INNOVATIVE COMPUTING LABORATORY

2002 REPORT

COMPUTER SCIENCE DEPARTMENT
UNIVERSITY OF TENNESSEE

INNOVATIVE COMPUTING LABORATORY 2002 REPORT

Edited by Scott Wells

Designed by David Rogers

The University of Tennessee - Knoxville
EEO/Title IX/Section 504/ADA Statement

The University of Tennessee does not discriminate on the basis of race, sex, color, religion, national origin, handicap, or veteran status in provision of educational opportunities and benefits. This policy extends both to employment by and admission to the University.

UT does not discriminate on the basis of race, sex, or disability in the education programs and activities which it operates, pursuant to requirements of Title VI of the Civil Rights Act of 1964, Title IX of the Educational Amendments of 1972, section 504 of the Rehabilitation Act of 1973, and the Americans with Disabilities Act (ADA) of 1990.

Inquiries and charges of violation concerning Title VI, Title IX, Section 504, ADA or the Age Discrimination in Employment Ace (ADEA), or any of the other above referenced policies should be directed to the Office of Diversity Resources and Educational Services (DRES), 1818 Lake Avenue, Knoxville, TN 37996-3560, telephone (865) 974-2498 (TTY available). Requests for accommodation of a disability should be directed to the ADA Coordinator at the Office of Human Resources Management, 600 Henley Street, Knoxville, TN 37996-4125.

2002 REPORT

CONTENTS

INTRODUCTION

MISSION STATEMENT

The guiding mission of the Innovative Computing Laboratory (ICL) is to push back the frontiers of discovery in high performance and distributed computing in the 21st Century, and to infuse leading edge scientific and general-purpose applications with the innovations flowing from this research. The pace of progress of scientific inquiry across a vast spectrum of disciplines is now intimately linked with advances in scientific computing and with the creation of new software tools that puts computing power into the hands of domain scientists and engineers. ICL, along with the richly talented group of people who have made it successful, stands committed to leadership in this new era of scientific simulation in which the cooperative efforts of computer and computational scientists propel scientific research to new and unparalleled knowledge of the world around us.

ICL aspires to be a world leader in enabling technologies and software for scientific computing. Our vision is to provide 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.

BACKGROUND

ICL was established in the fall of 1989 when Dr. Jack Dongarra came to the University of Tennessee (UT) from Argonne National Laboratory (ANL). Dr. Dongarra was given a dual appointment as Distinguished Professor in the Computer Science Department at the university and as Distinguished Scientist at Oak Ridge National Laboratory (ORNL). This dual position was established by the UT/ORNL Science Alliance, Tennessee's oldest and largest Center of Excellence, as a means for attracting top research scientists from around the country and the world to visit the university and collaborate. Subsequently, many post-doctoral researchers and professors from various research backgrounds such as mathematics, geology, chemistry, etc. visited the university. Many of these scientists have passed through UT as post-doctoral researchers and worked with Dr. Dongarra. In addition, these scientists were vital in helping him attract additional researchers as well as top graduate students. Below is a list of some of the original researchers who were instrumental in helping Dr. Dongarra with the establishment and growth of ICL.

- Zhaojun Bai of the University of California Davis
- Adam Beguelin of the Inktomi Corporation
- Susan Blackford of Myricom
- Jaeyoung Choi of Soongsil University (Korea)
- Andy Cleary of Data Digest
- Frederic Desprez of ÉNS Lyon
- Robert van de Geijn of the University of Texas Austin
- Antoine Petitet of Sun Microsystems France
- Roldan Pozo of NIST
- Françoise Tisseur of Manchester University
- Bernard Tourancheau of Sun Microsystems Lab
- Clint Whaley of Florida State University

Through interactions with colleagues at Rice University, ICL became an integral part of the Center for Research on Parallel Computation (CRPC), a National Science Foundation (NSF) Science and Technology Center established in 1989 and lead by Rice University. CRPC worked to make parallel computation accessible to industry, government, and academia and to educate a new generation of technical professionals. Through the 1990s, ICL worked on a number of efforts that have since become part of the basic fabric of scientific computing around the world. The basic technologies that our research has produced include the ATLAS, BLAS, LAPACK, ScaLAPACK, PVM, MPI, Netlib, and the NHSE. These successes are continuing along with current ICL efforts, such as Active Netlib, PAPI, HARNESS, NetSolve, SANS-Effort, and the TOP500. In the past eight years, four of our projects have earned R&D 100 awards; PVM in 1994, ATLAS and NetSolve in 1999, and PAPI in 2001.

Claxton Building - Home of ICL

Having linear algebra as its original focus, the group has evolved and expanded to encompass many progressive areas of high performance computing research such as distributed computing, software repositories, and performance evaluation. Currently a University Distinguished Professor, Dr. Dongarra continues as director of ICL. As such, he not only serves as principal investigator (PI) for many of our projects, but he also maintains a level of participation in all projects.

PROFILE

Located at the heart of the UT campus in Knoxville, ICL is an internationally recognized academic leader in high performance computing (HPC) research. We have recently moved into new facilities at the university, due in large part to our incredible growth over the last few years. Located in the recently built Claxton Complex, ICL and the UT Computer Science Department occupy nearly ¾ths of the approximately 70,000 square foot facility.

Operating under grants totaling nearly $5 million annually, ICL is also recognized by senior UT administration as one of UT's top three research centers. According to Dr. T. Dwayne McCay, Vice President of the UT Office for Research and Information Technology, "Jack Dongarra and the students and staff of the Innovative Computing Laboratory have been leading our university and our nation in high performance computing and information technology research for more than a decade now. Looking out at the decade to come, we're really excited by the prospect of what their creativity and their new discoveries will bring us. The work of Dr. Dongarra's team shows why we are so determined to be a top 25 research university. The efficient solution to modern scientific and technological problems requires the most modern tools and the best trained minds, and both are attracted to the great research universities by the opportunity to work with someone like Jack Dongarra. The rich opportunities presented by such environments provide the fuel that drives our economy and raise the standard of living for all our citizens."

DIRECTOR'S STATEMENT

In 2002, the Innovative Computing Laboratory is celebrating 13 years of leadership in enabling technologies for high performance computing. Looking back over the 13-year period, the evolution and growth of the technology for computing has been truly astonishing. In an environment where technology changes every 18 months, ICL cannot afford to stand still. In 1989 the speed of a supercomputer was measured in gigaflops and in gigabytes. Today our measures are teraflops for speed and terabytes for memory, a thousand-fold increase over the standards of a decade ago. The research that ICL has undertaken in the past decade has followed a natural progression and growth from our original tread of numerical linear algebra to performance evaluation, to software repositories, and to distributed computing.

ICL Director
Dr. Jack Dongarra

The Innovative Computing Laboratory 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.

We are building from a firm foundation. Over the past 13 years, we have developed robust research projects, attracted some of the best and brightest students and researchers, and created leading-edge research programs. The ICL staff's ongoing ability to apply the latest technologies to provide advanced services and solutions for the scientific computing community underscores our leadership role. Standards and efforts such as PVM, MPI, LAPACK, ScaLAPACK, BLAS, ATLAS, Netlib, NHSE, TOP500, and the LINPACK Benchmark have all left their mark on the scientific community. We can be proud of the recognition and use our tools receive. We are continuing these efforts with PAPI, NetSolve, RIB, the TOP100 Clusters, HARNESS, Active Netlib, and the Self-Adapting Numerical Software (SANS) Effort, as well as other innovative computing projects.

