Clusters and Computational Grids for Scientific Computing

CCGSC 2006

September 10-13, 2006

Highland Lake Inn

Flat Rock, North Carolina

Sponsored by:

Message from the Program Chairs

This proceeding gathers information about the participants of the Workshop on Clusters and Computational Grids for Scientific Computing that was held at Highland Lake Inn at Flat Rock, North Carolina on September 10-13, 2006. This workshop is a continuation of a series of workshops started in 1992 entitled Workshop on Environments and Tools for Parallel Scientific Computing. These workshops have been held every two years and alternate between the U.S. and France. The purpose of this the workshop, which is by invitation only, is to evaluate the state-of-the-art and future trends for cluster computing and the use of computational grids for scientific computing. This workshop addresses a number of themes for developing and using both cluster and computational grids. In particular, the talks covered:

- Message Passing and High Speed networks
- Grid Infrastructure and Programming Environment
- Cluster Based Computing and Tools
- Languages, Libraries, and Fault Tolerance.

Speakers presented their research in the above four areas and interacted and worked with all the participants on the future software technologies that will provide for easier use of parallel computers. This workshop was made possible thanks to sponsorship from Intel, Sun Microsystems Inc., Microsoft Inc., Myricom Inc., and HP, with the scientific support of INRIA (French National Institute for Research in Computer Science), the Computer Science Departments of the University of Tennessee in Knoxville (UTK) and University Claude Bernard (UCB-Lyon 1).

Jack Dongarra, Knoxville University, Tennessee, USA.
Bernard Tourancheau, Lyon, France
Frédéric Desprez, INRIA, ENS Lyon, France
Laurent Lefevre, INRIA, ENS Lyon, France

Organized by:

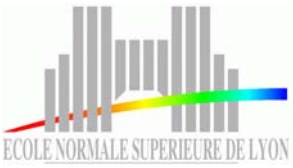

With sponsored from:

Sunday
September 10, 2006

Introduction and Welcome
Jack Dongarra, ICL, CS, U of Tennessee
Bernard Tourancheau, University of Lyon

Sunday		
5:30 – 7:30	Session Chair: Jack Dongarra, U Tennessee	Grand Ole Hall (5 talks - 20 minute each)
5:30	Giri Chukkapalli, Sun	SUN's Petascale Strategy
5:50	Dan Fay, Microsoft	Mystery and Suspense
6:20	Patrick Geoffray, Myricom	Trends in High-Performance Networking: The Good, The Bad And The Very Ugly
6:40	Rob Schreiber, HP	Parallel Programming Research at HP Labs
7:00	David Scott, Intel	HPC @ Intel: Platforms and the Technologies That Go Into Them
8:00 pm - 10:00 pm	Dinner	Season's Restaurant
9:30 pm –	Fire Pit and/or Refreshments in Woodward House	

Monday
September 11, 2006

Monday		
7:30 - 8:30	Breakfast	Season's Restaurant
8:30 - 10:30	Session Chair: Bernard Tourancheau	Grand Ole Hall (4 talks – 30 minutes each)
8:30	Fran Berman, UCSD/SDSC	Beyond Branscomb
9:00	Tony Hey, Microsoft	e-Science and Cyberinfratructure
9:30	Pete Beckman, ANL/UC	Urgent Computing, Sharing Grid Resources, and Elastic Computing
10:00	Joel Saltz, OSU	Architecture of a Strongly Typed Grid and Lessons Learned from the Cancer Research Community
10:30 -11:00	Coffee/Cookies	Grand Ole Hall
11:00 - 1:00	Session Chair: Graham Fagg, U Tennessee	Grand Ole Hall (4 talks – 30 minutes each)
11:00	Franck Cappello, U of Paris	An Update of Grid'5000 and a Focus on a Fault Tolerant MPI Experiment
11:30	Rich Graham, LANL	Open MPI - A High Performance Fault Tolerant MPI Library
12:00	Rusty Lusk, ANL	DARPA's High Productivity Language Systems Project: One Language To Rule Them All, or ADA Strikes Back?

