (GTC) development team, and ICL researchers have contributed to significant performance improvement of this code.

SCALASCA

As a parallel trace analysis tool for message passing applications, based on the KOJAK approach, SCALASCA has a primary goal of scalability for the massive levels of parallelism of today's leadership class computers (tens of thousands of processors and beyond). Instead of sequentially analyzing a single global trace file, SCALASCA analyzes separate local trace files in parallel by replaying the original communication on as many CPUs as have been used to execute the target application itself. SCALASCA v0.5 was released as a technology preview edition in January 2007. V0.9, released in August 2007, incorporates a runtime summarization mode. Both summarization and tracing with SCALASCA have demonstrated to be effective tools up to 22,000 CPUs. Both SCALASCA and KOJAK are collaborative efforts with Forschungszentrum Jülich, Germany.

TOP500

For nearly 15 years now, a ranking of the top 500 fastest computers in the world has been compiled biannually with publication occurring in June and November. The basis for this list is computer performance running the numerically intensive LINPACK benchmark developed by ICL. While other benchmarks, including HPCC, have been developed to measure performance of HPC systems, the TOP500 still relies on the LINPACK 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 2007.

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 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 OpenMPI as part of a collaborative effort involving several institutions including ICL. By the end of 2007, the fault tolerant mechanisms of FT-MPI will also be added to OpenMPI as a runtime user selectable module. Simultaneously, additional fault tolerance modules are being developed, shaping the OpenMPI library into the first MPI library supporting all existing models of fault tolerance: coordinated, uncoordinated and FT-MPI.

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. A new release of GridSolve includes improved scheduling capabilities, new client interfaces (Octave and IDL), MySQL compatibility, performance enhancements, and bug fixes.

vGrADS

We are involved in the Virtual Grid Application Development Software (VGrADS) project, another multi-institution effort that addresses fundamental problems in effectively programing and using highly dynamic Grid systems. In order to adapt to changing Grid resources, an application describes its resource requirements (processors, connectivity, etc) as an abstract virtual grid. At runtime, the VGrADS Execution System (vgES) schedules the actual resources and monitors the execution. 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. In recent work, a fault-tolerant scheduler has been added to enable deadline-based scheduling in the face of dynamic and unreliable resources and services.

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 problem solving environments. 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. 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 at Knoxville (ICL), and the University of Wisconsin at Madison.

PEOPLE

ICL GROUP, 2007

Remaining at the forefront of research and discovery requires that we employ some of the brightest and most determined people. Therefore, 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. Invariably, one of our greatest strengths is the diversity of our full and part-time staff, which comprises individuals from all over the world including Austria, Bulgaria, China, France, Italy, Poland, and Thailand, as well as the U.S. Such diversity allows us to apply a variety of unique and innovative approaches to the challenges and problems inherent in world-class computing.

In addition, as part of an egineering college at a large research university, we have a responsibility to combine teaching and research. We have been, and will continue to be, very proactive in securing graduate and undergraduate internships and assistantships for those students who are motivated, hard working, and willing to learn. It is not uncommon for us to support more than a dozen students in an academic year.

ICL Staff and Students


Wes Alvaro
STUDENT ASSISTANT


Thara Angskun
GRADUATE RESEARCH ASSISTANT

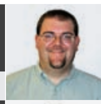
Barry Britt
IT ADMINISTRATOR I


George Bosilca
RESEARCH SCIENTIST

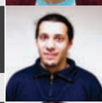
Aurelian Bouteiller
SENIOR RESEARCH ASSOCIATE

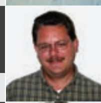
David Cronk
RESEARCH DIRECTOR

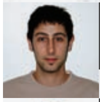
Alfredo Buttari
SENIOR RESEARCH ASSOCIATE

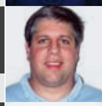
Tom Cortese
RESEARCH SCIENTIST - CE LEAD

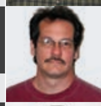
Don Fike
RESEARCH ASSISTANT


Jack Dongarra
UNIVERSITY DISTINGUISHED PROFESSOR

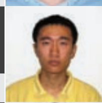
Peng Du
GRADUATE RESEARCH ASSISTANT

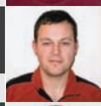
Karl Fuerlinger
SENIOR RESEARCH ASSOCIATE

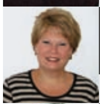
Teresa Finchum
ADMINISTRATIVE SERVICES ASSISTANT

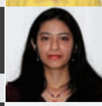
Erika Fuentes
GRADUATE RESEARCH ASSISTANT

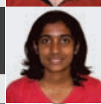
Madhuri Kasam
GRADUATE RESEARCH ASSISTANT

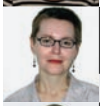
Jan Jones
PUBLICATIONS COORDINATOR

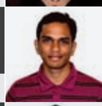
Venkata Kakani
STUDENT

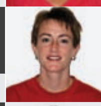
Tracy Lee
ACCOUNTING SPECIALIST III

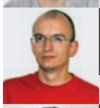
Jakub Kurzak
SENIOR RESEARCH ASSOCIATE

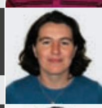
Julie Langou
RESEARCH ASSOCIATE

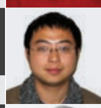
Teng Ma
GRADUATE RESEARCH ASSISTANT


Yinan Li
GRADUATE RESEARCH ASSISTANT

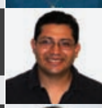
Daniel Lucio
GRADUATE RESEARCH ASSISTANT


Terry Moore
ASSOCIATE DIRECTOR

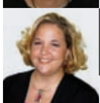
Leighanne Mills
ADMINISTRATIVE SUPPORT ASSISTANT

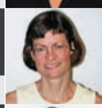
Shirley Moore
ASSOCIATE DIRECTOR

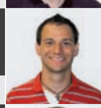
Paul Peltz
IT ADMINISTRATOR II

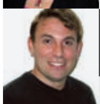
Phil Mucci
RESEARCH CONSULTANT

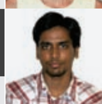
Dilip Patlolla
GRADUATE RESEARCH ASSISTANT