We continue to grow in terms of the resources we have at our disposal. We have ongoing efforts to strengthen our organization and to ensure the proper balance and integration of research and projects. The pace of change will continue to accelerate in the coming years.

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 ICL one of the best centers for enabling technology in the world.

-Jack Dongarra

ICL RESEARCH

As ICL has grown over the years, the range and diversity of the research performed by our staff and students has increased in parallel. In the past year alone, we supported or participated in more than 20 significant projects. Our large and wide-ranging portfolio of research projects has evolved over the course of more than a decade, beginning from a narrow but solid foundation.

The original focus of ICL was Dr. Dongarra's work in numerical linear algebra and the numerical libraries that encode its operations in software. But driven by the relentless demand for enabling technology in the computational science community, ICL built upon its successes in the area of numerical libraries and the growing strength of its personnel to break new ground in the areas of high performance and distributed computing. Similarly, our work with numerical libraries created a strong area of expertise in performance evaluation and benchmarking for high-end computers. The enormous investments by both government and private industry in high performance computing have made our ability to do research in this area correspondingly important. Finally, as a by-product of a long tradition of delivering high quality software produced from our research, we have helped to lead the movement to build robust, comprehensive, and well-organized software repositories.

With the phenomenal growth over the last several years in parallel computing technology and the demands placed on such technology by government and private business, we are consistently challenged to apply expert-level understanding to each of our research efforts. The areas of distributed and network computing are no exception as we have learned to harness enormous computing power to quickly and efficiently solve mathematical problems that would take humans years or decades to solve by hand.

ACKNOWLEDGEMENT

In the next few pages, the focus of our enabling technology research - numerical libraries, high performance distributed computing, performance evaluation and benchmarking, and software repositories - is described in more detail. Evidence of the perceived value of our work and the importance of our research is apparent in the range of agencies and organizations that have funded, and continue to fund, our efforts. Our main source of support has been federal agencies that are charged with investing the nation's research funding: the National Science Foundation (NSF), the Department of Energy (DOE), the Department of Defense (DoD), the Defense Advanced Research Projects Agency (DARPA), and the National Aeronautics and Space Administration (NASA). However, 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:

AMD	Hewlett-Packard	Myricom
Compaq	IBM	Scali
Dolphin Interconnect Solutions	Intel	Sun Microsystems
	MathWorks	
	Microsoft	

NUMERICAL LINEAR ALGEBRA

ICL has long been a leader in the area of numerical linear algebra algorithms and software for high performance computers. Linear algebra operations form the core of an overwhelming number of scientific applications. Having efficient algorithms and implementations for these operations is of utmost importance in achieving good performance for these applications. In collaboration with other researchers and with industry, ICL has led successful efforts to standardize library interfaces for the Basic Linear Algebra Subroutines (BLAS) as well as for higher level dense linear algebra routines such as those for solving linear systems. Our groundbreaking research in efficient algorithms and implementations for these routines has been widely adopted and refined by industry to produce high efficient linear algebra implementations for most architectures, with the result that applications that rely on the standard libraries can achieve superior performance while remaining portable across multiple platforms. Dense linear algebra software produced by our researchers have included LINPACK portable software, LAPACK for shared memory multiprocessors, and ScaLAPACK for distributed memory multiprocessors. To address the problem of efficiently generating core software for new architectures, we have developed the Automatically Tuned Linear Algebra Software (ATLAS) system for the automatic generation and optimization of linear algebra software. ATLAS has been widely adopted by vendors to produce efficient BLAS for their machines in a fraction of the time required for hand coding. Recent ICL linear algebra research has focused on sparse linear algebra in the areas of iterative methods and parallel preconditioning and on smart libraries for partial automation of the choice of method and preconditioners.

To keep pace with the shift in the milieu for scientific computing from desktop computers and dedicated parallel machines to clusters of workstations and distributed computational grids, our research has begun targeting mathematical software towards these new environments. Java has become widely accepted as a portable language for multi-threaded Internet-based distributed applications. Java source code is compiled into machine independent byte code, which is then interpreted by a Java Virtual Machine (JVM). It is a common and somewhat erroneous belief that Java will always be too slow for scientific computing. The ICL Fortran to Java (F2J) project is addressing the question of the feasibility of scientific computing via Java through development of a translator that converts programs written in a subset of Fortran 77 into a form that can be executed on Java Virtual Machines (JVMs). This translator makes it possible for a Java application or applet to use established legacy numerical code that was originally written in Fortran. Such access to efficient numerical routines, in addition to use of just-in-time (JIT) compilers that translate Java byte code to native code a run-time and thus increase execution speed, have the potential of making Java a viable programming language for distributed scientific computing.

Also targeting new computational environments, ICL is leading the collaborative Self Adaptive Numerical Software (SANS) effort which seeks to optimize numerical software at different levels in clustered and distributed execution environments. SANS systems are intended to help manage the complex grid computing environment by delivering to the scientist the full power of flexible compositions of the available algorithmic alternatives. SANS aims to provide a dynamic computational environment in which the most effective library components are automatically selected based on problem characteristics, attributes of input data, and the state of the computational grid. A SANS system accepts user data and performs numerical and structural analysis on it to determine the algorithms and data structures that are feasible and to predict the expected performance of the various alternatives. For a non-expert user, a SANS system functions as an expert system, fully taking on the burden of searching for an optimal solution. For users able to supply information about input data and algorithmic choices, a SANS system serves as a testbed to aid in searching for the most efficient solution and methods.

An example of a SANS system is the ICL LAPACK for Clusters (LFC) project that merges the ease of use of LAPACK with the parallel processing capabilities of ScaLAPACK. LFC is a self-contained package with built-in knowledge of how to run linear algebra software on a cluster. Users are responsible for stating their numerical problems but can assume they are working in a serial environment. The LFC middleware assesses the possibility of solving the problem faster in parallel on some subset of the available resources, based on information describing the state of the system. If parallel execution is expected to be faster, the user's data are distributed over the selected processors, the problem is executed in parallel, and the solution is returned to the user. Experimental results compare favorably with the performance obtained by an expert user in the same environment.

Our SANS effort also includes development of the Self-Adapting Linear Solver Architecture (SALSA), which acts as an intermediary between the user and the computational environment for solving linear systems. It utilizes an 'intelligent agent' component that heuristically investigates properties of the user data that influence the decision making process of finding the best solver. User interaction comprises both the possibility of guiding the computation by user input and educating the user on the numerics of the input problem through feedback of the analysis and performance results. Over time, the intelligent agent builds up a database of performance data that allows tuning of the heuristics.

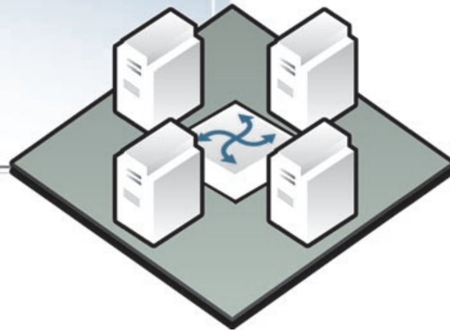

DISTRIBUTED COMPUTING

Distributed computing has been a core area of research at ICL for more than a decade. An early and highly successful example of this work is the Parallel Virtual Machine (PVM) system, developed in collaboration with researchers at ORNL, which enables scientific applications to run in parallel across networks of workstations. Today, our portfolio of distributed computing projects continues to build on this strong tradition, especially in the area of Grid computing. The most prominent effort focuses on NetSolve. NetSolve is Grid middleware that combines the power of hardware and software aggregation, which is characteristic of Grid computing, with the ease of use of general purpose Scientific Computing Environments (SCEs) (e.g. Matlab and Mathematica), which are essential tools for the majority of working computational scientists and engineers. Built around a form of the classic remote procedure call (RPC) mechanism with enhanced intelligence to enable it to deal with the complexities of Grid environments, NetSolve allows domain scientists to discover, remotely access, and utilize a huge variety software modules (e.g. advanced numerical libraries, large software simulators, non-portable codes) that would otherwise be unavailable via Matlab or Mathematica, and to access and aggregate hardware resources needed to run them. With a solid base of production level software already in place, NetSolve's user community is growing quickly.