12:30	Thomas Sterling, LSU/ORNL/Caltech	*ParalleX*: Towards a New Parallel Execution Model for Scalable Programming and Architectures
1:00 - 2:00	Lunch	Season's Restaurant
2:00 – 4:30	Free time	
4:30 – 5:00	Coffee/Cookies	Grand Ole Hall
5:00 – 5:30	Frederica Darema, NSF	Clusters, Computational Grids, and Beyond
5:30 - 7:30	Panel Chair: Bill Gropp, ANL	Grand Ole Hall (Panel Session)
	Geoffrey Fox, IU	"Three Questions You Should Ask" - Quantitative measures for improving the users' experience in using clusters and grids
	John Morrison, UCC	
	Phil Papadopoulos, SDSC	
	Jeff Hollingsworth, UM	
8:00 – 10:00	Dinner	Season's Restaurant
9:30 pm -	Fire Pit and/or Refreshments in Woodward House	

Tuesday
September 12, 2006

Tuesday		
7:30 - 8:30	Breakfast	Season's Restaurant
8:30 - 10:30	Session Chair: Mark Baker, Reading U	Grand Ole Hall (4 talks – 30 minutes each)
8:30	Ken Kennedy, Rice U	The Role of Performance Models in Efficient Grid Scheduling
9:00	Laurent Lefevre, INRIA, ENS Lyon	Challenges In Designing High Performance Autonomic Gateways In Large Scale Grids And Distributed Environments
9:30	Philippe D'anfray, Renater	Renater Dark Fiber Architecture: User's Feedback
10:00	Yves Robert, ENS, Lyon	Matrix Product/LU with Limited Memory
10:30 -11:00	Coffee/Cookies	Grand Ole Hall
11:00 - 1:00	Session Chair: Jean-Yves Berthou, EDF	Grand Ole Hall (4 talks – 30 minutes each)
11:00	Satoshi Matsuoka, TiTec	The Hatching of TSUBAME and Its Siblings --- One Possible Future Scenario for the Japanese Petascale Grid in the 2010s ---
11:30	Miron Livny, U Wisc	Submitting Locally And Running Globally - The GLOW And OSG Experience
12:00	Frédéric Desprez, INRIA, ENS Lyon	Scheduling for Network Enabled Servers Systems
12:30	Jim Plank, UTK	Erasure Coding Research for Reliable Distributed and Cluster Computing

1:00 - 2:00	Lunch	Season's Restaurant
2:00 – 4:00	Free time	
4:00 – 5:00	Coffee/Cookies	Grand Ole Hall
5:00 - 7:30	Session Chair: Thomas Ludwig, Heidelberg	Grand Ole Hall (5 talks – 30 minutes each)
5:00	Wuchun Feng, VT	Global Climate Warming? Yes ... in the Machine Room
5:30	Al Geist, ORNL	Counting on Failure
6:00	Dieter Kranzlmuller, Kepler U	Interactive Videostreaming Visualization on Clusters and Grids
6:30	Andrew Lumsdaine, IU	High-Performance Distributed Memory Graph Computations
7:00	Bart Miller, U Wisc	Scalable Tool Design for Large-Scale Applications
8:00 – 10:00	Dinner	Season's Restaurant
9:30 pm -	Fire Pit and/or Refreshments in Woodward House	

