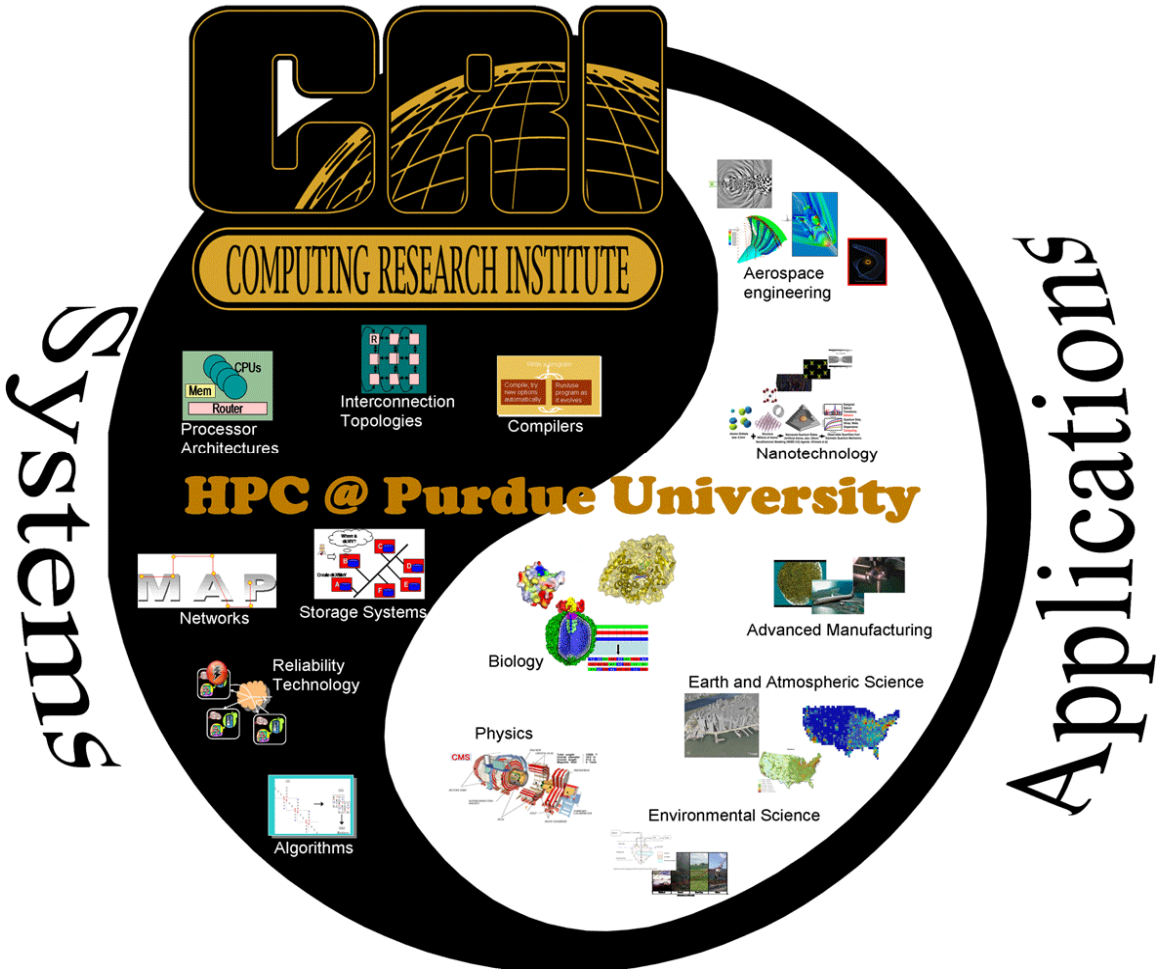

Computing Research Institute
Purdue University
Annual Report 2007

302 Wood Street
West Lafayette, IN 47907

Computing Research Institute
Purdue University
Annual Report 2007

Computing Research Institute
Executive Summary
2007

1. Executive Summary

Simulation-based Science and Engineering: CRI aims to revolutionize research through computational simulation. In the vision of many funding agencies, simulation-based science and engineering has become a major research thrust, through which future projects across many fields of science and engineering will be conducted. Computational simulation represents the *3rd pillar of science*, complementing theory and real experiments. While real experiments will always have an important place in validating computational results, simulation can literally "go where no one has gone before." Through simulation, we can overcome laws of physics, explore extreme space and time scales, break financial barriers, and exceed human capacities. Playing a leadership role in this revolution is of great strategic importance for Purdue.

High-performance computing (HPC): HPC technology is an exciting and key enabler of simulation-based science and engineering. While the power of today's workstations - and even laptops - is high enough to enable computational experiments in many research areas, advanced simulations of physical phenomena, such as investigations at atomic and molecular levels, will need powerful petascale platforms, with hundreds of thousands of processors. HPC means more than just high-end computing. Parallel processing technology has been developed through HPC research. Today, parallel processing is becoming a mainstream technology, as multicore architectures replace single processors in all computer systems. Parallel software techniques are of crucial importance for these architectures, just as they are for massively parallel, high-end systems. Parallel computing technology also exemplifies the forefront position of HPC in information technology. Software and hardware techniques that have been developed over the past two decades are now becoming essential for mainstream computing technology. As the forefront of this technology is advancing rapidly, it is also important to consider the branch of HPC that *creates* the parallel and high-end computer systems of the future. Hence, for CRI, both computational application research and HPC systems research are essential.

HPC Funding Opportunities: Funding agencies have recognized the importance of simulation-based research and HPC technology. The National Science Foundation (NSF) has committed nearly $500M to building up the national HPC computer infrastructure. Similarly large investments are expected in HPC software technology - both applications and systems. Plans to create a simulation-based NSF program are being formed. Similarly there have been efforts at the Department of Energy to create a program that pursues exascale computing - 1000x the compute power of the most

powerful computer today. These plans show the importance and opportunities that HPC research bears in the clearest terms.