Several complementary projects are also under way. For example, ICL is working with the Global Grid Forum to develop GridRPC, which will be a community standard RPC mechanism for Grid environments; NetSolve's code base provides the foundation for that effort. And both NetSolve and GridRPC are encompassed by GridSolve, which is a new project sponsored by the NSF Middleware Initiative, that aims to harden and deploy the positive results of Grid computing research for the benefit of the science and engineering community, generally. GridSolve is building on the NetSolve code base in order to combine NetSolve's ease of use for computational scientists with a variety of de facto standard Grid technologies, including Globus and Condor compute services, the Network Weather Service (NWS) for performance forecasting, storage services via the Internet Backplane Protocol (IBP), and the Grid Security Infrastructure for authentication and access control.

Another project, called HARNESS (Heterogeneous Adaptable Reconfigurable Networked SystemS), is an experimental Grid computing framework that leverages the success of PVM. Siminar to PVM, HARNESS is built around the concept of a "virtual machine" (VM) that allows a set of

DISTRIBUTED COMPUTING

NetSolve · GridSolve/GridRPC · HARNESS · FT-MPI · GrADS · SInRG

diverse computing resources to be viewed as a single, large, distributed memory computing resource. But HARNESS seeks to remove the limitations discovered in PVM's monolithic architecture in order to create a completely different approach to building, modifying, and using multiple VMs that can facilitate a new generation of collaborative and computation environments. Developed in collaboration with researchers at ORNL and Emory University, HARNESS builds on the services of modular, highly customizable and reconfigurable distributed virtual machines (DVM), which consist of tightly coupled computational resources that provide a flexible environment to manage and coordinate parallel application execution and dramatically improve scalability.

Our HARNESS effort is especially focused on providing fault tolerance for long Message Passing Interface (MPI) programs. The MPI plug-in for HARNESS developed by ICL, called FT-MPI (Fault Tolerant MPI), provides support for fault-tolerant applications that is crucial for large, long running simulations. Currently, fault tolerant applications cannot be built with MPI because MPI is unable to handle failures gracefully. Under FT-MPI, applications have flexible control over how failures are handled, allowing message passing applications to be built that can survive node failures without the need to continuously make expensive state and checkpoint dumps to disk and supporting the development of algorithms that can adapt to failures, which is currently not possible with other implementations of MPI.

Grid computing efforts such as NetSolve and HARNESS are motivated by, and lend themselves to, inter-institutional collaboration, and the Grid Application Development System (GrADS) in which ICL participates carries this collaboration to new heights. GrADS is a leading edge research project led by Rice University that includes more than a dozen top computer scientists from nine different universities. The GrADS project aims to address the key scientific and technical problems that must be solved in order to make it relatively easy to develop Grid applications for real problems and to tune those applications for high performance. This goal involves challenges in a wide spectrum of research areas, including software architectures, wide area information systems, languages, compilers, performance monitoring and analysis, numerical algorithms, and so on.

All of our Grid computing research is supported by the Scalable Intracampus Research Grid (SInRG), an NSF funded research infrastructure established by the UT computer science department under ICL's leadership. It mirrors within the boundaries of the Knoxville campus both the underlying technologies and the interdisciplinary research collaborations that are characteristic of the national and international technology Grid. SInRG, which in terms of hardware, at present, consists of six large compute clusters connected to UT's high performance campus network, provides a technological and organizational microcosm where ICL and UT researchers can address the key research challenges of grid-based computing using the advantages of local communication and local control.

PERFORMANCE EVALUATION

In addition to producing software that helps achieve high performance on parallel computers, ICL has been a leader in benchmarking and performance evaluation efforts that measure and report performance on theses machines. Our researchers have developed a number of benchmark codes. The LINPACK benchmark is a numerically intensive test that has been used for years to measure the floating-point performance of computers. Performance on this benchmark is the basis of the semi-annual TOP500 list that ranks the fastest 500 computers in the world. The LINPACK benchmark code solves a dense system of linear equations. For the TOP500, the version is used that allows the user to scale the size of the problem and to optimize the software to achieve the best performance for a given machine. In order to obtain uniformity across computers in performance reporting, the restriction is made that the algorithm used must conform to the standard operation count for LU factorization with partial pivoting. LINPACK performance does not reflect the overall performance of a given system, as no single number ever can. However, it does reflect the performance of a dedicated system for solving a dense system of linear equations. Because the problem is very regular, the performance achieved is quite high and the performance numbers give a good estimate of peak performance. In addition to the semi-annual TOP500 list published by ICL and the University of Mannheim, ICL Director Jack Dongarra and colleague Aad J. van der Steen publish a yearly Overview of Recent Supercomputers which reflects the technical state of the supercomputer arena as accurately as possible.

ICL researchers have led the development of a portable high-performance implementation of the LINPACK benchmark for distributed memory parallel computers, called High Performance LINPACK, or HPL. HPL contains many possible variants for the various operations, so as to leave the user with the opportunity of experimentally determining an optimal set of parameters for a given machine configuration. State-of-the-art algorithms are used, including recursive panel factorization with pivot search and column broadcast combined and a bandwidth-reducing swap broadcast algorithm. The algorithms are scalable in the sense that their parallel efficiency is maintained constant with respect to the per processor memory usage. In addition to providing a means for obtaining LINPACK results, HPL exemplifies the coding practices needed to obtain near peak performance on modern parallel architectures

Reflecting the strong emerging trend of cluster computing in high performance computing (HPC), the TOP500 team is developing a similar list to rank the world's top 100 cluster

PERFORMANCE EVALUATION

LINPACK Benchmark/HPL · PAPI · SparseBench · TOP500 Supercomputers/TOP100 Clusters

computing systems. Currently there is no publicly available basis that would allow the compilation of statistics about different technologies and the application areas of cluster computing. The coverage of cluster computing by the TOP500 is not sufficient to produce specialized statistics about this increasingly important HPC segment. To promote the development of this new list, the TOP500 team has decided to start the collection of data about high-performance clusters and rank them initially by peak-performance only. At the same time, the TOP500 team is discussing with the IEEE Task Force on Cluster Computing (IEEE TFCC) the proper choice of a benchmark for ranking clusters.

ICL researchers have also helped develop the SparseBench benchmark suite for iterative methods on sparse matrices. Solution of sparse linear systems, such as those derived from Partial Differential Equations (PDEs), form an important problem area in numerical analysis as well as being the basis of computational problems in a number of application areas, including computational fluid dynamics and structural mechanics. Unlike in the case of dense linear systems, solution of sparse systems does not entail much reuse of data. Thus, algorithms for sparse matrices will be more bound by memory speed than by processor speed. SparseBench uses common iterative methods, preconditioners, and storage schemes to evaluate machine performance on typical sparse operations.