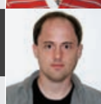
Keith Seymour
SENIOR RESEARCH ASSOCIATE

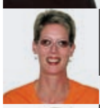
Tracy Rafferty
COORDINATOR II

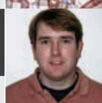
David Rogers
GRAPHIC DESIGNER


Dan Terpstra
RESEARCH LEADER


Gwang Son
GRADUATE RESEARCH ASSISTANT

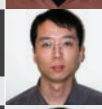
Fengguang Song
GRADUATE RESEARCH ASSISTANT


Scott Wells
ASSISTANT DIRECTOR FOR COMMUNICATIONS


Joe Thomas
IT ANALYST III

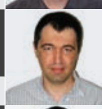
Stanimire Tomov
RESEARCH SCIENTIST


Asim YarKhan
SENIOR RESEARCH ASSOCIATE


Haihang You
RESEARCH ASSOCIATE

Visitors

Because our research is highly collaborative and because we have so many research collaborators and partners, we have numerous opportunities to host visitors from around the globe. Since our group was founded, we have routinely hosted many researchers, 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

Over the 18 years of its existence, ICL has been proud to employ many outstanding students and staff. From undergraduate students to experienced research scientists, more than 150 individuals have passed through our doors, many of whom have gone on to both contribute and achieve great things in academia, government, and the private sector. In fact, 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 AOL, Cray, Hewlett Packard, Hitachi, IBM, Intel, Microsoft, Myricom, NEC, SGI, and many others. In addition, many of our former staff and students are now faculty at some of the most prominent academic institutions around the world.

ICL ALUMNI WITH LAST KNOWN AFFILIATION

Carolyn Aebischer	1990-1993	Antonin Bukovsky	1998-2003	Victor Eijkhout TEXAS ADVANCED COMPUTING CENTER	1992-2005
Sudesh Agrawal	2001-2006	Greg Bunch	1995	Brett Ellis OAK RIDGE NATIONAL LABORATORY	1995-2005
Bivek Agrawal	2004-2006	Henri Casanova UNIVERSITY OF HAWAII AT MANOA	1995-1998	Shawn Ericson	2004
Ed Anderson EPA	1989-1991	Ramkrishna Chakrabarty	2005	Zachary Eyler-Walker	1997-1998
Daniel Andrzejewski UT EECS	2007	Sharon Chambers SIMULATION DYNAMICS	1998-2000	Lisa Ezzell PELLISSIPPI STATE TECHNICAL COMMUNITY COLLEGE	2003-2004
Papa Arkhurst	2003	Zizhong Chen JACKSONVILLE STATE UNIVERSITY	2001-2006	Graham Fagg MICROSOFT CORPORATION	1996-2006
Dorian Arnold UNIVERSITY OF WISCONSIN, MADISON	1999-2001	Jaeyoung Choi SOONGSIL UNIVERSITY, KOREA	1994-1995	Markus Fischer HPC CLUSTER SOLUTIONS	1997-1998
Zhaojun Bai UNIVERSITY OF CALIFORNIA, DAVIS	1990-1992	Eric Clarkson	1998	Xiaoquan Fu	2003-2004
Ashwin Balakrishnan	2001-2002	Andy Cleary LAWRENCE LIVERMORE NATIONAL LABORATORY	1995-1997	Megan Fuller EAST TENNESSEE STATE UNIVERSITY	2006
Richard Barrett OAK RIDGE NATIONAL LABORATORY	1992-1994	Camille Coti INRIA (LRI)	2007	Edgar Gabriel UNIVERSITY OF HOUSTON	2003-2004
Alex Bassi	2000-2001	Jason Cox	1993-1997	Lynn Gangwer	2000-2001
Micah Beck LOCI LAB, UT EECS	2000-2001	Cricket Deane	1998-1999	Tracy Gangwer	1992-1993
Adam Beguelin FORMERLY OF AOL, CURRENTLY RETIRED	1991	Frederic Desprez ÉNS LYON, FRANCE	1994-1995	Nathan Garner CLEVELAND STATE COMMUNITY COLLEGE	2001-2006
Annamaria Benzoni	1991	Jun Ding	2001-2003	Kelley Garner	1998
Scott Betts	1997-1998	Jin Ding	2003	Jonathan Gettler	1996
Nikhil Bhatia OAK RIDGE NATIONAL LABORATORY	2003-2005	Martin Do	1993-1994	Eric Greaser	1993
Noel Black	2002-2003	Leon Dong	2000-2001	Stan Green	1992-1996
Susan Blackford MYRICOM	1989-2001	David Doolin	1997	Alice Gregory ACTS FLEET MAINTENANCE	2004-2006
Fernando Bond	1999-2000	Andrew Downey	1998-2003	Hunter Hagewood NEVOA NETWORKS	2000-2001
Randy Brown	1997-1999	Mary Drake	1989-1992	Christian Halloy ORNL/JICS	1996-1997
Cynthia Browne HARVARD UNIVERSITY	2005	Julio Driggs	2002-2004	Sven Hammarling RETIRED FROM NAG	1996-1997
Murray Browne TURNER BROADCASTING	1998-1999	Brian Drum	2001-2004	Hidehiko Hasegawa UNIVERSITY OF TSUKUBA, JAPAN	1995-1996