Interdisciplinary Research: A key role of CRI, as a center, is to foster interdisciplinary projects. HPC and simulation-based research efforts at individual groups at Purdue are strong. Forming interdisciplinary project teams can leverage this strength significantly. In creating such interdisciplinary teams, good relationships of CRI with other centers that have similar missions are essential. In January 2007, CRI merged with Purdue's Cyber Center (CC), forming CC's high-end computing branch and establishing a clear link with this center. The relationship with ITaP's Rosen Center for Advanced Computing (RCAC) is also relevant. CRI does not plan to acquire HPC hardware resources for the community. As such resources and related services are essential for advanced computational projects, a strong relationship with RCAC is key to CRI's mission. To this end CRI - together with the CC - have moved to co- locate with RCAC in December 2007. Planning to strengthen collaboration further is underway.

Summary of CRI 2007 Activities

In pursuit of the described mission and opportunities, in 2007, CRI has engaged in four thrusts. Community building activities took a central role. The goal was to leverage the significant strength of individual research groups that are already engaged in simulation-based research and HPC technology at Purdue. The second thrust aimed at advancing HPC projects at Purdue, enhancing competitiveness for the significant funding opportunities that lay ahead. For advanced projects, external partnerships with research organizations and industry are of further importance. A third thrust was to foster such partnerships. Fourth, CRI saw a need to communicate Purdue successes in HPC research both on campus and externally.

CRI Successes

Among the successes achieved in 2007, the Purdue-Industry HPC workshop ranks high. The workshop aimed to initiate partnerships between Purdue faculty and industrial participants. The event exceeded expectations, attracting a large audience showing interest in HPC technology. Another success was the external funding received by the HPC community. Counting only 2007 proposal activities of faculty and staff directly involved in CRI and its leadership, the received funding was in the order of $20M. CRI has received acknowledgment from faculty reporting that it has helped move their research "light-years ahead." The community building activities resulted in faculty writing joint proposals that would otherwise not have been possible; several of these proposals received funding. Last but not least, the 2007 activities have led to better awareness among Purdue faculty and staff of CRI's role and of opportunities that exist to revolutionize research through simulation-based science and engineering. These successes have strengthened Purdue’s reputation in HPC and our ability to secure funding for big computational projects.

2. CRI Overview

2.1 CRI Organizational Structure

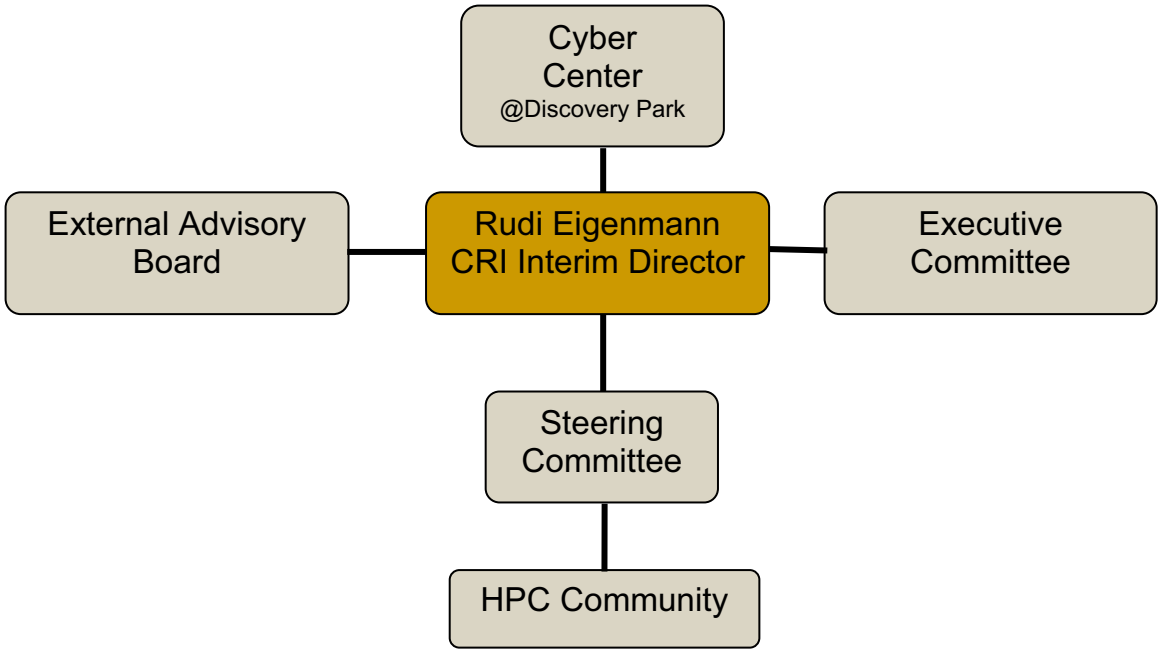

www.cri.purdue.edu

2.2 CRI Staff

CRI staff consists of

- Prof Rudi Eigenmann, interim Director of CRI, professor of ECE.
- Barb Fossum, Managing Director of the Cyber Center
- Faisal Saied, Senior Research Scientist. This position supported 50% by CRI and 50% by the Rosen Center for Advanced Computing.
- Mohamed Sayeed, Research Programmer. This position supported 50% by CRI and 50% by the Rosen Center for Advanced Computing.
- Melanie Lindsay, Secretary.

CRI has one part time work-study undergraduate and one part time graduate student.

2.3 Steering Committee



Noah Diffenbaugh, Earth & Atmospheric Sciences
Rudolf Eigenmann, Electrical and Computer Engineering (Interim Director)
Bernard Engel, Agricultural and Biological Engineering



Maarten V. De Hoop, Mathematics
Gerhard Klimeck, Electrical and Computer Engineering
Anastasios S. Lyrintzis, Aeronautics and Astronautics



Sam Midkiff, Electrical and Computer Engineering
Jayathi Murthy, Mechanical Engineering
Norbert Neumeister, Physics



Robert Skeel, Computer Science
Alejandro Strachan, Materials Engineering
T.N. Vijaykumar, Electrical and Computer Engineering

2.4 Executive Committee

Ahmed Elmagarmid, Director of Cyber Center
Jon Harbor, Associate Vice President for Research for Centers and Institutes
Leah Jamison, Dean of Engineering
Gerry McCartney, Vice President of Information Technology and CIO
Jeffrey Vitter, Dean of Science
Randy Woodson, Dean of Agriculture

2.5 HPC Community

CRI has identified nearly 200 faculty in over 30 departments in 7 Purdue colleges who list High Performance Computing (HPC) as one of their research areas. Many of these researchers use simulation-based methods to advance their disciplines; they are the primary users of HPC technology. The HPC Community also includes researchers that build the next-generation of parallel computer software and hardware systems. All these faculty are listed on CRI's web site: http://www.cri.purdue.edu/facultyadvis.cfm.