In addition to developing benchmarks, our researchers 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. These counters exist as a small set of registers that count events, which are occurrences of specific signals related to the processor's function. Monitoring these events facilitates correlation between the structure of source/object code and the efficiency of the mapping of that code to the underlying architecture. This correlation has a variety of uses in performance analysis including hand tuning, compiler optimization, debugging, benchmarking, monitoring and performance modeling. The interface, called the Performance API, or PAPI, not only provides a standard set of routines for accessing counter data, but also defines a common set of performance metrics considered relevant and useful for application performance tuning. PAPI provides two interfaces to the underlying counter hardware; a simple, high level interface for the acquisition of simple measurements and a fully programmable, low level interface directed towards application and tool developers with more sophisticated needs. The PAPI library interface and reference implementations have been widely adopted by performance tool developers, including the SvPablo project at the University of Illinois, the TAU project at the University of Oregon, and the VProf project at Sandia National Laboratories. By using PAPI, tool developers can devote most of their efforts to tool design and implementation rather than to low-level details of re-implementing access to hardware counters.

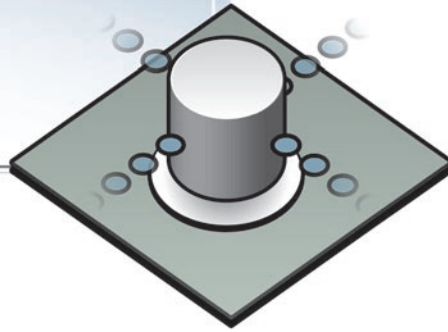

SOFTWARE REPOSITORIES

In the early eighties the rapid increase in the importance of computational science in several core disciplines created a powerful need for the timely and efficient dissemination of crucial numerical software then scattered around the community, a need that could no longer be addressed by a combination of word of mouth, magnetic tapes and the postal service. This led to the creation of one of the first open source efforts; the highly successful Netlib repository of freely available mathematical software, documents, and databases, which to date has recorded more than 180 million requests. Mirrored across more than a dozen sites around the world, Netlib has been, and continues to be, an important fixture of the scientific computing community. Using the highly successful Netlib mathematical software repository as a model, the National High Performance Software Exchange (NHSE) was formed by several academic and government institutions in the mid 1990s (including ICL). NHSE's goal was to establish discipline-oriented software repositories that could be contributed to and maintained by experts in their respective fields. Because of the need to share software between organizations and across disciplines, repository interoperability was an important goal. These efforts formed the cornerstone of our work in the realm of software repositories.

To enable the establishment and interoperability of discipline-oriented software repositories, the NHSE project spawned the development of the Repository in a Box (RIB) toolkit. RIB was originally conceived as a tool to facilitate the exchange and reuse of high performance applications within the high performance computing (HPC) community, and this remains its primary use today. However, under continued development by ICL, RIB has evolved in such a way that it now supports the creation and exchange of nearly any digital object. Each repository created with RIB contains a data model to which the cataloged metadata conforms. The data models supported by RIB are entity-relationship models and are extremely flexible. RIB uses a default data model that has been standardized by the Institute of Electrical and Electronics Engineers (IEEE). This standard, IEEE Std. 1402 - Basic Interoperability Data Model (BIDM), defines the minimal set of information necessary for the exchange of digital library objects between libraries.

As computationally intensive modeling and simulation have become staples of scientific life across many disciplines, the problem of acquiring and sustaining the necessary expertise in scientific computing is becoming increasingly acute. Building on the success of both RIB and Netlib, we are addressing these problems through the creation of Active Netlib, which provides an active collection of high-quality mathematical software resources in the context of an inquiry-based learning environment for computational science and engineering education. Active Netlib essentially takes advantage of the existing Netlib resources and extends them through the use of RIB, which will allow the Netlib collection to be selectively mirrored and contributed to by multiple participants. Furthermore, Active Netlib leverages NetSolve, a client-server system for accessing hardware and software resources over a network that provides an active interface to the contents of Netlib. NetSolve essentially constructs network-accessible objects with executable content from the software packages in Netlib. By making the subroutines housed in Netlib available over the network on computational servers, NetSolve enables access to up-to-date mathematical software from a variety of client interfaces running on users' workstations, without requiring the users to download and install the software themselves.

Another way in which we continue the progression of our repository efforts is our work on NetBuild, which is a project to make it easier for authors and users of scientific computation software to utilize standardized mathematical software libraries, many of which may be located in the Netlib repository or similar collections. Specifically, NetBuild attempts to eliminate the need for authors and users to locate, download, configure, compile, and install each of the mathematical software libraries that are required by a program. In contrast to Active Netlib, which provides facilities for remote execution of computations on the user's data using Netlib servers, NetBuild securely downloads the code necessary to perform those computations on the user's computer. NetBuild will attempt to find pre-compiled versions of those software libraries that are suitable for the target platform, automatically download those components, and incorporate them into the program. When no pre-compiled libraries are available, NetBuild will have the capability to download the source code and to compile the libraries from source. Though similar facilities exist on some platforms, notably Linux, FreeBSD, and NetBSD, NetBuild attempts to be usable on a wide range of platforms. It is also specifically intended for use with high-performance mathematical software, which, for good efficiency, often requires fine-grained tuning to specific characteristics of target platforms.

PEOPLE

As is the case with most organizations, our employees are instrumental in our success. Equally important are the working relationships we have established with individuals and organizations within the high performance computing (HPC) community. Our staff, our partners and collaborators, and the many commercial vendors with which we work have helped us create a strong foundation for fostering creative, original research.

We have been very successful at attracting experts and top researchers that comprise our staff. With a large number of staff, we are able to apply adequate people resources to the projects on which we work. Currently, we employ 14 students and 31 full or part-time staff, many of whom worked for us as students themselves. Because ICL is known internationally as a leading HPC academic research group, we have been successful in attracting research scientists from around the world. Proudly, our staff includes representatives from many countries. Our ability to attract such experts from around the world is only one reason ICL remains an HPC research leader. Table 1 (p. 16) provides information about our current staff.

Equally important to our group are our students. As part of the computer science (CS) department of a large university, ICL has access to both graduate and undergraduate students. With a CS program consisting of nearly 200 students, additional help with our projects is just a job posting away. These students represent a resource that is not readily available to many research groups, and we have been

ICL Group
August 2002

very proactive in securing internships and assistantships for those students who are motivated, hard working, and willing to learn. Table 2 (p. 17) provides a list of our current students.

In addition to our employees, we routinely host numerous visitors from around the globe. While many of our visitors stay briefly to give seminars or presentations, many remain with us for as long as a year collaborating, teaching, and learning. Though many of our visitors are professors from various international universities, we also host researchers and administrators from many research institutions. 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. In fact, many Ph.D. students from universities as far away as Japan have passed through our doors in an effort to broaden their understanding of linear algebra techniques and how we apply them to our research. The experience shared between our visitors and ourselves has been extremely beneficial to us, and we will continue providing opportunities for visits from our international colleagues in research. See Table 3 (p. 17) for the many guests who have stopped by in the last year to exchange ideas and share their expertise with us. We have worked hard to create and maintain many collaborative relationships and are always eager to open doors to new opportunities for sharing research endeavors.

As proud as we are of our current students and staff, we are equally proud that many of our former students and staff have moved on to do further interesting and useful research at places around the world. Many of our alumni have gone on to apply the knowledge gained from their time at ICL with companies such as Hewlett-Packard, Hitachi, IBM, Inktomi, Intel, Microsoft, Myricom, NEC, SGI, Sun Microsystems, and many others. Table 4 (p. 18) shows a list of past students, staff, and visitors to our group.