Wednesday
September 13, 2006

Wednesday		
7:30 - 8:30	Breakfast	Season's Restaurant
8:30 - 10:30	Session Chair: Vaidy Sunderam, Emory	Grand Ole Hall (4 talks – 30 minutes each)
8:30	Thilo Kielmann, Vrije U	Grid Programming Models: Requirements and Approaches
9:00	Michel Cosnard, INRIA	Resource Discovery in the Arigatoni Overlay Network
9:30	Rich Wolski, UCSB	Predicting Bounds on the Batch Queuing Delay Experienced by User Jobs in Real Time
10:00	Hans Zima, Caltech	High-Level Programming Models for Clusters: Issues and Challenges
10:30 -11:00	Coffee/Cookies	Grand Ole Hall
11:00 - 1:00	Session Chair: Micah Beck U Tennessee	Grand Ole Hall (4 talks – 30 minutes each)
11:00	David Walker, Cardiff U	Portal Interfaces for Grid-Based Scientific Computing
11:30	Chuck Seitz, Myricom	Ethernet Will Eradicate Ethernot
12:00	Jeff Vetter, ORNL	Application Accelerators: Deus ex Machina?
12:30	Emmanuel Jeannot, Loria INRIA-Lorraine	Improved Scheduling Strategies for Agent-Client-Server Middleware
1:00 - 2:00	Lunch	Season's Restaurant
2:00	Depart	

Talks, Titles, and Abstracts

Pete Beckman
Urgent Computing, Sharing Grid Resources, and Elastic Computing

High-performance modeling and simulation are playing a driving role in decision making and prediction. For time-critical emergency support applications such as severe weather prediction, flood modeling, and influenza modeling, late results can be useless. With HPC computing and distributed Grid resources becoming more widely available, the community can build global systems and policy frameworks for utilizing high-end computational resources in support of emergency computation. A specialized infrastructure is needed to provide computing resources quickly, automatically, and reliably. SPRUCE is a system to support urgent or event-driven computing on both traditional supercomputers and distributed Grids. Currently, SPRUCE is deployed at several large supercomputer centers. With the rapid growth of the web services industry, companies such as Amazon are beginning to provide virtual machines on demand, for an hour at a time. How will services such as Amazon's Elastic Computing Cloud (EC2) change Grids and community resources? Could it become important for Urgent Computing?

Fran Berman
Beyond Branscomb

In 1993, the Branscomb Pyramid provided a conceptual framework for thinking about high performance computing. Although the magnitudes of each level are different today, the framework has continued to provide a convenient and elegant way of thinking about the scale, costs, and use of HPC infrastructure.

In this talk, we expand the Branscomb organizational structure to other kinds of infrastructure, particularly data. We discuss the challenges and opportunities for leadership and competitiveness at all levels of the Branscomb Pyramid. We discuss the need for a paradigm shift in business models which support research as the use of enabling technologies and infrastructure becomes more and more a ubiquitous part of the research and education landscape.

Giri Chukkapalli
SUN's Petascale strategy

Franck Cappello
An Update of Grid'5000 and a Focus on a Fault Tolerant MPI Experiment

Grid'5000 is a nation-wide infrastructure for computer science research in Grid and P2P computing. We shortly recall the motivations, design, architecture of Grid'5000 and present update information as well as some early results obtained with this platform. Then we describe an example of experiment on a fault tolerant MPI implementation, comparing at large scale two fault tolerant protocols: blocking and non blocking coordinated checkpointing.

Michel Cosnard
Resource Discovery in the Arigatoni Overlay Network

with Raphael Chand, Luigi Liquori INRIA, France

Arigatoni is a lightweight Overlay Network for dynamic and generic Resource Discovery. Entities in Arigatoni are organized in Colonies. A colony is a simple virtual organization composed by exactly one leader, offering some broker-like services, and some set of Individuals. Individuals are subcolonies of individuals, or basic units called Global Computers. Global computers communicate by first registering to the colony and then by mutually asking and offering services. The leader, called Global Broker, has the job to analyze service requests/responses coming from its own colony or arriving from a surrounding colony, and to route requests/responses to other individuals. After this discovery phase, individuals get in touch with each others without any further intervention from the system, typically in a P2P fashion. Communications over the behavioral units of the overlay network are performed by a simple Global Internet Protocol. Arigatoni provides fully decentralized, asynchronous and scalable resource discovery, that can be used for various purposes from P2P applications to more sophisticated Grid applications. The main focus of this paper is to present the resource discovery algorithm used in Arigatoni, that is reminiscent to some algorithms employed in the publish/subscribe paradigm. We show some simulations that show that resource discovery in Arigatoni is efficient and scalable.