In pursuing its mission, CRI interacts most closely with this group of faculty members.

3. Building the Purdue HPC Community

Building the HPC community at Purdue has been a focus activity of CRI in 2007. There is strong presence of computational technology in many science and engineering fields at Purdue. Faculty in many departments are conducting research that uses computational simulation. The goal of community-building activities is to create synergy. The expected outcomes are stronger and more interdisciplinary projects, increased competitiveness in grant competitions, increased efficiency by sharing resources, accelerated research by exchanging ideas, and better demonstration of the strong research thrust in HPC that exists at Purdue.

CRI's community building activities included monthly lunch discussions with faculty, a seminar series in the topic area, poster sessions, and working seminars. CRI was also engaged in educational activities with students and staff.

Good relationships with campus organizations that pursue similar goals is essential. Through the faculty involved in CRI's leadership, there are natural links with the Department of Computer Sciences and the School of Electrical and Computer Engineering. Through CRI's new home in the Cyber Center, there is a strong link to other activities that involve interdisciplinary research with computing. A strong collaboration with RCAC was also considered essential. Means to this end were the shared staff positions (see CRI Staff), and an effort to make events, such the HPC luncheons and workshops, co-sponsored activities by CRI and RCAC.

3.1 CRI HPC Luncheons

CRI organized a very successful series of monthly HPC Luncheons through 2007 that contributed to building an HPC community at Purdue. At each luncheon, a few Purdue faculty members spoke about their research and engaged in discussions that could lead to collaborations, joint projects, and developing interdisciplinary proposals.

These meetings were attended by representatives from ITaP and the Rosen Center for Advanced Computing. In particular, Gerry McCartney, Purdue's Vice President for Information Technology, regularly updated the faculty on new developments in computing infrastructure at ITaP.

3.2 CRI Seminar Series

CRI has an active seminar series that brings in speakers from a broad range of disciplines on topics that are computational and interdisciplinary in nature. CRI has worked to invite leaders in relevant fields to visit Purdue, speak in the CRI seminar, and interact with Purdue faculty. In 2007, the following speakers visited Purdue and gave presentations.

Hector Klie, UT at Austin	Deflation Physics-based Solvers for Multiscale Porous Media Flow Applications
Luiz DeRose, Cray	Cray Adaptive Supercomputing: The Road to Petaflops
Sun Kim, IU	Mining Gene Clusters by Comparing Multiple Genomes
Max Gunzburger, Florida State University	Least-squares finite element methods
Joseph Grcar, LBL	John von Neumann and the Origins of Scientific Computing.
Eric Jakobsson, NCSA	The Computational Challenge of Soft Condensed Matter
Mikhail Ivanov, Institute of Theoretical and Applied Mechanics of the Siberian Branch of Russian Academy of Sciences	Computational Analysis of Hysteresis Phenomena in Shock Wave Reflection.
Alan Qi, Computer Science Department, Purdue University	Efficient Bayesian inference and learning for graphical models
Arjan van der Vaart, Arizon State University	Unfolding upon binding: Computer simulations of the unusual binding mechanism of the Ets-1 transcription factor
Henri Calandra, TOTAL	Seismic Processing Challenge and Computational Roadmap
Doug Arnold, Institute for Mathematics and its Applications, University of Minnesota	The Institute for Mathematics and its Applications
Tayfun Tezduyar, Rice University	Fluid-Structure Interaction Modeling in Arterial Fluid Mechanics
Laxmikant Kale, University of Illinois Urbana-Champaign	Petascale and Multicore programming Models: What is needed

3.3 CRI Poster Sessions

The first HPC Poster session was held on April 5, 2007 in the MSEE Atrium. Twenty-six research groups presented their work, related to high-performance computing. There were posters from many Engineering departments, from Math, Computer Sciences, Agriculture, Biology, Pharmacy, Chemistry, Technology, and from ITaP. It was felt that the first event was so successful that CRI would hold a poster session per semester. The second CRI Poster session was held on October 25, 2007 in Lawson. Twenty-four posters were presented at this session. Following the poster session, there were winners for Best Posters at each event where students were awarded cash prizes.

3.4 CRI Working seminars

The goal of working seminars is to get an overview from faculty about their computational applications/research, challenges and the needs for implementing these applications on parallel HPC systems. The seminars we bring together CRI staff and a few computer science/engineering faculty and students who might help address these challenges. Ideally, this would result in some form of collaboration between the faculty research group and the attendees on the computer systems side. It may also lead to future joint proposals. Some of the working seminar presentations in 2007 were given by-

Xiao Chuan, Biological Sciences
Aleena Alexenko, Aeronautics and Astronautics
Indrajeet Chaubey, Agricultural and Biological Engineering
Brian Pijanowsky, Forestry and Natural Resources
Wen Jiang, Biological Sciences
Ding Li, Aeronautics and Astronautics
Ganesh Subbarayan, Mechanical Engineering
Rabi Mohtar, Agricultural and Biological Engineering

3.5 Educational Activities

Supporting the CS&E Program
The Computational Science and Engineering (CS&E) Program is an interdisciplinary graduate specialization that provides students with the opportunity to study a specific science or engineering discipline along with computing in a multi-disciplinary environment. CRI has a close relationship with the CS&E Program. The CRI seminars are co-sponsored by the CS&E Program and graduate students enrolled in CS&E regularly attend these seminars. CRI has also provided clerical support to CS&E in 2007

Sponsorship for SIAM Student Chapter
CRI contributed as sponsor to the first year of the newly founded SIAM Student Chapter at Purdue.

HPC Training Workshops
A High Performance Computing (HPC) Training Workshop was held at Purdue, August 6-7, 2007, in the Envision Center. This workshop was jointly sponsored by the Computing Research Institute (CRI) and the Rosen Center for Advanced Computing (RCAC).