2007 VISITORS TO ICL

Ayaz Ali	Yves Dolce	Heike Jagode	Anna Matsekh
University of Houston	Microsoft Corporation	Dresden University of Technology	Los Alamos National Lab
Marc Baboulin	Victor Eijkhout	Emmanuel Jeannot	Ricardo Ortiz
CERFACS, France	Texas Advanced Computing Center	INRIA, France	University of Iowa
Adam Beguelin	Riadh Fezzani	Myung Ho Kim	Craig Rasmussen
AOL Video	CERFACS, France	Soongsil University, Seoul, South Korea	Los Alamos National Lab
Daniel Becker	Marc Garbey	Julien Langou	Ken Roche
Forschungszentrum Jülich, Germany	University of Houston	University of Colorado at Denver	Oak Ridge National Laboratory
Zizhong Chen	John Goodhue	Hatem Ltaief	Felix Wolf
Jacksonville State University	SiCortex	University of Houston	Forschungszentrum Jülich, Germany
Camille Coti	Joshua Hursey	Piotr Luszczek	Lamia Youseff
INRIA, France	Indiana University	MathWorks	University of California, Santa Barbara

Satomi Hasegawa HITACHI	1995-1996	Robert Manchek MEDIAONE	1990-1996	Majed Sidani J. P. MORGAN	1991-1992
Chris Hastings	1996	Tushti Marwah	2004	Shilpa Singhal	1996-1998
David Henderson KNOX. VOLUNTEER EMERGENCY RESCUE	1999-2001	Paul McMahan IBM TIVOLI SOFTWARE	1994-2000	Thomas Spencer	1999-2001
Greg Henry INTEL	1996	Eric Meek UT MEDICAL SCHOOL	2003-2006	Erich Strohmaier NERSC	1995-2001
Sid Hill	1996-1998	Jeremy Millar US AIR FORCE	1998-2002	Martin Swany UNIVERSITY OF DELAWARE	1996-1999
George Ho	1998-2000	Michelle Miller	1999-2003	Daisuke Takahashi UNIVERSITY OF TOYKO, JAPAN	2002
Jeff Horner VANDERBILT UNIVERSITY	1995-1999	Cindy Mitchell BAPTIST HEALTH SYSTEM	2001-2002	Judi Talley	1993-1999
Yan Huang	2000-2001	Keith Moore	1987-2007	Yuan Tang	2005-2006
Chris Hurt	2002	Steven Moulton	1991-1993	Keita Teranishi	1998
Paul Jacobs	1992-1995	Shankar Narasimhaswami	2004-2005	John Thurman	1998-1999
Weizhong Ji	1999-2000	Peter Newton	1994-1995	Francoise Tisseur MANCHESTER UNIVERSITY, ENGLAND	1997
Weicheng Jiang	1992-1995	Caroline Papadopoulos	1997-1998	Bernard Tourancheau UNIVERSITY OF LYON, FRANCE	1993-1994
Song Jin	1997-1998	Leelinda Parker	2002	Lauren Vaca	2004
Balajee Kannan DISTRIBUTED INTELLIGENCE LAB, UTK	2001	Antoine Petitet ESI	1993-2001	Sathish Vadhiyar INDIAN INSTITUTE OF SCIENCE, INDIA	1999-2003
David Katz	2002	Jelena Pjesivac-Grbovic GOOGLE	2007	Robert van de Geijn UNIVERSITY OF TEXAS, AUSTIN	1990-1991
Youngbae Kim	1992-1996	James S. Plank UT EECS	1991-1992	Scott Venckus	1993-1995
Michael Kolatis	1993-1996	Roldan Pozo NIST	1992-1994	Reed Wade CATALYST IT LIMITED	1990-1996
Coire Kyle	2005	Farzona Pulatova	2005-2006	Michael Walters	2001-2005
Amanda Laake	2003-2004	Tammy Race	1999-2001	R. Clint Whaley UNIVERSITY OF TEXAS, SAN ANTONIO	1991-2001
Julien Langou UNIVERSITY OF COLORADO AT DENVER	2003-2006	Ganapathy Raman MICROSOFT	1998-2000	Susan Wo	2000-2001
Jeff Larkin OAK RIDGE NATIONAL LABORATORY	2003-2005	Kamesh Ramani	2003	Felix Wolf FORSCHUNGSZENTRUM JÜLICH, GERMANY	2003-2005
DongWoo Lee	2000-2002	Mei Ran	1999-2004	Jiayi Wu AMAZON.COM	2004-2007
Todd Letsche	1993-1994	Yves Robert ÉNS LYON	1996-1997	Qiu Xia	2004-2005
Sharon Lewis	1992-1995	Ken Roche OAK RIDGE NATIONAL LABORATORY	1999-2004	Tinghua Xu	1998-2000
Xiang Li	2001	Tom Rothrock	1997-1998	Tao Yang	1999
Weiran Li	2002	Tom Rowan	1993-1997	Yuanlei Zhang MICROSOFT	2001-2005
Chaoyang Liu	2000	Kiran Sagi MICROSOFT	2001-2005	Yong Zheng	2001
Kevin London MICROSOFT	1996-2005	Evelyn Sams	1998-1999	Luke Zhou MICROSOFT	2000-2001
Matt Longley	1999	Farial Shahnaz	2001	Min Zhou	2002-2004
Richard Luczak	2000-2001	Zhiao Shi VANDERBILT UNIVERSITY	2001-2007		
Piotr Luszczek THE MATHWORKS	1999-2007	Sergei Shinkarev AMAZON.COM	2005-2007		