Philippe D'anfray
Renater Dark Fiber Architecture: User's Feedback

The dark fiber "project architecture" has been deployed in 2006 by RENATER to support research projects with high network resources requirements. Quantifying and measuring user’s needs would help to design the next generation of network infrastructure. This implies to set up some "application network benchmark" and exhibit relevant related metrics.

Frederica Darema
Clusters, Computational Grids, and Beyond

The talk will address research and technology advances for optimized and dependable execution in large scale heterogeneous computing environments. Applications in nearly all sectors, scientific, engineering, and commercial, are becoming more encompassing and new concepts, such as the Dynamic Data Driven Applications Systems (DDDAS), create new capabilities in physical and engineered systems, but also new challenges and new opportunities in terms of the systems software required to support such environments. Efficient and effective development of emerging applications and computational environments, optimized use of the computational resources, and guaranteeing quality of service and dependability at all layers of the computational system, requires systems software advances, such as in programming environments, application composition systems, advanced compiler technologies for optimized application mapping and dynamic runtime support, debugging and check-pointing methods, and performance-engineered hardware and software capabilities at all layers. Such research opportunities will be discussed together with the overarching consideration, that these advances need to be made in a synergistic and integrated manner, taking a systems-view in developing these enabling technologies, rather than advancing each of the individual technologies in an isolated manner.

Frédéric Desprez
Scheduling for Network Enabled Servers Systems

Dan Fay
Mystery and Suspense

Wu Feng
Global Climate Warming? Yes ... In the Machine Room

For decades now, the notion of "performance" has been synonymous with "speed." Unfortunately, this particular focus has led to the emergence of supercomputers that consume massive amounts of electrical power and produce so much heat that extravagant cooling facilities must be built to ensure proper operation. In addition, the emphasis on speed as the performance metric has adversely affected other performance metrics, e.g., reliability. As a consequence, all of the above has contributed to a marked increase in the total cost of ownership (TCO) of a supercomputer. Therefore, we espouse the importance of being green in high-performance computing and even argue for a complementary list to the TOP500: The Green500 List.

Geoffrey Fox

Panel on "Three Questions You Should Ask" - Quantitative measures for improving the users' experience in using clusters and grids

Al Geist
Counting on Failure
The best kind of fault is the one that never happens. Fault prevention through active monitoring and dynamic fault prediction is in its infancy, but will have to grow up fast in the next few years as computer facilities, such as Oak Ridge National Lab, install petaflop systems. ORNL's 54 TF system already illustrates that facilities can count on failure as I/O systems, and communication fabrics are scaled up to tens of thousands processors. This talk will explore the options in a world where system faults are a given. Topics include: active monitoring, fault prediction, run through, checkpoint, recalculation, and fault tolerant MPI recovery options.

Patrick Geoffray
Trends in High-Performance Networking: The Good, The Bad And The Very Ugly
High Performance networking is a moving target, trade-offs change constantly. Concepts that looked like good ideas a few years back, such as kernel-bypass and zero-copy, may not be as relevant today. New approaches such as TCP Offload Engine (TOES) are not going anywhere and RDMA-based protocols get a reality check. Let's take a look at where Cluster Computing is going in the next couple of years. One thing is sure, it can't be worse than Grid Computing.

Rich Graham
Open MPI - A High Performance Fault Tolerant MPI Library
Open MPI is a collaborative open-source project developing a high- performance, fault-tolerant, production grade MPI implementation of the MPI standard. This talks with give a overview of the collaboration of some technical aspects of this project, and present some recent performance data.