As HPC is playing an increasingly important role in many areas of research in Science and Engineering, there is an urgent need to train more researchers in the use of this technology. The agenda for this workshop was designed to address this need. A new feature of the workshop was the inclusion of a presentation on multicore architectures by Prof Vijaykumar from Purdue's ECE Dept, who is a nationally recognized expert in computer architecture. Multicore architectures present a new challenge and an opportunity in all areas of computing including HPC, and including this topic in the workshop was very timely. Other presentations by faculty members covered Parallel Programming Models (Prof. Sam Midkiff, ECE), HPC architectures (Prof Mithuna Thottethodi, ECE) and OpenMP (Prof Rudi Eigenmann, ECE). Additional topics were covered by RCAC staff members, on topics that included MPI programming, the use of the TotalView parallel debugger, the Condor system for accessing unused cycles, and the use of software tools for analyzing performance and increasing programmer productivity.

The hands-on lab sessions were designed to give the workshop attendees a practical feel for the use of HPC resources, and parallel programming techniques. The Network for Computational Nanotechnology (NCN) that operates the nanoHUB provided technical support for videotaping the workshop, which has been made available on the nanoHUB.

The response to the workshop was very strong. The hands-on aspect limited the number of attendees to 18. There were over 25 people wait-listed.
As a result, CRI offered the workshop a second time September 10-11, 2007.

More details about these workshops can be found on
http://www.cri.purdue.edu/PurdueHPCTrainingWorkshop.cfm
http://www.cri.purdue.edu/PurdueHPCTrainingWorkshop_Sept2007.cfm

4. Driving HPC Projects at Purdue

One of the most important activities towards the goal of simulation- based science and engineering is to create the simulation tools that support advanced research projects in diverse fields. To this end, (i) CRI was able to sponsor a number of seed projects, (ii) CRI staff was involved directly in several efforts to create such simulation tools, (iii) other projects were enabled through CRI cost sharing and travel funds, and (iv) there was an initiative to create "superprojects".

4.1 CRI/SIRG Graduate Fellowship Awards

CRI was able to make eight 1-year SIRG graduate fellowship awards in areas that are close to the CRI mission. These awards serve as seed grants far advanced HPC projects and future funding opportunities. Several CRI/SIRG projects have resulted in proposals, that are listed with each project. In the call for proposals, CRI had asked for two P.I.s to be associated with each project, preferably from different departments, to encourage interdisciplinary collaborations. Most of the selected CRI/SIRG projects have reported that they submitted proposals for external funding for research related to their SIRG project.

1. Alina Alexeenko, AAE, Jayathi Murthy, ME, “Large-Scale simulation of coupled phonon-gas thermal transport in nanoscale geometries".

Related to this work the P.I.’s have submitted a preliminary proposal for NSF Cyber-Enabled Discovery and Innovation program: Proposal title: "CDI-Type I: Large-Scale Simulation of Coupled Phonon-Electron-Gas Nanoscale Thermal Transport".

2. Steve Dong, Math, Steve Frankel, ME, "A high order coupled fluid-structure solver for simulating blood flow artery wall interactions".

Related to this work, the P.I.’s in this project have submitted a proposal entitled "An efficient high-order method for fluid-structure interactions" to NSF-DMS. They also have an MRAC grant from TeraGrid for computer time ("Modeling physical and biomechanical phenomena: simulation and control"), involving certain activities related to the award.

3. Prof G. Klimeck, ECE/NCN, Prof Edwin Garcia, MSE, “Reaching petascale computing with a widely known, Purdue-centered simulation code in nanoelectronics for applications in solid-state lighting”.

The have submitted to NSF’s PetaApps program, and won an award.
Gerhard expects to write a proposal for the NSF PRAC program, in March.

4. Sam Midkiff, ECE, James Caruthers, ChE, "A parallel system for describing and analyzing complex chemical systems".

This work has led to a proposal submitted to NSF: CDI-Type II: Cyber-Enabled Model Development and Discrimination for Complex Chemical and Biological Processes.

5. Prof B. Pijanowski, FORS, Prof B. Waldorf, AG ECON, "Advanced coupled climate-land-society modeling at Purdue: Using the Teragrid to move scientific "Light Years" ahead".

Related to this work, the PI's submitted a NASA fellowship proposal for the involved student's work using the Watu model. They proposed to couple it to RAMS and the LTM. This is exactly what the model is being designed to do. Another proposal was submitted to NASA that focuses on using the model to examine the impacts of monsoons on migration patterns in the South Pacific. It was written and submitted by their postdoc Deepak Ray.

6. Bob Skeel, CS, Carol Post, MCMP, "Computing transition pathways for biomolecular systems".

7. Alejandro Strachan, MSE, Gerhard Klimeck, ECE/NCN, "Predictive modeling of the atomic level structure and electronic properties of semiconductor nanostructures" .

8. T. N. Vijaykumar, ECE, C. Bouman, ECE/BME, T. Talavage, ECE/BME, "Scalable parallelization of high-performance applications in medical imaging".

4.2 Projects with Direct Involvement and Cost Sharing

Sections 4.3 and 4.4 will list Purdue projects in HPC. These projects are conducted by the HPC Community described in Section 2.2. The following list includes projects in which CRI staff was directly involved.

- 3-D Nanoelectronic Modeling, Faculty: Gerhard Klimeck, ECE
- SPIKE: A parallel banded linear solver, Faculty: Ahmed Sameh, CS
- HPC Benchmarking using Real Applications, Faculty: Rudi Eigenmann, CRI, ECE
- Numerical Nuclear Reactor, Faculty: Prof. Thomas Downar, NE
- 3D Structure Reconstruction of Viruses using PFT, Faculty: Prof. Michael Rossmann, Biological Sciences
- LTM (Land Transformation Modeling), Faculty: Brian Pijanowsky, Forestry
- GEMS project,Faculty: Prof. Charles Merkle, AAE
- Condor simulations to evaluate BMP scenarios considering stochasticity of weather, Faculty: Indrajeet Chaubey

CRI was also able to provide cost sharing on a number of proposals. Of those proposals, the project "Accelerating nano-scale Transistor Innovation through Petascale Simulation", (details given in Section 4.6) was funded through NSF's PetaApps Program.