TABLE 1. CURRENT ICL STAFF

Sudesh Agrawal Research Associate	MS - University of Tennessee NetSolve	**Shirley Moore** Associate Director	PhD - Purdue University ICL Administration Active Netlib, NHSE, PAPI, RIB
Antonin Bukovsky Research Assistant	BS - University of Tennessee FT-MPI, HARNESS	**Keith Moore** Senior Research Associate	MS - University of Tennessee HARNESS, NA-Net, NA-Digest, NetBuild, NetSolve
Tom Cortese Research Scientist	PhD - University of Illinois at Urbana-Champaign CE Onsite Lead	**Terry Moore** Associate Director	PhD - UNC Chapel Hill ICL Administration NetSolve, RIB, SInRG
David Cronk Research Leader	PhD - College of William & Mary Parallel I/O	**Phil Mucci** Research Consultant	MS - University of Tennessee PAPI
Jack Dongarra Univ. Distinguished Professor	PhD - University of New Mexico ICL Director	**Paul Peltz** Computer Information Specialist	Systems Administration TORC
Victor Eijkhout Research Assistant Professor	PhD - Catholic Univ., Netherlands NetSolve, SANS Effort, SALSA	**Tracy Rafferty** Business Manager	ICL Administration
Brett Ellis Senior Computer Systems Specialist	BS - University of Tennessee Systems Administration GrADS, SInRG, TORC	**Ken Roche** Research Associate	MS - University of Tennessee SANS Effort, LFC
Graham Fagg Research Associate Professor	PhD - University of Redding, UK FT-MPI, HARNESS, MPI_Connect, Parallel I/O	**David Rogers** Graphic Designer	BFA - University of Tennessee Graphic Design
Don Fike Research Assistant	BS - Illinois State University Active Netlib, NHSE, RIB	**Keith Seymour** Senior Research Associate	MS - University of Tennessee F2J, PAPI
Teresa Finchum Administrative Services Assistant	ICL Administration	**Dan Terpstra** Research Leader	PhD - Florida State University PAPI
Nathan Garner Research Consultant	MS - University of Tennessee NetSolve, Active Netlib	**Joe Thomas** Research Associate	MS - University of Tennessee Systems Administration
Jan Jones Publications Coordinator	MS - University of Tennessee	**Scott Wells** Assistant Director	MS - University of Tennessee Active Netlib, NHSE, RIB ICL Administration
Tracy Lee Senior Budget Assistant	BA - University of Tennessee ICL Administration	**Asim YarKhan** Senior Research Associate	MS - Penn State GrADS, NetSolve
Kevin London Research Assistant	PAPI, MPI_Connect	**Haihang You** Research Associate	MS - University of Tennessee PAPI
Michelle Miller Senior Research Associate	MS - University of Utah NetSolve		

TABLE 2. CURRENT ICL STUDENTS

Noel Black Graduate Research Assistant	Active Netlib
Zizhong Chen Graduate Research Assistant	Sparse Libraries and Algorithms
Jun Ding Graduate Research Assistant	Sparse Libraries and Algorithms
Andrew Downey Graduate Research Assistant	Parkbench, F2J
Brian Drum Undergraduate Research Assistant	NetSolve
Piotr Luszczek Graduate Research Assistant	Sparse Libraries and Algorithms
Eric Meek Graduate Research Assistant	NetSolve
Mei Ran Graduate Research Assistant	Active Netlib
Kiran Sagi Graduate Research Assistant	NetSolve
Zhiao Shi Graduate Research Assistant	NetSolve
Sathish Vadhiyar Graduate Research Assistant	FT-MPI, NetSolve
Michael Walters Undergraduate Student	NHSE, Netlib
Yuanlei Zhang Graduate Research Assistant	SciDAC
Min Zhou Graduate Research Assistant	PAPI

TABLE 3. RECENT VISITORS TO ICL

Ashok Adiga	United Devices - US
Vassil Alexandrov	University of Reading - United Kingdom
Ed Anderson	Lockheed Martin Services, Inc. - US
Richard Barrett	Los Alamos National Laboratory (LANL) - US
Iain Duff	Rutherford Appleton Laboratory - United Kingdom
Gene Golub	Stanford University - US
Tomo Hiroyasu	Doshisha University - Japan
Gene Ice	Oak Ridge National Laboratory (ORNL) - US
David Keyes	Old Dominion University - US
Chris Moulding	Enterasys Networks, Inc. - US
Rick Muller	California Institute of Technology - US
Roldan Pozo	National Institute of Standards and Technology (NIST) - US
Martin Swany	University of California, Santa Barbara (UCSB) - US
Daisuke Takahashi	University of Tsukuba - Japan
Jikku Venkat	United Devices - US
Jon Weissman	University of Minnesota - US
Felix Wolf	Research Centre Jülich - Germany
Junlong Zhao	Joint Institute for Computational Science - US

TABLE 4. ICL ALUMNI

Carolyn Aebischer	Graduate Student	1990-1993
Ed Anderson	Research Associate	1989-1991
Dorian Arnold	Research Associate	1999-2001
Zhaojun Bai	Post Doc	1990-1992
Ashwin Balakrishnan	Graduate Research Assistant	2001-2002
Richard Barrett	Graduate Research Assistant	1992-1994
Alex Bassi	Research Associate	2000-2001
Micah Beck	Research Associate Professor	1998-2001
Adam Beguelin	Post Doc	1991-1994
Annamaria Benzoni	Visiting Research Associate	1991
Scott Betts	Undergraduate Student	1997-1998
Susan Blackford	GRA, Research Associate	1989-2001
Fernando Bond	Principal Secretary	1999-2000
Randy Brown	Undergraduate Student	1997-1999
Cynthia Browne	Undergraduate Student Assistant	2001
Murray Browne	Research Associate	1998-1999
Greg Bunch	Artist	1995
Henri Casanova	GRA, Post Doc	1995-1998
Sharon Chambers	Undergraduate Student	1998-2001
Jaeyoung Choi	Post Doc	1994-1996
Eric Clarkson	Artist	1998-1999
Andy Cleary	Post Doc	1995-1997
Jason Cox	Graduate Student	1993-1997
Cricket Deane	Research Associate	1998-1999
Frederic Desprez	Post Doc	1994-1995
Martin Do	Graduate Research Assistant	1993-1994
Leon Dong	Graduate Research Assistant	2000-2001
David Doolin	Research Associate	1997
Mary Drake	Office Supervisor	1989-1992
Zack Eylar-Walker	Undergraduate Student	1998
Markus Fischer	Visiting Student	1997-1998
Erika Fuentes	Graduate Research Assistant	2000-2001
Lynn Gangwer	Principal Secretary	2000-2001
Tracy Gangwer		
Jonathan Gettler	Graduate Research Assistant	1996
Eric Greaser	Graduate Research Assistant	1993
Stan Green	GRA, Senior Research Associate	1992-1996
Hunter Hagewood	Graduate Research Assistant	2000-2001
Christian Halloy	Associate Director	1996-1997
Sven Hammarling	Visiting Professor	1996-1997
Hidehiko Hasegawa	Visiting Research Associate	1995-1996
Satomi Hasegawa	Visiting Research Associate	1995-1996
Chris Hastings	Research Associate	1996
David Henderson	Undergraduate Student	1999-2001
Greg Henry	Post Doc	1996
Sid Hill	Undergraduate Student	1996-1998
George Ho	GRA, Research Consultant	1998-2000
Jeff Horner	Undergrad, Research Associate	1995-1999
Yan Huang	Graduate Research Assistant	2000-2001
Chris Hurt	Graduate Research Assistant	2002
Paul Jacobs	Undergraduate Student	1992-1995