Bill Gropp
Panel on "Three Questions You Should Ask" - Quantitative measures for improving the users' experience in using clusters and grids

Tony Hey
e-Science and Cyberinfratructure

The Internet was the inspiration of J.C.R.Licklider when he was at the Advanced Research Projects Agency in the 1960's. In those pre-Moore's Law days, Licklider imagined a future in which researchers could access and use computers and data from anywhere in the world. He funded an elite group of Computer Science Departments in the USA – which he called his 'InterGalactic Computing Group' - to explore how to realize his vision. Today, as everyone knows, the killer applications of the Internet were email in the 1970's and Tim Berners-Lee's World Wide Web in the 1990's which was developed initially as a collaboration tool for the particle physics academic community. In the future, frontier research in many fields will increasingly require the collaboration of globally distributed groups of researchers needing access to distributed computing, data resources and support for remote access to expensive, multi-national specialized facilities such as telescopes and accelerators or specialist data archives. There is also a general belief that an important road to innovation will be provided by multi-disciplinary and collaborative research – from systems biology and bio-informatics to earth systems science and chemo-informatics. In the context of science and engineering, this is the 'e-Science' agenda. Robust middleware services will be widely deployed on top of the academic research networks to constitute the necessary 'Cyberinfrastructure' to provide a collaborative research environment for the global academic community. This talk will review the elements of this vision and describe how the scientists and engineers are collaborating with computer scientists and the IT industry to create the new e-Infrastructure. When mature, it is clear that such an infrastructure will support the creation of dynamic 'Virtual Organizations' and collaborative environments for many types of application in both academia and industry. This new Cyberinfrastructure will clearly be of relevance to more than just the research community and will support both the e-learning and digital library communities as well as many business applications. This technology is likely also to change the nature of scientific publication with institutional or subject repositories linked to digital archives containing the primary research data.

Jeff Hollingsworth

Panel on "Three Questions You Should Ask" - Quantitative measures for improving the users' experience in using clusters and grids

Emmanuel Jeannot
Improved Scheduling Strategies for Agent-Client-Server Middleware

Agent-client--server middleware like Ninf, DIET or GridSolve are efficient way to access distributed resources through different kind of networks. For such middleware the

scheduling service is one of the key services to achieve performance. In this talk, we will discuss how we have improved the scheduling strategy of GridSolve. We have tackled three different issues: how to evaluate efficiently the duration of the services, how accurately measuring the communication time between the client and the server and how to simultaneously fulfill each actor (agent, client server) objectives.

Ken Kennedy
The Role of Performance Models in Efficient Grid Scheduling

For the past several years, the *Virtual Grid Application Development (VGrADS) Project*, has been pioneering new software strategies to make the Grid easier to use. One of the key developments has been the introduction of generic schedulers that use performance models as surrogates for workflow step running times, thus permitting scheduling to be carried out in advance. We have demonstrated that in-advance scheduling can reduce the time to complete a workflow by integer factors.

A potential disadvantage of this approach is the difficulty of constructing accurate performance models by hand. The VGrADS project has addressed this by providing tools to automatically construct performance models for an application running on different platforms.

This talk will review the construction and use of performance models within VGrADS and will show how universal performance models can lead to strategies that not only improve the efficiency of individual applications, but also improve the efficiency of resource usage across the entire Grid.

Thilo Kielmann
Grid Programming Models: Requirements and Approaches

For programming a single computer, widely used programming models are object-oriented and component-based programming, a choice driven by their high abstraction level, leading to high programmer productivity. For parallel computers, the winner turned out to be message passing, providing by far not the highest-possible abstraction level, but the closest match between machine architecture and programming model, leading to efficient program execution.

For grids, the race is still open. Here, additional non-functional properties like fault-tolerance, malleability, and platform independence come into play. In this talk, we explore the scope of grid programming problems and argue for a palette of programming abstractions, each suitable for its respective problem domain.