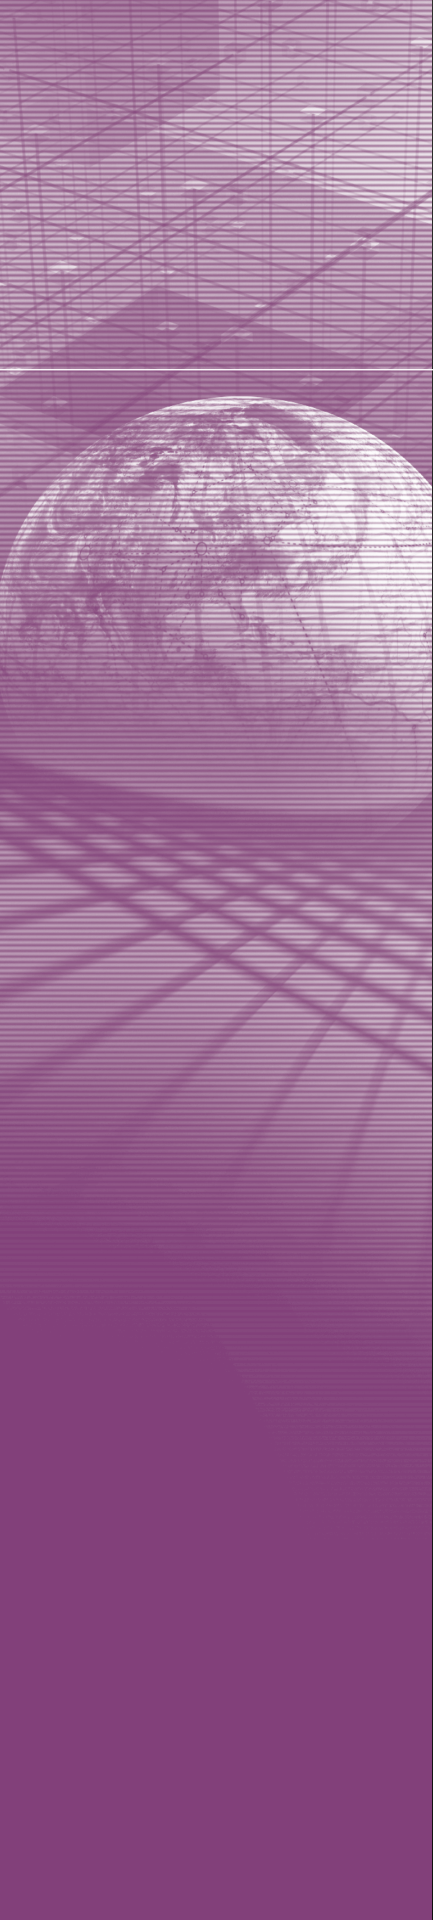

PARTNERSHIPS

Over the past 18 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.

Leveraging the incredible number of collaborators that we have worked with over the years, we have built a strong portfolio of shared resources, both material and intellectual. In addition, we have forged numerous lasting relationships with application and tool vendors, many of who 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. Below are lists that highlight 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

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
Emory University	
Google	
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
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
NIST	National Institute of Standards and Technology
NPACI	The National Partnership for Advanced Computational Infrastructure
NSF	The National Science Foundation
ORNL CSMD	Oak Ridge National Laboratory Computer Science and Mathematics Division
Rice University	
SDSC	San Diego Supercomputing Center
SGI	Silicon Graphics Incorporated
Sun Microsystems	
TACC	Texas Advanced Computing Center
University of California, Berkeley	
University of California, San Diego	
University of California, Santa Barbara	
University of Illinois, Urbana-Champaign	
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
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

The heterogeneous nature of high performance computing requires inexhaustible knowledge about the hardware, software, and interconnects that comprise such machines. Thus, such a myriad of architectures 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 Intel EM64T cluster connected with Myrinet 2000	64 node AMD Opteron cluster connected with Myrinet 2000
24 node AMD Opteron cluster connected with Silverstorm, Mellanox, and Myricom	16 node Intel Xeon connected with Gigabit Ethernet
8 node Intel Core2 Duo cluster with Gigabit Ethernet	Intel Itanium clusters

In addition to these resources, we have access to several server class machines and several HPC clusters in other departments. 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.

In addition, exclusive access to many remote resources, some that are regularly found in the Top500 list of the worldís 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 remote systems and architectures that we utilize:

Cray X1, X1E, XT3, and XT4	HP XC system
IBM Power 4, 4+, 5, 5+, 6, Cluster 1600, BlueGene/L, and the Cell	Several large (256+ proc) Linux Clusters
SGI Altix	

RECENT PUBLICATIONS

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

Angskun, T., Bosilca, G., Dongarra, J. "**Self-Healing in Binomial Graph Networks**," *2nd International Workshop On Reliability in Decentralized Distributed Systems (RDDS 2007)*, Vilamoura, Algarve, Portugal, November, 2007.

Angskun, T., Bosilca, G., Dongarra, J. "**Binomial Graph: A Scalable and Fault- Tolerant Logical Network Topology**," *Proceedings of The Fifth International Symposium on Parallel and Distributed Processing and Applications (ISPA07)*, Niagara Falls, Canada, Springer, August 29-30, 2007.

Angskun, T., Bosilca, G., Fagg, G., Pjesivac-Grbovic, J., Dongarra, J. "**Reliability Analysis of Self-Healing Network using Discrete-Event Simulation**," *Proceedings of Seventh IEEE International Symposium on Cluster Computing and the Grid (CCGrid '07)*, IEEE Computer Society, 437-444, May, 2007.

Angskun, T., Bosilca, G., Vander Zanden, B., Dongarra, J. "**Optimal Routing in Binomial Graph Networks**," *The International Conference on Parallel and Distributed Computing, applications and Technologies (PDCAT)*, Adelaide, Australia, IEEE Computer Society, December 3-6, 2007.

Baboulin, M., Dongarra, J., Gratton, S., Langou J. "**Computing the Conditioning of the Components of a Linear Least Squares Solution**," ICL Technical Report, UT-CS-07-604, (also LAPACK Working Note 193), September, 2007.

Bouteiller, A., Bosilca, G., Dongarra, J. "**Retrospect: Deterministic Relay of MPI Applications for Interactive Distributed Debugging**," *Recent Advances in Parallel Virtual Machine and Message Passing Interface, 14th European PVM/MPI Users' Group Meeting*, Paris France, 2007, pp 297-306, LNCS, Volume 4757, Springer Berlin / Heidelberg, 2007.

Buttari, A., Dongarra, J., Husbands, P., Kurzak, J., Yelick, K. "**Multithreading for synchronization tolerance in matrix factorization**," *Journal of Physics: Conference Series*, SciDAC 2007, Volume 78, 2007, June, 2007.