TABLE 4. (CONTINUED) **ICL ALUMNI**

Weizhong Ji	Graduate Research Assistant	1999-2000
Weicheng Jiang	Post Doc	1992-1995
Song Jin	Graduate Student	1998
Balajee Kannan	Graduate Research Assistant	2000-2001
David Katz	Graduate Research Assistant	2002
MyungHo Kim	Visiting Scholar	1998
Youngbae Kim	Graduate Research Assistant	1992-1996
Michael Kolatis	GRA, Research Assistant	1993-1996
DongWoo Lee	Visiting Scholar	2000-2002
Todd Letsche	Graduate Student	1993-1994
Sharon Lewis	GRA, Manager	1992-1995
Xiang Lee	Graduate Research Assistant	2001
Chaoyang Liu	Graduate Research Assistant	2000
Weiran Li	Graduate Research Assistant	2002
Matt Longley	Undergraduate Student	1999
Richard Luczak	ASC MSRC Onsite Lead	2000-2001
Robert Manchek	Research Associate	1990-1996
Paul McMahan	GRA, Program Director	1994-2000
Cindy Mitchell	Principal Secretary	2001-2002
Steven Moulton	Graduate Research Assistant	1991-1993
Peter Newton	Post Doc	1994-1995
Caroline Papadopoulos	Graduate Student	1997-1998
Leelinda Parker	Undergraduate Student	2002
Antoine Petitet	GRA, Post Doc, Research Scientist	1993-2001
Roldon Pozo	Post Doc	1992-1994
Tammy Race	Graduate Research Assistant	1999-2001
Ganapathy Raman	Graduate Research Assistant	1998-2000
Yves Robert	Visiting Professor	1996-1997
Tom Rothrock	Undergraduate Student	1998
Tom Rowan	Collaborating Scientist	1993-1997
Evelyn Sams	Principal Secretary	1998-1999
Farial Shahnaz	Undergraduate Student	2001
Majed Sidani	GRA, Post Doc	1990-1992
Shilpa Singhal	Undergraduate Student	1996-1998
Thomas Spencer	Undergraduate Student	2000-2001
Erich Strohmaier	Post Doc Research Associate	1995-2001
Martin Swany	Research Associate	1996-1999
Judi Talley	Sr. Computer Systems Specialist	1993-1999
Keita Teranishi	Undergraduate Student , GRA	1998
John Thurman	Graduate Student	1998-1999
Françoise Tisseur	Post Doc	1997
Bernard Tourancheau	Post Doc	1993-1994
Robert van de Geijn	Post Doc	1990-1991
Scott Venckus	Graduate Research Assistant	1993-1995
Reed Wade	Research Associate	1990-1996
Clint R. Whaley	GRA, Senior Research Associate	1991-2001
Susan Wo	Graduate Research Assistant	2000-2001
Tinghua Xu	Graduate Research Assistant	1998-200
Tao Yang	Graduate Research Assistant	1999
Yong Zheng	Graduate Research Assistant	2001
Luke Zhou	Graduate Research Assistant	2000-2001

HARDWARE

Research in the many areas of high performance computing requires access to various hardware resources. As a research group, we take pride in our hardware assets. Currently, we manage more than 47 desktop computers ranging from Windows/Linux machines to Sun Sparc workstations, many of which are configured as servers. In addition, we also have a host of parallel computing machines, which include architectures such as two IBM Power 3s, two commodity-based clusters consisting of 24 total Dual Intel-based commodity machines, and a non-production cluster consisting of six machines of various configurations. We also have an SGI Octane and two Compaq Alphas. The newest cluster in our group is a cluster of seven Itanium-based machines. There are two Single 900Mhz Itanium2s, three Dual 800Mhz Itaniums, and two Single 750Mhz Itanium processor machines. Our close working relationship with hardware vendors has proven to be very beneficial due to the fact that nearly all of our desktops and servers were either donated or are on loan.

As part of the UT computer science department, we share a 2500 square foot, state-of-the-art machine room dedicated to housing servers, clusters, and other single and multiple processor, high performance computing (HPC) architectures exceeding 200 total machines. The department's newest addition is a cluster of 32 Dual P4 2.4Ghz machines with a Dolphin interconnect that was donated for

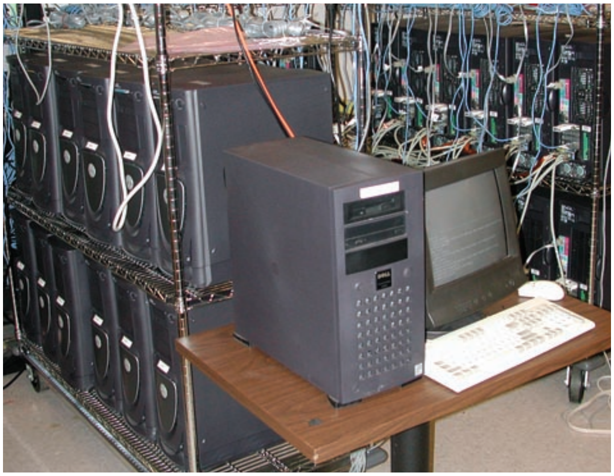

Cluster of 32 Dual P4 2.4 Ghz machines

State-of-the-art machine room

that cluster. Our clusters are arranged in the classic Beowulf configuration in which machines are connected by low latency, high-speed network switches. Having such variety of architectures at our disposal positions us to take advantage of the rapid changes in the dynamic field of computational research.

In addition to our local resources, we are also fortunate to have access to many remote resources. Due to the many organizations and institutions with which we collaborate, a wide range of hardware architectures are made available to our research staff. Remaining at the forefront of computational research requires that our staff have access to the latest computing technology. Below are some of the architectures currently available for our use:

Compaq ESs, GSs, and SCs	SGI Origin 2000s, 3000s, and 3800s
Cray SV-1s and T3Es	
IBM SPs	Sun E10000s

Our ability to harness the computing power of multiple architectures allows us to perform comprehensive software development and testing. In addition, we have the heterogeneous resources necessary to parallelize many applications that previously ran only sequentially.

To further our collaboration practices, we are also currently installing an Access Grid, which consists of various interfaces and environments on the grid that support distributed meetings, lectures, tutorials, and other collaborative efforts. The Access Grid comprises multiple video cameras, speakers, projectors, and PCs to form a seamless resource for conducting timely, online collaborative activities.

PARTNERSHIPS

For more than a decade, ICL has enjoyed many mutually beneficial working relationships with institutions all over the globe. The high performance computing (HPC) community consists of academic institutions, research centers, branches of the federal government, and various other public and private organizations, both domestic and international. We attribute much of our growth to the solid partnership foundations we have established within this community. Our collaborative initiatives have strengthened our research efforts by allowing us to share both material and intellectual resources. Table 1 highlights many of our domestic partners and collaborators. In addition to our numerous government and academic partners, we have also enjoyed a strong working relationship with many commercial software vendors. Included in the list of software vendor collaborators are Etnus, Inc., developer of the TotalView debugger; Kuck and Associates, Inc., developer of the KAP/Pro toolset; and Pallas, developer of the Vampir performance visualization and analysis tool.

Figure 1 shows the geographical location of many of the domestic and international partners and collaborators in research with whom we continue to work.