CRI also sponsored a number of travel activities to support faculty project presentations, discussions of collaborations, and proposal meetings.
Travel grants to the SC'07 conference will be listed in section 6.2. Another travel grant went to Professor Steve Frankel, who attended meetings at IBM and Argonne National Labs to prepare for submitting a Petascale proposal. Jointly with the Cyber Center and RCAC, CRI staff also participated in a visit to the National Science Foundation to meet with funding officers.

4.3 HPC Projects at Purdue: Computational Applications

The following is a list of Purdue efforts that develop and use computational applications for advanced research projects.

Quiet Supersonic Jet Engine Analysis

Computational aerodynamic and aeroacoustic analysis of a quiet supersonic business jet, focusing on the inlet and nozzle aerodynamics, and acoustics of the fan and nozzle.

Researchers: A.Lyrintzis and G. Blaisdell
http://cobweb.ecn.purdue.edu/~lyrintzi/

High-Speed CFD

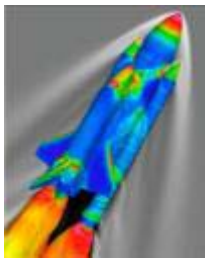

Computational predictions of turbulent heat transfers in high speed boundary layer flows.

Researchers: G. Blaisdell and A. Lyrintzis
http://cobweb.ecn.purdue.edu/~blaisdel/

Large Eddy Simulation of Jet Noise

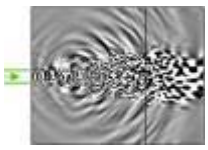

We are looking at analyzing jet noise reduction using large-eddy simulations. Our 3-D large eddy simulation (LES) code is part of a Computational Aeroacoustics (CAA) methodology that couples surface integral acoustics techniques such as Kirchhoff's method and the Ffowcs Williams-- Hawkings method with LES for the far field noise estimation of turbulent jets. The LES code employs high-order accurate compact differencing together with implicit spatial filtering and state-of-the-art non-reflecting boundary conditions.

Researchers: P.-T. Lew, S. C. Lo, G. A. Blaisdell, and A. S. Lyrintzis
http://cobweb.ecn.purdue.edu/~blaisdel/CRI_0607/CRI_0607.htm

High Altitude Plumes

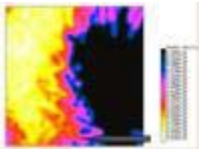

High-altitude plume-atmosphere interactions produces complex three-dimensional chemically reacting flows.

Researcher: Alina Alexeenko
http://web.ics.purdue.edu/~alexeenk/

Computational Multi-Physics

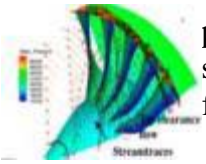

The group led by Prof. Merkle is involved in Multi-scale and Multi-physics Computational fluid dynamics (CFD). The group is working on several projects. Please visit our website for detailed information. The figure shows the air flow around the blades of a jet engine fan.

Researchers: Charles L. Merkle, Ding Li, Guoping Xia et.al
http://tspcpc110.ecn.purdue.edu/

Distributed Capillary Fluids Modeling

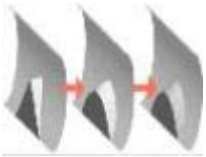

Our groups research focus is in low gravity and MEMS 3-D real-world capillary fluids problems. We use the freely available powerful "Surface Evolver" code for the computational simulations.

Researcher: S. H. Collicott
http://www.ufluids.net/

Spacecraft Mission Design and Analysis

The spacecraft mission design and analysis research with very large data sets and manifold structures includes 1. optimal design and analysis of spacecraft missions including trajectory design, 2) missions involving space-based interferometry: including searches for exosolar planets, and 3) identification of black holes, and spectral characterization of distant stars.

Researcher: K.C. Howell
https://engineering.purdue.edu/people/kathleen.howell.1/

Analysis of Pharmaceutical Coating Process

Our goal is to improve pharmaceutical coating process critical to pharmaceutical manufacturing and commercial drug product performance.

Researcher: Carl Wassgren
http://www.purdue.edu/dp/cam

Product Life Cycle Management

Often new products raise serious nterdisciplinary problems, combining electrical, mechanical, and embedded software elements and subsystems, necessitating an integration of diverse domains of knowledge and interoperations of diverse software systems. Our research focus is to address these complex issues using interdisciplinary PLM research, identifying business value and total system integration research.

Researcher: Craig Miller
http://www.purdue.edu/dp/cam

Laser Assisted Manufacturing

The goal is to develop critical or novel technologies in laser-based manufacturing and materials manufacturing.

Researcher: Yung Shin
http://www.purdue.edu/dp/cam

Extreme hot events in the 21st century

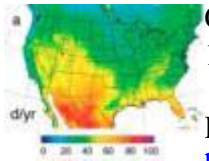

Changes in extreme hot events in the 21st century show an increase of 100 to 560% in occurrence of the hottest temperatures.

Researcher: Noah Diffenbaugh
http://www.purdue.edu/eas/people/faculty/diffenbaugh.html

What is Hestia?

High resolution, coupled models for global quantification of carbon dioxide emissions.

Researcher: Kevin Gurney
http://www.purdue.edu/eas/carbon/hestia.html

Ecological Acoustics Research (EAR) - A new science of sound

This research is in the area of coupled natural-human systems through the development and use of spatial models. We take a multidisciplinary approach to studying natural resource problems, combining expertise in GIS, remote sensing, simulation modeling and natural resource management.

Researcher: Bryan C. Pijanowski
http://www.human-environment.org

Seismic Migration Imaging

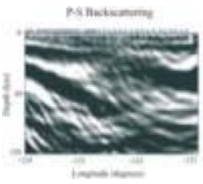

Seismic receiver function array imaging of the subduction zone in the Pacific Northwest beneath Oregon. Figure shows the ray migration imaging using a single teleseismic event.