Buttari, A., Dongarra, J., Kurzak, J. "**Limitations of the Playstation 3 for High Performance Cluster Computing**," *University of Tennessee Computer Science Technical Report*, *UT-CS-07-597 (Also LAPACK Working Note 185)*, 2007.

Buttari, A., Dongarra, J., Kurzak, J., Langou, J., Langou, Jn., Luszczek, P., Tomov, S. "**Exploiting Mixed Precision Floating Point Hardware in Scientific Computations**," In *High Performance Computing and Grids in Action* (to appear), Lucio Grandinetti eds. Amsterdam, IOS Press, 2007.

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 TOMS* (to appear December 2008).

Buttari, A., Dongarra, J., Langou, J., Langou, Jn., Luszczek, P., Kurzak, J. "**Mixed Precision Iterative Refinement Techniques for the Solution of Dense Linear Systems**," *International Journal of High Performance Computer Applications* (to appear), August, 2007.

Buttari, A., Langou, J., Kurzak, J., Dongarra, J. "**Parallel Tiled QR Factorization for Multicore Architectures**," *Concurrency and Computation: Practice and Experience* (submitted), July 2007.

Buttari, A., Langou, J., Kurzak, J., Dongarra, J. "**Parallel Tiled QR Factorization for Multicore Architectures**," *Seventh International Conference on Parallel Processing and Applied Mathematics, PPAM 2007*, Gdansk, Poland, September 9-12, 2007.

Buttari, A., Langou, J., Kurzak, J., Dongarra, J. "**A Class of Parallel Tiled Linear Algebra Algorithms for Multicore Architectures**," *Parallel Computing* (submitted), October 2007.

Buttari, A., Luszczek, P., Kurzak, J., Dongarra, J., Bosilca, G. "**SCOP3: A Rough Guide to Scientific Computing On the PlayStation 3**," *University of Tennessee Computer Science Dept. Technical Report, UT-CS-07-595*, April 17, 2007.

Chen, Z., Yang, M., Francia, G., Dongarra, J. "**Self Adapting Application Level Fault Tolerance for Parallel and Distributed Computing**," *Proceedings of Workshop on Self Adapting Application Level Fault Tolerance for Parallel and Distributed Computing at IPDPS*, pp. 1-8, March 26-30, 2007.

Dongarra, J. "**Performance of Various Computers Using Standard Linear Equations Software (Linpack Benchmark Report)**," *University of Tennessee Computer Science Dept. Technical Report CS-89-85*, 2007.

Dongarra, J., Luszczek, P. "**High Performance Development for High End Computing with Python Language Wrapper (PLW)**," *International Journal for High Performance Computer Applications*, Fall 2007, Volume 21, No. 3, pp 360-369.

Dongarra, J., Chen, Z., Bosilca, G., Langou, J. "**Disaster Survival Guide in Petascale Computing: An Algorithmic Approach**," in *Petascale Computing: Algorithms and Applications*, David Bader, Ed., Chapman & Hall / CRC Press, pp 263-288, 2007.

Dongarra, J., et al., "**Remembering Ken Kennedy**," *SciDAC Review*, Issue 5, Fall 2007, ISSN 1935-0570, Fall 2007, 2007.

Dongarra, J., Golub, G., Moler, C., Moore, K. "**Netlib and NA-Net: building a scientific computing community**," In *IEEE Annals of the History of Computing* (to appear), August, 2007.

Dongarra, J., Jeannot, E., Saule, E., Shi, Z. "**Bi-objective Scheduling Algorithms for Optimizing Makespan and Reliability on Heterogeneous Systems**," *19th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)* (submitted), San Diego, CA, June 9 - 11, 2007.

Dongarra, J., Luszczek, P. "**How Elegant Code Evolves With Hardware: The Case Of Gaussian Elimination**," in *Beautiful Code Leading Programmers Explain How They Think*, Andy Oram and Greg Wilson, eds. Chapter 14, pp 229-252, June 2007.

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

Dongarra, J., Pineau, J-P., Robert, Y., Shi, Z., Vivien, F. "**Revisiting Matrix Product on Master-Worker Platforms**," *International Journal of Foundations of Computer Science (IJFCS)* (to appear), 2007.

Eijkhout, V. "**Numerical Metadata API Reference**," *Innovative Computing Laboratory Technical Report*, February 2nd, 2007.

Fuerlinger, K., Dongarra J. "**On Using Incremental Profiling for the Performance Analysis of Shared Memory Parallel Applications**," *Proceedings of the 13th International Euro-Par Conference on Parallel Processing (Euro-Par '07)*, Rennes, France, Springer LNCS, August, 2007.

Fuerlinger, K., Gerndt M., Dongarra J. "**Scalability Analysis of the SPEC OpenMP Benchmarks on Large-Scale Shared Memory Multiprocessors**," *Proceedings of the 2007 International Conference on Computational Science (ICCS 2007)*, Yong Shi, Geert van Albada, Jack Dongarra, Peter M.A. Sloot eds. Springer LNCS, Volumes 4487-4490, 815-822", May 27-30, 2007.

Fuerlinger, K., Moore, S. "**Continuous Runtime Profiling of OpenMP Applications**," *Proceedings of the 2007 Conference on Parallel Computing (PARCO 2007)*, Juelich and Aachen, Germany, September, 2007.

Gerndt, M., Fuerlinger, K. "**Specification and detection of performance problems with ASL**," *Concurrency and Computation: Practice and Experience*, John Wiley and Sons Ltd., Vol. 19, no 11, 1451-1464, 2007.

Graham, R., Bosilca, G., Pjesivac-Grbovic, J. "**A Comparison of Application Performance Using Open MPI and Cray MPI**," *Cray User Group, CUG 2007*, May, 2007.