TABLE 1. **DOMESTIC PARTNERS AND COLLABORATORS**

Argonne National Laboratory	Lawrence Berkely National Laboratory	National Partnership for Advanced Computational Infrastructure (NPACI)
California Institute of Technology	Lawrence Livermore National Laboratory (LLNL)	National Science Foundation (NSF)
Center for Advanced Computing Research (CACR)	Los Alamos National Laboratory (LANL)	Oak Ridge National Laboratory (ORNL)
Defense Advanced Research Projects Agency (DARPA)	Microsoft Research	Rice University
Department of Defense (DoD)	Morehouse College	San Diego Supercomputing Center (SDSC)
Department of Energy (DOE)	National Aeronautics and Space Administration (NASA)	Silicon Graphics Incorporated (SGI)
Emory University	National Computational Science Alliance (NCSA)	Sun Microsystems
Information Sciences Institute (ISI)	National HPCC Software Exchange (NHSE)	The Salk Institute
Intel	National Institute of Standards and Technology (NIST)	University of California, Berkeley
International Business Machines (IBM)		University of California, Davis
Internet2		University of California, San Diego
Joint Institute for Computational Science (JICS)		University of California, Santa Barbara
		University of Kentucky

FIGURE 1. Geographical distribution of ICL partners and collaborators

TABLE 2. INTERNATIONAL PARTNERS AND COLLABORATORS

Danish Computing Centre for Research and Education LYNGBY, DENMARK	Laboratoire de l'Informatique du Parallelisme, École Normale Superieure de Lyon LYON, FRANCE
Department of Mathematical and Computing Sciences Tokyo Institute of Technology TOKYO, JAPAN	Laboratoire Réseaux Haut Débits et Support d'Applications Multimedia (RESAM) Jeune Equipe de l'Université Claude Bernard de Lyon LYON, FRANCE
Department of Mathematics, University of Manchester MANCHESTER, ENGLAND	Mathematical Institute, Utrecht University NETHERLANDS
Electrotechnical Laboratory, Computer Systems Division TSUKUBA, JAPAN	Monash University MELBOURNE, AUSTRALIA
European Centre for Research and Advanced Training in Scientific Computing (CERFACS) TOULOUSE, FRANCE	Rutherford Appleton Laboratory OXFORD, ENGLAND
Fakultät für Mathematik und Informatik, Universität Mannheim MANNHEIM, GERMANY	Ecole Polytechnique Fédérale de Lausanne LAUSANNE, SWITZERLAND
Institut für Wissenschaftliches Rechnen, ETH Zentrum - ZÜRICH, SWITZERLAND	Soongsil University SEOUL, SOUTH KOREA
Istituto per le Applicazioni del Calcolo "Mauro Picone" del C.N.R. - ROME, ITALY	The Numerical Algorithms Group Ltd. OXFORD, ENGLAND
Kasetsart University BANGKOK, THAILAND	University of Umeå UMEÅ, SWEDEN

2002 PUBLICATIONS

A complete bibliography of our publications and technical reports from 1999 to present can be found on our web site at http://icl.cs.utk.edu/publications/. Most of these are also downloadable.

Agrawal, S. "Hardware Software Server in NetSolve," *ICL Technical Report*, ICL-UT-02-02, March 11, 2002.

Agrawal, S., Arnold, D., Blackford, S., Dongarra, J., Miller, M., Sagi, K., Shi, Z., Seymour, K., Vahdiyar, S. "Users' Guide to NetSolve v1.4.1," *ICL Technical Report*, ICL-UT-02-05, June, 2002.

Bassi, A., Beck, M., Fagg, G., Moore, T., Plank, J., Swany, M., Wolski, R. "The Internet BackPlane Protocol: A Study in Resource Sharing," *Proceedings of the second IEEE/ACM International Symposium on Cluster Computing and the Grid* (*CCGRID 2002*), Berlin, Germany, May 21-24, 2002.

Beck, M., Arnold, D., Bassi, A., Berman, F., Casanova, H., Dongarra, J., Moore, T., Obertelli, G., Plank, J., Swany, M., Vadhiyar, S., Wolski, R. "Middleware for the Use of Storage in Communication," *Parallel Computing* (to appear), 2002.

Casanova, H., Bartol, T., Berman, F., Birnbaum, A., Dongarra, J., Ellisman, M., Faerman, M., Gockay, E., Miller, M., Obertelli, G., Pomerantz, S., Sejnowski, T., Stiles, J., Wolski, R. "The Virtual Instrument: Support for Grid-enabled Scientific Simulations," *Journal of Parallel and Distributed Computing* (submitted), April, 2002.

Cuenca, J., Giminez, D., Gonzalez, J., Dongarra, J., Roche, K. "Automatic Optimisation of Parallel Linear Algebra Routines in Systems with Variable Load," submitted to *EuroPar 2002*, Paderborn, Germany, August 27, 2002.

Dongarra, J. "Performance of Various Computers Using Standard Linear Equations Software (LINPACK Benchmark Report)," *University of Tennessee Computer Science Technical Report*, CS-89-85, 2002.

Dongarra, J., Eijkhout, V. "Self-adapting Numerical Software for Next Generation Applications," *ICL Technical Report*, ICL-UT-02-07, 2002.

Dongarra, J., Eijkhout, V., van der Vorst, H. "An Iterative Solver Benchmark," *Scientific Programming* (to appear), 2002.

Eijkhout, V. "Polynomial Acceleration of Optimised Multi-grid Smoothers; Basic Theory," *ICL Technical Report*, ICL-UT-02-03, *LAPACK Working Note 156*, March 28, 2002.

Hiroyasu, T., Miki, M., Shimosaka, H., Sano, M., Tanimura, Y., Mimura, Y., Yoshimura, S., Dongarra, J. "Truss Structual Optimization Using NetSolve System," *Meeting of the Japan*

Society of Mechanical Engineers, Kyoto University, Kyoto, Japan, October 12-13, 2002.

Hiroyasu, T., Miki, M., Shimosaka, H., Tanimura, Y., Dongarra, J. "Optimization System Using Grid RPC," *Meeting of the Japan Society of Mechanical Engineers*, Kyoto University, Kyoto, Japan, October 12-13, 2002.

Kelleher, Jr., M. "Development of the PICMSS NetSolve Service," *ICL Technical Report*, ICL-UT-02-04, April, 2002.

Kennedy, K., Broom, B., Cooper, K., Dongarra, J., Fowler, R., Gannon, D., Johnsson, L., Mellor-Crummey, J., and Torczon, L. "Telescoping Languages: A Strategy for Automatic Generation of Scientific Problem-Solving Systems from Annotated Libraries," *Journal of Parallel and Distributed Computing Volume 61, Number 12*, pp 1803-1826, December 2001.

Kennedy, K., Mazina, M., Mellor-Crummey, J., Cooper, K., Torczon, L., Berman, F., Chien, A., Dail, H., Sievart, A., Angulo, D., Foster, I., Gannon, D., Johnsson, L., Kesselman, C., Aydt, R., Reed, D., Dongarra, J., Vadhiyar, S., Wolski, R. "Toward a Framework for Preparing and Executing Adaptive Grid Programs," *Proceedings of NSF Next Generation Systems Program Workshop* (*International Parallel and Distributed Processing Symposium 2002*) (to appear), Fort Lauderdale, FL, April 15, 2002.

Lee, D., Dongarra, J., Ramakrishna, R. "visPerf: Monitoring Tool for Grid Computing," *Workshop on Grid Computing* (submitted), Baltimore, MD, November 18, 2002.

Moore, S. "A Comparison of Counting and Sampling Modes of Using Performance Modeling Hardware," *ICL Technical Report*, ICL-UT-02-01, 2002-01-02.

Moore, S., Baker, A., Dongarra, J., Halloy, C., Ng, C. "Active Netlib: An Active Mathematical Software Collection for Inquiry-based Computational Science and Engineering Education," *Journal of Digital Information special issue on Interactivity in Digital Libraries* (to appear), 2002.