Researcher: Robert Nowack
http://web.ics.purdue.edu/~nowack/

Computational Methods for Biomolecular Simulation

Fastest N-body solvers, long-time integrators, absolute free energy differences, transition pathways

Researcher: Robert D. Skeel
http://bionum.cs.purdue.edu/

Electron Cryo-Microscopy & 3-D Reconstruction

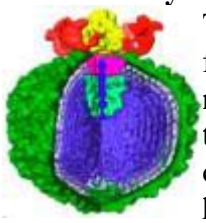

The research focuses on the development of cryo-EM technique to push for near atomic resolution and high throughput single particle 3-D reconstructions and to eventually transform this technique into a routine tool for functional studies of biological systems. To achieve these goals, our research involves development of new image processing algorithms, high performance computing, data collection automation and reliable sample preparation.

Researcher: Wen Jiang
http://jiang.bio.purdue.edu

Bioinformatics Infrastructure for Proteomics Analyses

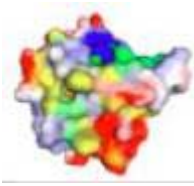

Our research focuses on - Protein function beyond BLAST search, Protein 3D structure prediction, fast protein 3D structure search and protein-protein docking.

Researcher: Daisuke Kihara
http://dragon.bio.purdue.edu

Pharmaceutical Informatics

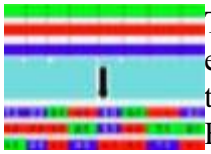

The pharmaceutical informatics research thrust area in our group aims to establish new mathematical, informatics and computational foundations to broaden the technological base for harvesting the fruits of genomics. For more information read this essay on pharmaceutical informatics - http://meweb.ecn.purdue.edu/~kim55/Pharma-essay.shtml

Researchers: Sangtae Kim and Venkat Venkatasubramanian
http://www.ipph.purdue.edu/faculty/?uid=venkatp

Multi-Scale Materials Modeling

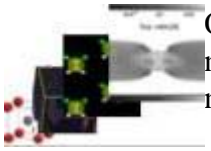

Our research group uses multi-scale modeling to predict behavior of materials from first principles and understand the fundamental mechanisms that govern materials behavior.

Researcher: Alejandro Strachan
https://engineering.purdue.edu/MSE/People/ptProfile?id=33239

3-D Nanoelectronic Modeling

The nanoelectronic modeling of multi-million atom electronic structure calculations using state of the art NEMO 3-D application is carried out in our group. It can calculate eigenstates in (almost) arbitrarily shaped semiconductor structures in the typical column IV and III-V materials.

Researcher: Gerhard Klimeck
http://cobweb.ecn.purdue.edu/~gekco/nemo3D/

CMS - Physics on a Global Scale

The Compact Muon Solenoid Detector (CMS) of the Large Hadron Collider (under construction at CERN) will enable new physics in the area of high energy physics. The CMS center at Purdue is a tier 2 center providing storage for the large amounts of data that will be generated from these experiments and also provide other services (including computing) to other researchers.

Researcher: Norbert Neumeister
http://neumeist.web.cern.ch/neumeist/

Mass Transfer Processes in Contaminated Heterogeneous Soils

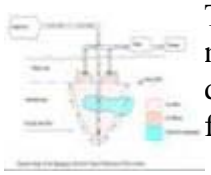

The goal of this project is to develop suitable tools to help understand mass transfer processes during air advective air flux, in order to aid in the design and operation of remediation systems that involve advective air flux such air sparging and Soil vapor Extraction (AS/SVE).

Researcher: Rabi H. Mohtar
http://cobweb.ecn.purdue.edu/~mohtar/

Novel Methods for Computer-Aided Drug Discovery

Our research aim is to develop novel computational concepts to identify feasible binding modes of protein-bound molecules and to compute their binding affinity. Our ongoing concept development is focused on the aspects of protein flexibility, solvation, and entropy.
Researcher: Markus A. Lill
http://people.pharmacy.purdue.edu/~mlill/

Ecosystems and Biogeochemical Dynamics

The Ecosystems and Biogeochemical Dynamics group is studying the interactions among atmosphere, biosphere, and human dimension in the context of climate change, chemical element cycles, and policy-making. To study these interactions, the data of biosphere and atmosphere obtained from field and in-situ measurements and satellite observations are fused with the numerical models of the ecosystems and biogeochemistry and the atmosphere.

Researcher: Qianlai Zhuang
http://www.purdue.edu/eas/ebdl/

4.4 HPC Projects at Purdue: Computer Systems Research

The following research projects aim at creating the next generation of parallel and high-end computer hardware and software systems, capable of executing computational applications efficiently.

HPC Architecture/Interconnection Network

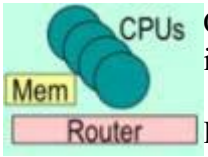

Our research group looks at addressing computation and communication issues in modern HPC architectures and interconnection networks.

Researcher: Mithuna Thottethodi
http://cobweb.ecn.purdue.edu/~mithuna/

Harnessing Unused Disk Space

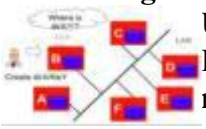

USFS is a peer-to-peer (p2p) enhancement for the widely-used Network File System (NFS). USFS harvests redundant storage space on cluster nodes and user desktops to provide a reliable, shared file system that acts as a large storage with normal NFS semantics

Researcher: Y. Charlie Hu
http://www.ece.purdue.edu/~ychu/usfs

Auto-tuning Compilers

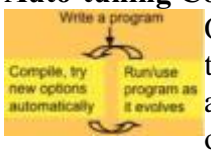

Over a number of years our research group has been working on auto-tuning compilers with the goal of making programs run optimally on wide array of architectures by optimally searching a huge space of optimization options, dynamically plug-in new code while ensuring the evolving program improves.

Researchers: Rudi Eigenmann, Sam Midkiff
http://www.ece.purdue.edu/ParaMount/

Why Put Up with Failures on the Grid?

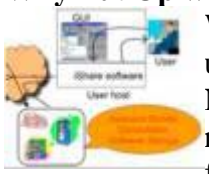

We are looking at reliability issues on grid. Our solution to these issues uses techniques such as predicting failure (using Semi-Markov model and Neural network), making scheduling decisions, perform runtime monitoring and Incremental failure prediction, migrate processes from failure-prone hosts using checkpoints

Researcher: Rudi Eigenmann, Saurabh Bagchi, Hugh Hillhouse
http://www.ece.purdue.edu/~dcsl/Projects/grid_reliability

Exploiting CMP's Latency-Capacity Tradeoff

The goals of this project is to exploit chip multi-processor (CMP) cores by exploiting the tradeoff of latency and capacity. With more CMP cores we need both large capacity and fast access from on-chip caches.