Graham, R., Brightwell, R., Barrett, B., Bosilca, G., Pjesivac-Grbovic, J. "**An Evaluation of Open MPI's Matching Transport Layer on the Cray XT**," *EuroPVM/MPI 2007*, September, 2007.

Jeannot, E., Seymour, K., YarKhan, A., Dongarra, J. "**Improved Runtime and Transfer Time Prediction Mechanisms in a Network Enabled Servers Middleware**," *Parallel Processing Letters*, Volume 17, Number 1, pp. 47-59, March, 2007.

Kurzak, J., Buttari, B., Dongarra, J. "**Solving Systems of Linear Equations on the CELL Processor Using Cholesky Factorization**," *UT Computer Science Technical Report (Also LAPACK Working Note 184)*, UT-CS-07-596, May, 2007.

Kurzak, J., Dongarra, J. " **Implementation of Mixed Precision in Solving Systems of Linear Equations on the Cell Processor**," *Concurrency and Computation: Practice and Experience*, Volume 19, Issue 10, pp 1371-1385, July, 2007.

Langou, J., Chen, Z., Bosilca, G., Dongarra, J., "**Recovery Patterns for Iterative Methods in a Parallel Unstable Environment**," *SIAM SISC* (to appear), May, 2007.

Mellor-Crummey, J., Beckman, P., Dongarra, J., Miller, B., Yelick, K. "**Creating Software Technology to Harness the Power of Leadership-class Computing Systems**," *DOE SciDAC Review*, Issue 5, Fall 2007 (Unreviewed).

Mellor-Crummey, J., Beckman, P., Cooper, K., Dongarra, J., Gropp, W., Lusk, E., Miller, B., Yelick, K. "**Creating Software Tools and Libraries for Leadership Computing**," *CTWatch Quarterly*, Vol. 3, Number 4, November 2007.

Pineau, J.-F., Robert, Y., Vivien, F., Shi, Z., Dongarra, J., "**Revisiting Matrix Product on Master-Worker Platforms**," *9th Workshop on Advances in Parallel and Distributed Computational Models in conjunction with 21st IEEE International Parallel and Distributed Processing Symposium*,Long Beach, California. March 26th, 2007.

Pjesivac--Grbovic, J., Bosilca, G., Fagg, G., Angskun, T., Dongarra, J. "**MPI Collective Algorithm Selection and Quadtree Encoding**," *Parallel Computing,* Volume 33, Number 9, pp 613-623.

Pjesivac-Grbovic, J., Angskun, T., Bosilca, G., Fagg, G., Gabriel, E., Dongarra, J. "**Performance Analysis of MPI Collective Operations**," *Cluster computing*, Springer Netherlands, Volume 10, Number 2, 127-143, June, 2007.

Pjesivac-Grbovic, J., Bosilca, G., Fagg, G., Angskun, T., Dongarra, J. "**Decision Trees and MPI Collective Algorithm Selection Problem**," *Euro-Par 2007*, Rennes, France, Springer, 105--115, August, 2007.

Shende, S., Malony. S., Moore, S., and Cronk, D. "**Memory Leak Detection in Fortran Applications using TAU**," *Proc. DoD HPMCP Users Group COnference (HPMCP-UGC'07)*, Pittsburgh, PA, IEEE Computer Society, June, 2007.

Song, F., Moore, S., Dongarra, J. " **Feedback-Directed Thread Scheduling with Memory Considerations**," *IEEE International Symposium on High Performance Distributed Computing*, Monterey Bay, CA, June 27-29, 2007.

Song, F., Moore, S., Dongarra, J. "**L2 Cache Modeling for Scientific Applications on Chip Multi-Processors**," *Proceedings of the 2007 International Conference on Parallel Processing*, Xi'an, China, IEEE Computer Society, September 10-14, 2007.

Supinski, B., Hollingworth, J., Moore, S., Worley, P. "**Results of the PERI survey of SciDAC applications,**" *Journal of Physics: Conference Series*, SciDAC 2007, Volume 78, 2007, June, 2007.

Vomel, 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*, Volume 223, pp. 774-782, 2007.

Vomel, C., Tomov, S. Z., Marques, O. A., 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* (submitted), May 2007.

Wolf, F., Mohr, B., Dongarra, J., Moore, S. "**Automatic Analysis of Inefficiency Patterns in Parallel Applications**," *Concurrency and Computation: Practice and Experience,* Volume 19, Issue 11, pp 1481-1496, August 2007.

YarKhan, A., Dongarra, J., Seymour, K. "**GridSolve: The Evolution of Network Enabled Solver**," *Grid-Based Problem Solving Environments: IFIP TC2/WG 2.5 Working Conference on Grid-Based Problem Solving Environments* (Prescott, AZ, July 2006), Patrick Gaffney, James C. T. Pool eds. Springer, 215-226, 2007.

You, H., Seymour, K., Dongarra, J., Moore, S. "**Automated Empirical Tuning of a Multiresolution Analysis Kernel**," *ICL Technical Report, ICL-UT-07-01*, 10, February 13, 2007.

You, H., Seymour, K., Dongarra, J., Moore, S. "**Empirical Tuning of a Multiresolution Analysis Kernel using a Specialized Code Generator**," *ICL Technical Report, ICL-UT-07-02*, March 5, 2007.

Cyberinfrastructure Technology Watch

Since 2004, ICL has a led a broad ranging publication effort called Cyberinfrastructure Technology Watch (CTWatch) 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). The goal of CTWatch has been to serve both as a forum for ideas and opinion on issues of importance to the cyberinfrastructure community as well 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 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, we created *CTWatch Quarterly*, an on-line serial publication modeled on a more traditional academic journal. Along a more experimental line, we 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.

FEBRUARY 2007

MAY 2007

AUGUST 2007

NOVEMBER 2007

CTWatch Quarterly is designed to be published on-line and is made available in both HTML and in a high quality PDF 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 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).