Nakada, H., Matsuoka, S., Seymour, K., Dongarra, J., Lee, C., Casanova, H. "GridRPC: A Remote Procedure Call API for Grid Computing," *Workshop on Grid Computing* (submitted), Baltimore, MD, November 18, 2002.

Roche, K., Dongarra, J. "Deploying Parallel Numerical Library Routines to Cluster Computing in a Self Adapting Fashion," *Parallel Computing* (submitted), April, 2002.

Seymour, K., Nakada, H., Matsuoka, S., Dongarra, D., Lee, C., Casanova, H. "GridRPC: A Remote Procedure Call API for Grid Computing," *ICL Technical Report*, ICL-UT-02-06, June 8, 2002.

Shende, S., Malony, A., Moore, S., Mucci, P., and Dongarra, J. "Integrated Tool Capabilities for Performance Instrumentation and Measurement," *2003 International Symposium on Performance Evaluation of Computer and Telecommunication Systems* (*ISPASS 2003 - International Symposium on Performance Analysis of Systems and Software*) (submitted for 2003).

Vadhiyar, S., Dongarra, J. "A Framework for Migrating Applications Under Changing Load Conditions in the Grid," *Workshop on Grid Computing* (submitted), Baltimore, MD, November 18, 2002.

Vadhiyar, S., Dongarra, J. "A Metascheduler For The Grid," *HPDC 2002* (submitted) , 2002.

Yarkhan, A., Dongarra, J. "Experiments with Scheduling Using Simulated Annealing in a Grid Environment," *Workshop on Grid Computing* (submitted), Baltimore, MD, November 18, 2002.

CENTERS *of* EXCELLENCE

The University of Tennessee is focused on becoming one of America's top 25 public research universities. In 2000 the university embarked upon a five-year, $335 million Tennessee Plan for Academic Excellence, including the establishment of nine new research centers of excellence, scholarships to keep the best and brightest students in Tennessee, and an infusion of dollars to improve academic programs and establish additional smaller centers of excellence.

The university's nine research centers of excellence - five in Knoxville and four at the Health Science Center in Memphis - promise 1,000 new jobs and as many as 20 new companies for Tennessee. The centers represent a $280 million investment, including $56 million from the university and the state and the balance primarily from grants. Below is a list of each of these centers including their directors.

CENTERS BASED IN KNOXVILLE

Advanced Materials Center

DIRECTOR: DR. WARD PLUMMER

Center for Information Technology Research

DIRECTOR: DR. JACK DONGARRA

Environmental Biotechnology Center

DIRECTOR: DR. GARY SAYLER

Food Safety Center

DIRECTOR: DR. STEPHEN P. OLIVER AND DR. ANN DRAUGHON

Structural Biology Center

DIRECTOR: DR. ENGIN SERPERSU

CENTERS BASED IN MEMPHIS

Connective Tissues Diseases Center

DIRECTOR: DR. ANDREW H. KANG

Genomics and Bioinformatics Center

DIRECTOR: DR. DAN GOLDOWITZ

Neurobiology and Imaging of Brain Disease Center

DIRECTOR: DR. S. T. KITAI

Vascular Biology Center

DIRECTOR: DR. LISA JENNINGS

CITR

BACKGROUND

The Center for Information Technology Research (CITR) was established in the spring of 2001 in order to drive the growth and development of leading edge Information Technology Research (ITR) at UT. Information Technology Research is a broad, multi-disciplinary area that investigates ways in which fundamental innovations in Information Technology affect and are affected by the research process.

The mission of CITR is to build up a thriving, well-funded community in basic and applied ITR at the University of Tennessee in order to help the university capitalize on the rich supply of research opportunities that now exist in this area. To carry out this mission CITR has implemented a two pronged strategy. First, CITR is investing in a diverse group of ITR laboratories, each one led by an established researcher or an emerging leader in some significant area of ITR. Second, CITR is developing a complimentary set of university-wide programs that can serve to foster innovative research ideas in the university community, seed the creation of new CITR laboratories, and help UT exploit the broadest possible spectrum of ITR opportunities.

CITR LABORATORIES

CITR is making strategic investments in several ITR Laboratories, targeting researchers with the potential for generating well-funded research and also with an eye toward diversifying the range of research opportunities and funding sources that UT can address. The first CITR laboratory, and the model for others, is the ICL. Dr. Dongarra serves as Director of both ICL and CITR. In addition to ICL, two more CITR Laboratories have been formed: The Logistical Computing and Internetworking (LoCI) Lab, co-directed by Professors Micah Beck and Jim Plank from the Computer Science Department, and The Institute for Environmental Modeling (TIEM) run by Professor Lou Gross from the Departments of Ecology and Evolutionary Biology and Mathematics. Combined, these three CITR laboratories have added more than six and a half million dollars in new research funding to UT's ITR community.

UNIVERSITY WIDE PROGRAMS

CITR has created several university wide programs to complement its ITR laboratory thrust. The first and most significant program offers Challenge Grants to IT researchers who will be applying for agency funding in the near future. The goal of the program is to provide UT faculty with seed money to support the pursuit of new funding opportunities in diverse areas of ITR. Challenge Grants represent an investment in the overall work of particular IT investigators or groups for a given year, an investment that can be used to hire a promising graduate student, support relevant travel, purchase special equipment, develop early prototypes, help in the proposal process, and so on.

CITR also provides start-up funding to support the recruitment of outstanding new faculty in different areas of ITR, scholarships to attract Ph.D. students of exceptional promise to UT, and supplemental funding for ITR related conferences at UT.

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

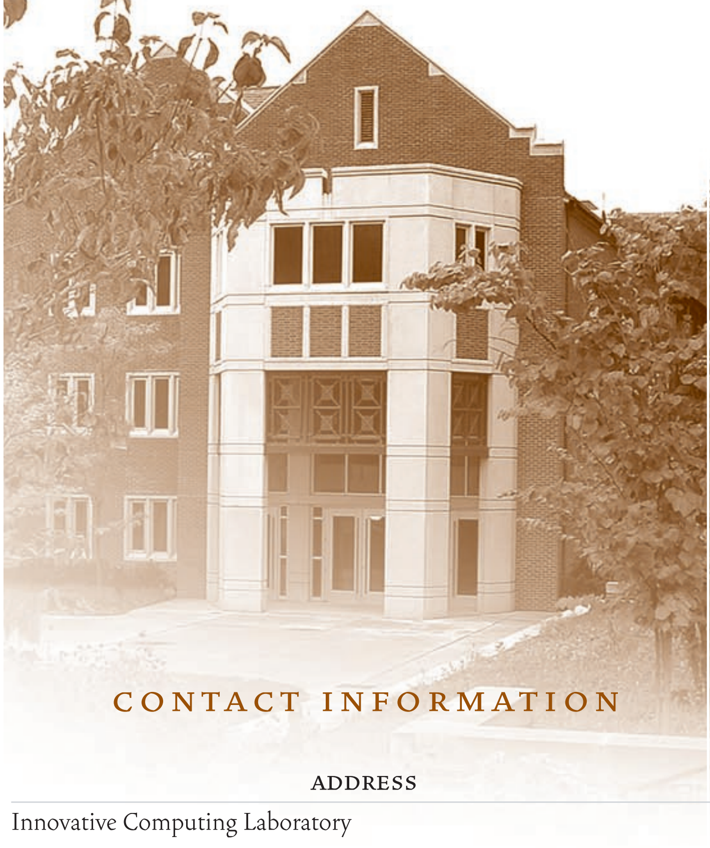

CONTACT INFORMATION

ADDRESS

Innovative Computing Laboratory
Suite 203
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/