Dieter Kranzlmueller

Interactive Videostreaming Visualization on Clusters and Grids

Problem solving and scientific visualization are tightly coupled, since users' comprehension is needed to enable understanding of scientific content. Consequently, a number of scientific visualization approaches have been proposed, which provide visualization of different data types on workstations or frontends of parallel machines. However, with the increased availability of clusters and grids, the coupling between the users' location and the generation of the output has been loosened. This talk presents a scientific visualization approach using video-streaming technology on clusters and grids. The GVid prototype represents a combination of well-known components to enable the generation and transportation of graphical data from parallel machines via possibly long-distance connections to the users' output device. The talk reports on the latest developments of GVid, it's connection to the Grid Visualization Kernel GVK and the interactivity tool glogin, and the future directions of this research project.

Laurent Lefevre
Challenges in Designing High Performance Autonomic Gateways in Large Scale Grids and Distributed Environments

This talk will describe our current works on the design on network gateways inside Grids for supporting new generation of embedded network services (monitoring, analysis, protocol adaptation...). Autonomic and performances aspects will be described. I will also present some current experimental works on large scale Grid5000 platform and more futuristic scenario for an Inter-Planetary Grid.

Miron Livny
Submitting Locally and Running Globally - The GLOW and OSG Experience

Andrew Lumsdaine
High-Performance Distributed Memory Graph Computations

We present the Parallel Boost Graph Library, a library of high-performance reusable software components for distributed graph computation. Like the sequential Boost Graph Library (BGL) upon which it is based, the Parallel BGL applies the paradigm of generic programming to the domain of graph computations. Emphasizing efficient generic algorithms and the use of concepts to specify the requirements on type parameters, the Parallel BGL also provides flexible supporting data structures such as distributed adjacency lists and external property maps. The generic programming approach simultaneously stresses flexibility and efficiency, resulting in a parallel graph library that can adapt to various data structures and communication models while retaining the

efficiency of equivalent hand-coded programs. Performance data for selected algorithms are provided demonstrating the efficiency and scalability of the Parallel BGL.

Rusty Lusk
DARPA's High Productivity Language Systems Project: One Language To Rule Them All, or ADA Strikes Back?

DARPA is currently supporting IBM, Cray, and Sun to develop "high productivity language systems," to run on the new HPCS machines and elsewhere) and to deliver application developers from the pain and agony of MPI (the last great leap forward in parallel programming methodology :-)). This talk will begin with some of the reasons one might be skeptical of such a project, but will conclude with a positive status report on recent progress in the design and implementation of Suns' Fortress, IBM's X10, and Cray's Chapel. All of these languages introduce interesting for scalable parallel programming. Suggestions on how the HPC community might aid in the birthing of a new, widely used, productive parallel programming environment will be solicited.

Satoshi Matsuoka
The Hatching of TSUBAME and Its Siblings --- One Possible Future Scenario for the Japanese Petascale Grid in the 2010s

The Japanese s-called “Keisoku” or the “10-Petascale” project has recently been receiving considerable attention recently. What has not been reported, however, is that CSTP (Council for Science and Technology Policies), the supreme council for overseeing the overall research activities in Japan, has explicitly recommended that there shall be a CyberScience grid infrastructure that accompanies the said machine, and moreover, that major grid & supercomputing centers will host petascale resources in themselves to populate the overall petascale grid infrastructure in the 2010s. The talk will discuss the future outlook on such an infrastructure, based on our experiences in the NAREGI national Japanese research project, as well as those in TSUBAME, the first machine to beat the Earth Simulator in Japan, and reflecting on its future with its siblings that will likely constitute the overall petascale grid.

Bart Miller
Scalable Tool Design for Large-Scale Applications

I will address the problem of tools for large scale environments. We are especially interested in systems, clusters and distributed collections of 1000's and even 10,000's of nodes. The infrastructure that we have developed to address this problem is called MRNet, the Multicast/Reduction Network. MRNet is a tree-based overlay network (TBON) that allows for efficient request distribution and flexible data reductions.