Researcher: T. N. Vijaykumar, Zeshan Chishti
http://www.ece.purdue.edu/~vijay

SPIKE: A parallel banded linear solver

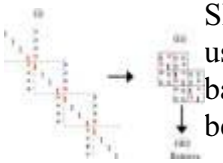

SPIKE is a robust hybrid parallel banded linear system solver that can be used as a direct solver or pre-conditioner for outer iterative scheme. The banded system could be dense or sparse. We have seen performance better than Scalapack package in different architectures.

Researchers: Ahmed Sameh, Eric Polizzi, Mohamed Sayeed, Faisal Saied
http://www.cs.purdue.edu/people/faculty/sameh/

HPC Benchmarking using Real Applications

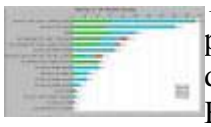

The goal of this project is to address challenging questions related to performance of real application benchmarks such as SPEC MPI 2007 on different high performance computing platforms such as Cray-XT4, IBM Blue Gene, SGI Altix, IBM JS21 linux cluster etc., We use these benchmarks for benchmarking and performance characterization of new architectures. We also maintain a rank list based on such real application benchmarks.

Researchers: M. Sayeed, Hansang Baeh, Faisal Saied, Rudi Eigenmann
http://www.purdue.edu/taplist/

4.5 Super Projects

The CRI "Superproject" initiative aims at identifying project groups that conduct simulation-based research and pursue common goals. By creating an umbrella project encompassing these projects, the effort expects to create synergy, resulting in accelerated progress and increased competitiveness in grant competitions. By combining multiple computational projects, the effort also aims at creating multi-system simulators that can exploit some of the largest HPC platforms. These project initiatives are still at the beginning, but a short description is warranted.

Two superprojects were initiated in 2007. The goal of the first project combines simulation-based research efforts that investigate parts of the human body. The long-term vision of the umbrella project is to create a system capable of simulating the entire human anatomy. The project is expected to have many applications in science and medicine.

Another umbrella project aims at combining simulation-based efforts that can contribute to virtual expeditions to other planets. The umbrella project aims to combine research efforts, such as the study of propulsion systems, exploration of atmoshperic phenomena, computation of trajectories taken by space crafts and planets, and the study of biological processes on other planets.

4.6 Proposals

The following list includes proposals with direct involvement of CRI staff and those related to CRI seed grants. This list represents a small part of the funding activity of Purdue's HPC community.

Awarded:

CSR-AES: Adaptive Optimization for Dynamically Discovered Hardware and Software Sources, $300,000, P.I. Rudi Eigenmann, ECE

CRI: CRD-Supporting the Cetus Compiler Infrastructure for the Community, $475,000, P.I. Rudi Eigenmann, ECE

PRISM: Center for Prediction of Reliability, Integrity and Survivability of Microsystems, DoE PSSAP proposal, over $16 million, P.I. Jayathi Murthy, ME

"Accelerating nano-scale Transistor Innovation through Petascale Simulation" NSF PetaApps proposal, $1,599,205, P.I. Gerhard Klimeck, ECE

“CRI: IAD: Accelerator-based High-Performance Computing” NSF, 570,000, P.I.Vijay Pai, ECE

The DAC (Development Allocation Committee) proposal submitted to NCSA/Teragrid for compute time on a linux cluster was accepted. This was to do performance analysis

and modeling of 30,000 service units and computer time SPEC MPI2007 benchmarks that was released recently in June 2007.
P.I. M. Sayeed, CRI/ITaP

In addition to the awarded proposals, CRI has participated in over $15 million of submitted proposals. The following are some of the proposed titles:

- An Expedition in NICA Computing.
- Knowledge Driven Watershed Modeling to Enhance Research, Education, and Decision Making.
- High-Fidelity Hybrid Finite Element-Discrete Element Method Modeling.
- Understanding Thunderstorm Evolution around Urban Areas.
- A Virtual Organization for Data Discovery, Composition, and Prototyping of Data-analysis Algorithms
- Performance Engineering for Petascale Applications
- Extending Purdue's BoilerGrid to a Broad-based Community through Acquisition of a Cluster Pool for Parameter-sweep Applications

5. Fostering Partnerships

CRI activities have included the creation of research partnerships and collaborations with industrial partners. A first, highly successful event of the "Purdue-Industry HPC workshop" was held in March 2007. Among the activities that facilitated research partnerships were the CRI seminar speaker program, which brought potential partners to Purdue, where they met faculty in their area of expertise (see section 3.2). CRI also organized a series of "futures meetings" at the Supercomputing 2007 conference and sponsored faculty travel to participate in the workshops that may lead to future collaborations.

Purdue-Industry HPC Workshop

The Purdue-Industry HPC Workshop, held on March 20, 2007 at the Burton Morgan Entrepreneurship Center of Purdue's Discovery Park, brought together experts and users of high-performance computing technology from industry and from Purdue.

The response to the workshop exceeded expectations. The capacity of the workshop venue, holding 80 attendees, had already filled before the workshop begun. Most attendees stayed throughout all morning talks and the luncheon presentation; the afternoon roundtable discussions similarly drew a large number of participants, filling the breakout rooms to capacity. Several industry sponsors supported the event.

The events started with a reception on the evening before the workshop, at which the Purdue Provost, Dr. Sally Mason, welcomed the guests and commented on the relevance of high-performance computing technology and collaborations between industry and Purdue.

The actual workshop program was opened by Professor Rudi Eigenmann, Interim Director of the Computing Research Institute, followed by welcoming remarks of Gerry McCartney, Interim Vice President of Information Technology and Professor Bob Bernhard, Associate Vice President for Research.