2007 marks the third and final year of the CTWatch effort. Over the past three years, 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 collaborators at NCSA and SDSC, the guest editors, the article authors, the editorial board, and the entire production staff for successfully meeting the original goals of the *CTWatch Quarterly* and for playing an enormous role in increasing cyberinfrastructure visibility within academia, government, and industry.

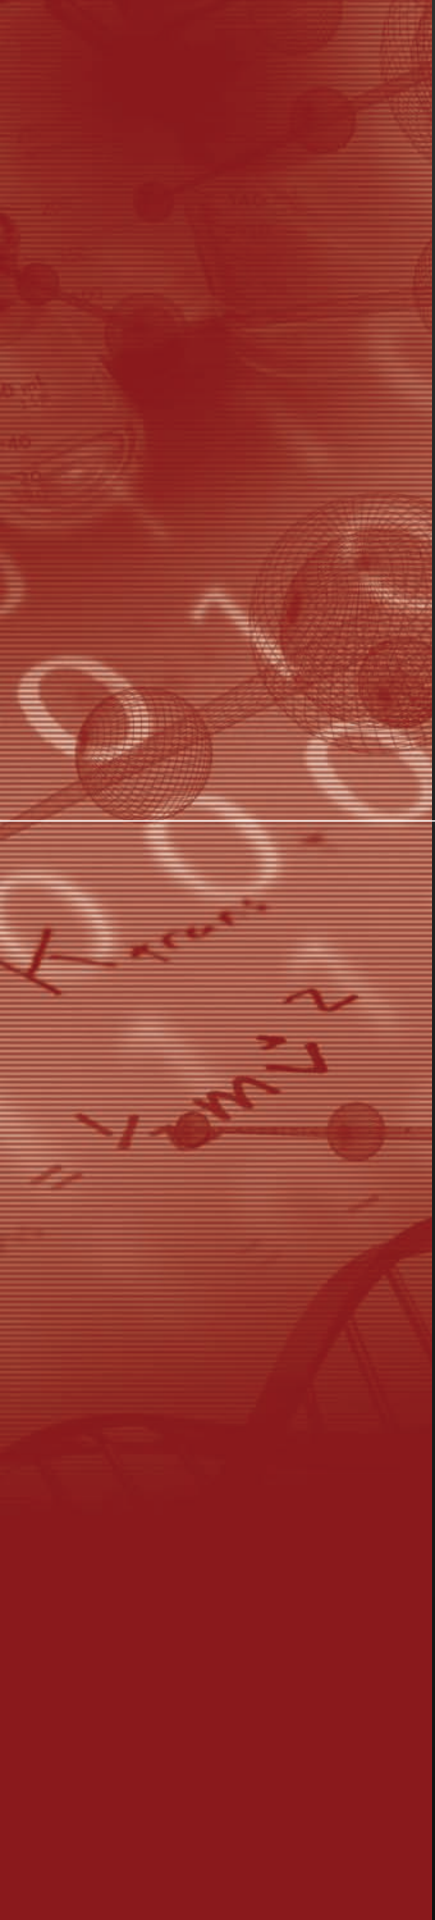

CITR

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.

COMPUTATIONAL SCIENCE: INHERENTLY INTERDISCIPLINARY

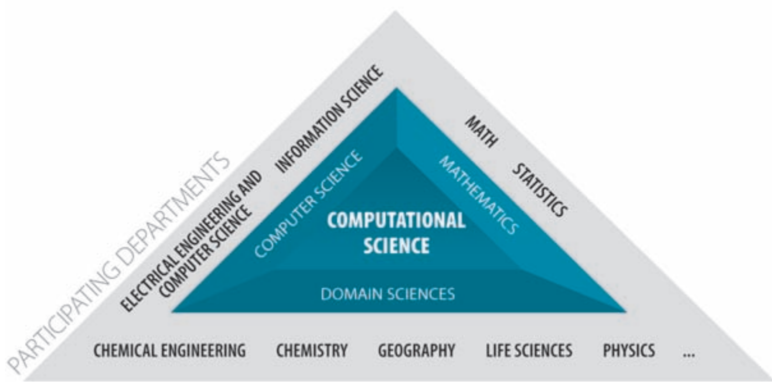

IGMCS

Since first rate students and staff are indispensable to the success of such a research 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. In 2007, CITR worked 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 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.

With enormous progress in IT capabilities during this first decade of the 21st century, Computational Science will continue to emerge as an essential component to academic 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 more academic disciplines begin to realize the incredible potential Computational Science provides, the IGMCS program is expected to grow by adding new disciplines, new courses, and new faculty. At the beginning of Fall semester in 2007, there were 12 different departments contributing more than 75 courses to the program representing three of the largest colleges on the UT campus: Arts and Sciences; Communication and Information; and Engineering.

For more information about the IGMCS program, visit http://igmcs.utk.edu/.

Acknowledgments

Since ICL's inception, we have earned the trust and support of the 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 investing the nationís computational research funding. Therefore we acknowledge the following for their support of our efforts past and present:

NSF DOE DoD DARPA ONR NIH

In addition to the support of the federal government, we have solicited strong support from private industry, which has also played a significant role in our success and growth. Some organizations have targeted specific ICL projects, while others have made contributions to our work that are more general and open-ended. We gratefully acknowledge the following for their generosity and their significance to our success:

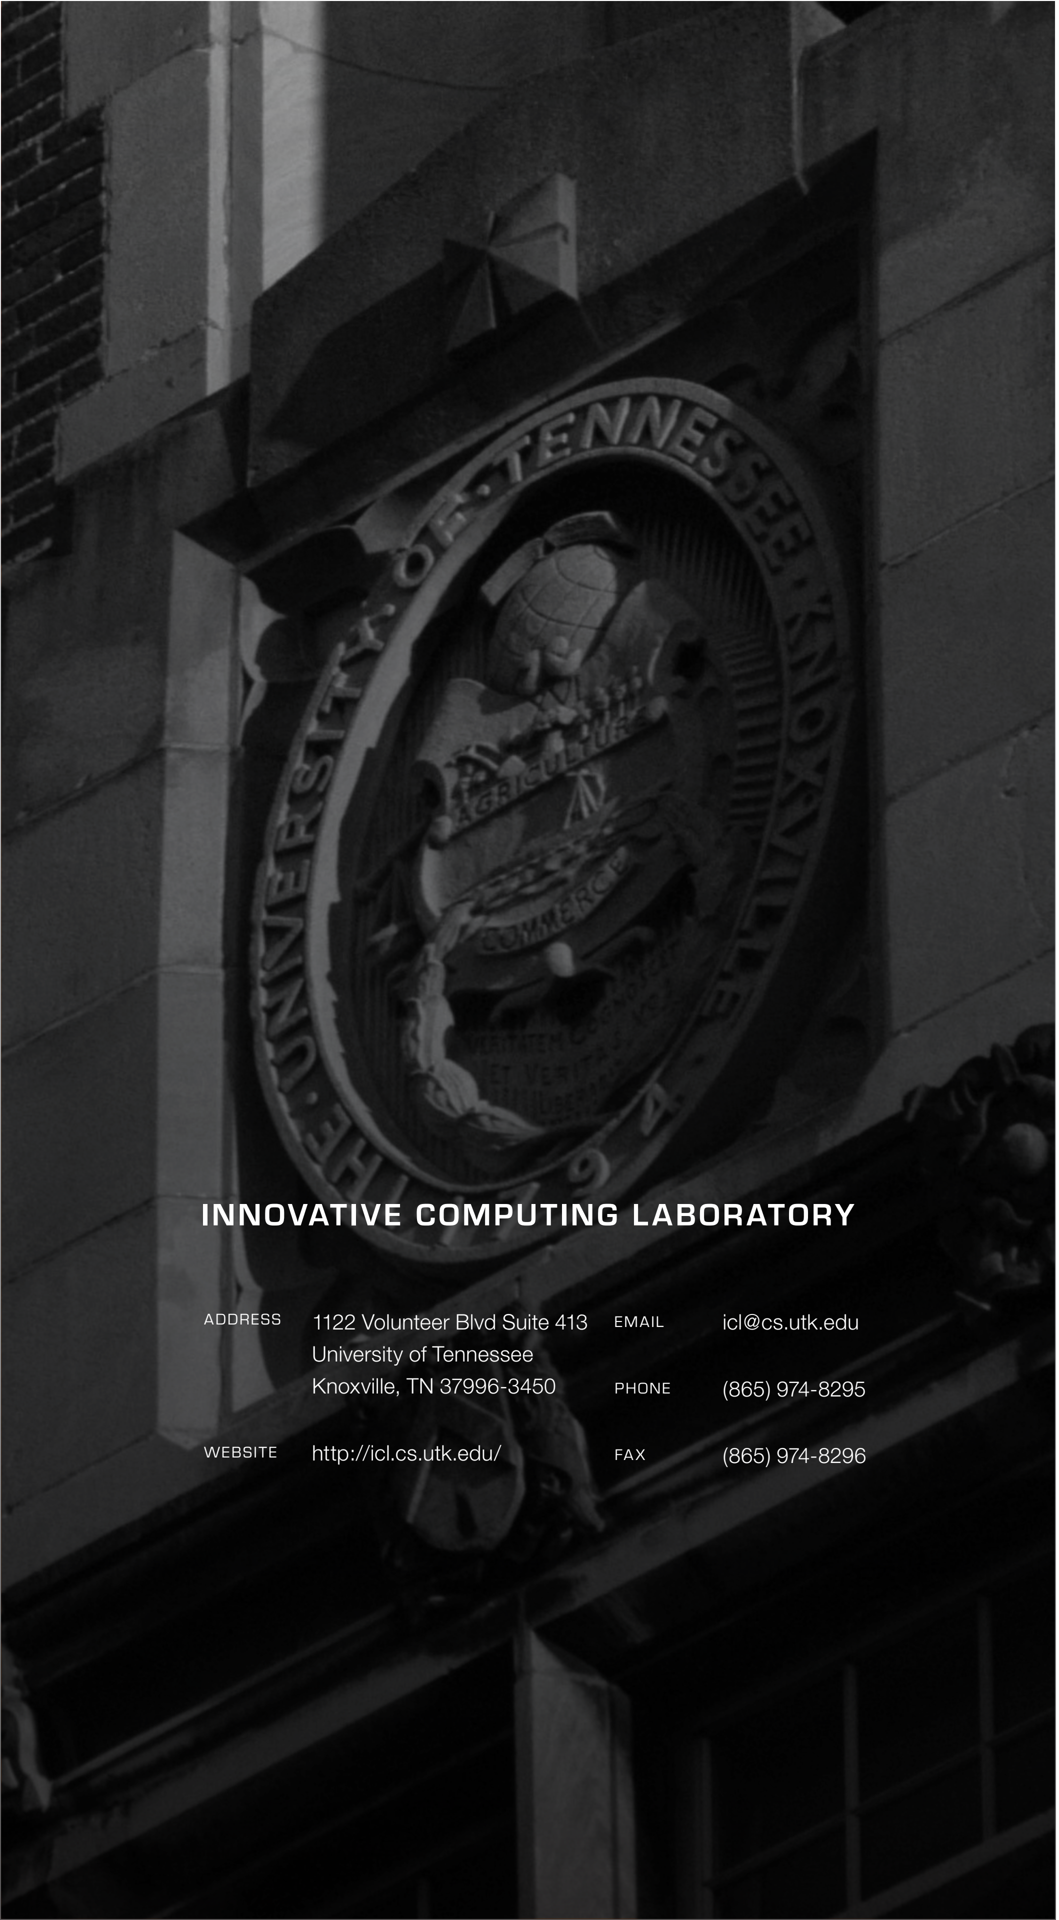

INNOVATIVE COMPUTING LABORATORY

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

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