I will present an overview of the MRNet design, architecture, and computational model and then discuss several of the applications of MRNet. The applications include scalable

automated performance analysis in Paradyn, a vision clustering application and, most recently, an effort to develop our first petascale tool, a scalable stack trace analyzer for the IBM BG/L system**.**

John Morrison

Panel on "Three Questions You Should Ask" - Quantitative measures for improving the users' experience in using clusters and grids

Phil Papadopoulos

Panel on "Three Questions You Should Ask" - Quantitative measures for improving the users' experience in using clusters and grids

Jim Plank

Erasure Coding Research for Reliable Distributed and Cluster Computing

Unsurprisingly, this talk will concern erasure codes for storage applications. Extant for decades, this field has seen recent resurgences as the size of storage systems has grown, and as Grid, peer-to-peer and cluster applications have had a need for it. Regardless, it remains a mess. Our mission is to declutter the field, so that the programmers and system designers who need to employ these codes may do so without having to wallow through the mire of obtuse theory that often does not even apply to their applications.

In this talk, Prof. Plank will explain the current state of the art, the avenues for useful research, and the research/outreach projects that he is leading.

Yves Robert

Matrix Product/LU with Limited Memory

Joel Saltz

Architecture of a Strongly Typed Grid and Lessons Learned from the Cancer Research Community

caGrid is a strongly typed grid layered on top of a curated system used to represent application level data elements and semantics. caGrid objectives are ambitious: the infrastructure is intended to be the middleware substrate to support discovery, database federation, service invocation and workflow management across the entire US Cancer research community. This talk will provide: 1) an overview of the caGrid system architecture, 2) examine consequences of caGrid design tradeoffs, 3) discuss extensibility of caGrid design elements to other application domains and 4) address relevance of caGrid type systems to DDDAS and high end computing infrastructure.

Rob Schreiber
Parallel Programming Research at HP Labs

David Scott
HPC @ Intel: Platforms and the Technologies That Go Into Them

For the first time since 1996, Intel has launched an HPC-focused board and platform. And soon, a second one will be launched. These two HPC-designed platforms focus on higher performance (of course), lower power, and higher density. These platforms include a number of technologies including wider execution, faster SSE2 instructions, better power management, smarter caches, and more, which leads to dramatic improvements in both performance and performance/watt. But more importantly, the platforms are designed with balance in mind, where we've taken end-user product feedback and turned it into product definition.

Intel is also researching technologies that will be of benefit to future HPC platforms including advanced process technology, stacked DRAM, and silicon photonics. In this talk, I will review the new platforms, I will briefly review the technologies enveloped therein, and then close with a glimpse of what is to come over the next four years.

Chuck Seitz
Ethernet Will Eradicate Ethernot

Over the next decade, 'specialty networks' -- networks other than mainstream Ethernet -- will gradually fade away, even for HPC-cluster applications. There are both technical and business reasons for this prediction. On the technical side, I'll try to show how Ethernet NIC and switch implementations are evolving to achieve the same low latency and host-CPU utilization as HPC-cluster specialty networks. On the business side, technology convergence and standardization, together with business consolidation, has been evident in the computer industry over the past decade, and will continue to reduce the choices along with the costs in high-speed networks and in other computing technologies.

Thomas Sterling
ParalleX: Towards a New Parallel Execution Model for Scalable Programming and Architectures

Future HPC systems will exhibit dramatic increases in hardware parallelism and latency due to technology trends while future applications in some cases will impose more severe memory-oriented behavior with poor temporal locality aggravating bandwidth constraints and degrading cache hierarchy performance. The exceptionally successful message-passing model (e.g. MPICH, PVM) must either evolve and adapt to these trends or yield to future execution strategies explicitly devised to address these challenges. To explore possible methods for developing future scalable hardware/software systems, an experimental execution model, ParalleX, is being derived that integrates a set of semantic constructs and operational mechanisms to mitigate factors of performance degradation. ParalleX supports a message-driven split-phase multi-threaded transaction processing paradigm incorporating futures-based synchronization. It is postulated that such a model will exploit near fine-grain data parallelism implicit in sparse data structures (e.g., directed graphs) and provide intrinsic latency hiding while facilitating dynamic load balancing. This presentation will describe the ParalleX model and discuss its implications for parallel architecture and programming.