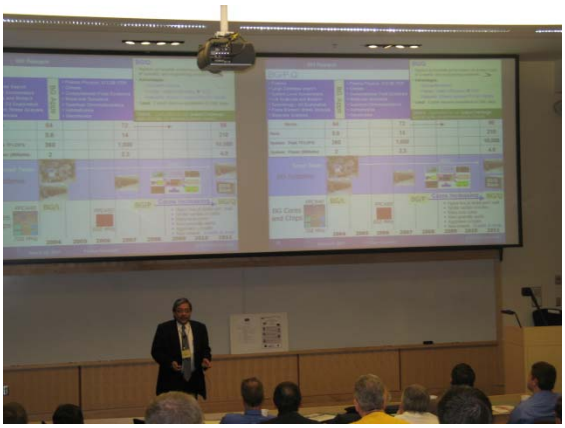

Dr. Tilak Agerwala, Vice President, Systems, IBM Research, gave the keynote presentation, entitled "High Performance Computing, Innovation, and Accelerating Discovery". A second invited speaker from industry was Steve Kirsch, Raytheon Space and Airborne Systems, giving a talk on "Multi-core Processors: The Dream of Infinite Processing Power or the Nightmare of Reality."

Two talks from Purdue faculty completed the morning program. Rudi Eigenmann, Professor of Electrical and Computer Engineering, talked about "High Performance Computing Going Mainstream" and Tasos Lyrintzis, Professor of Aeronautics and Astronautics Engineering, presented "Examples of High Performance Computing in Aerospace Engineering."

A luncheon address was given by Sangtae Kim, Professor of Chemical Engineering and Mechanical Engineering, entitled "Pharmaceutical Informatics and a Pathway to Personalized Medicines in the Peta-Scale Era."

In afternoon breakout sessions the attendees discussed the four topics "State of the Art of HPC", "Future of high-end computing", "Future of multicore", and "HPC Education." The workshop concluded with summaries of these sessions.

After-workshop events included a demo by Gerhard Klimeck, Professor of Electrical and Computer Engineering on "nanoHUB.org - serving thousands of researchers and educators with interactive HPC-based simulations in nanotechnology" and a tour to the Discovery Park's Birck Nanotechnology.

5.2 Futures Meetings

Collaboration meetings at the SC'07 conference.
CRI organized a series of Futures Meetings at the Purdue booth with representatives of the following:

1. Universities
    a. Louisiana State University
    b. University of Oregon
    c. University of Illinois National Center for Supercomputing Applications
    d. University of Notre Dame
2. National Labs and Centers
    a. Lawerence Berkeley National Lab
    b. Oak Ridge National Labs
    c. Argonne National Labs
    d. Artic Region Supercomputer Center
3. Funding Agencies
    a. National Science Foundation
    b. Department of Energy
4. Industry
    a. IBM Inc.
    b. Si Cortex Inc.
    c. Cray Inc.
    d. SGI Inc.

5.3 Workshops Participation

CRI staff attended a number of conferences and workshops in areas that are closely related to the mission of CRI.

- DoE Townhall meeting on DoE Exascale Computing Argonne – May 31, June 01, 2007 (Rudi Eigenmann, Faisal Saied).
- Great Lakes Consortium workshop – July 13 – Northwestern University (Faisal Saied).
- SciDAC 2007 conference in Boston (M. Sayeed).
- SC'07 – Reno – Nov 12-16, 2008 (Rudi Eigenmann, Faisal Saied, M. Sayeed).
- NSF Teragrid Petascale Workshop at Tempe AZ - Dec 11-12, 2007 (Faisal Saied).

6. Communicating Successes in HPC

The Web presence at www.cri.purdue.edu is the primary means for communicating CRI's mission, activities, and announcements to a broad audience. Through presentations at Purdue departments and centers, CRI presented its mission and services to Purdue faculty and staff. CRI also engaged in a "new faculty program". At an international level, Purdue's booth at the annual "Supercomputing" conference offered a window into CRI activities.

6.1 CRI Web Presence

In 2007, CRI's web site was updated with a number of areas. The "HPC Community" now lists the approximately 200 faculty members that conduct research related to CRI's mission. This community replaces the previous "CRI advisory board", which included some 30 core members of the community.

A projects web page now offers an overview of the large number of HPC- related projects. For each project, the web page gives a brief description, lists primary faculty contacts, and links to the project web pages at the home departments.

Workshops, luncheons, and poster sessions, are now listed as events and workshops on CRI's Web pages. They provide a view of future activities as well as a record of past events.

A news items category has been added, listing notable successes and other important developments.

6.2 SC'07 Participation and Functions

Purdue University had a strong presence in the SC'07 Conference in Reno, November 2007. SC07 is the premier international conference on high performance computing, networking, storage, and analysis. The SC conference series has grown to include scientists, researchers, software developers, network engineers, policy makers, corporate managers, CIOs and IT administrators from universities, industry and government from all over the world.

A number of units on campus contributed to the success of the Purdue booth. ITaP took the lead in organizing the booth. CRI contributed to shaping the message conveyed at the booth and increasing the effectiveness of the Purdue presence in several ways.

CRI funded a number of Purdue faculty members to travel to Reno and make presentations in the Purdue booth. This helped raise the visibility of HPC and cyberinfrastructure research at Purdue. The CRI funded presenters at SC07 were

- Norbert Neumeister, Physics, "The CMS Tier-2 Center at Purdue"

- Kevin Gurney, Earth and Atmospheric Sciences, “Hestia: The Global Metabolism of Greenhouse Gas Emissions”
- Rudi Eigenmann, Electrical and Computer Engineering, “Simulating Where No One Has Gone Before”
- Sam Midkiff, Electrical and Computer Engineering , “Aspen: A Language for Highly Concurrent Network Applications”
- Venkatesh Merwade, Civil Engineering “Cyberinfrastructure for Environmental Data and Modeling”

CRI helped increase the number of Purdue research projects showcased at the Purdue booth at SC07 by encouraging researchers who had participated in CRI’s on-campus Poster Session to contribute the posters for display at the Purdue booth. As a result, conference attendees were able to get a good sense of the breadth and depth of computing research at Purdue.

CRI was especially pro-active in using the booth to schedule meetings with leaders in the field who were attending SC07. These included Ed Seidel (LSU), Thom Dunning (NCSA), Horst Simon (NERSC), Al Maloney (U Oregon), Pat Worley (Oak Ridge).

7. Publications

The list of publications of the Purdue HPC community is post on the CRI website at www.cri.purdue.edu. It is available upon request as a separate report.