Jeff Vetter
Application Accelerators: Deus ex Machina?

Commodity computing systems are rapidly moving to homogenous multicore processors as a strategy to continue improving performance while confronting the constraints of power, heat, signaling, and instruction level parallelism. Enter application accelerators. Numerous hardware accelerators have recently appeared on the supercomputing scene: FPGAs, ClearSpeed, STI Cell, Graphical Processing Units, etc. Our initial investigations have revealed that these accelerators can dramatically improve the performance of specific algorithms; in one example, our acceleration of an ORNL protein-folding application with FPGAs has shown good speedups. Nevertheless, accelerators face numerous hurdles to widespread adoption, such as programmer productivity and unstable performance reactivity. We are tackling these challenges in the Siskiyou project at ORNL with new performance modeling tools and a software system targeted at tightly-coupled heterogeneous computing systems.

David Walker
Portal Interfaces for Grid-Based Scientific Computing

This talk will focus on portal frameworks for providing user-friendly GUIs for distributed scientific computing. Particular attention will be given to standards-based portlets for building portal interfaces, and their use in integrating distributed resources and presenting

them to the user will be illustrated using the Grid-enabled Computational Electromagnetics (GECEM) portal. The prospects for developing a portal builder that is both easily configurable and general purpose will be considered, and future work on the GECEM portal will be discussed.

Rich Wolski
Predicting Bounds on the Batch Queuing Delay Experienced by User Jobs in Real Time

In this talk, we present a new method for providing TeraGrid end-users with real-time predictions of the bounds on queuing delay individual jobs will experience when waiting to be scheduled to a machine partition. Predicting the delay users will experience while waiting for their jobs to be be scheduled is a problem that has been studied both by the academic and commercial HPC communities for some time. Our approach, based on a new statistical methodology, predicts bounds on the waiting time (upper or lower) that individual jobs will experience with quantified confidence measures. Thus the predictions made by this system constitute a statistical guarantee of best-case and worst-case waiting delay where the confidence measure quantifies the quality of the guarantee.

We have implemented this new methodology as part of the Network Weather Service and deployed it on TeraGrid where it currently provides real-time bounds predictions. In the talk we will report on the effectiveness of the system which has been in operation as a prototype for approximately 8 months. We will discuss the methodology and its evaluation using batch-queue logs spanning 10 years at the NSF and open DOE supercomputer centers. We will also demonstrate the web interface to the system and make "live" predictions of TeraGrid delay bounds during the presentation from the web page located at http://nws.cs.ucsb.edu/batchq and we will detail the operation of a set of command-line tools that are portable among all ETF architectures.

Our results show that it is possible to predict delay bounds with specified confidence levels for individual jobs in different queues, and for jobs requesting different ranges of processor counts and different maximum execution delays. Using these predictions, users with roaming allocations or with allocations at multiple TeraGrid sites can choose the machine that is most likely to minimize turn-around time. Users can also determine the probability that a job will meet a specified deadline in a particular queue. Finally, the system is portable to all ETF architectures making it possible for users to consider the use of heterogeneous resources, and to predict which is most likely to impose the shortest waiting time for their jobs.

Hans Zima
High-Level Programming Models for Clusters: Issues and Challenges

In today's dominating programming paradigm for clusters, users are forced to adopt a low-level programming style if they want to fully exploit the capabilities of these machines. This leads to high-cost software production and error-prone programs that are difficult to write, reuse, and maintain. Emerging large-scale architectures with many thousands of processors, and applications of growing size and complexity will further aggravate this problem.

This presentation will discuss the requirements for high-productivity programming languages that represent a viable compromise between the dual goals of human productivity and target code efficiency, and outline the related research challenges. We provide an overview of new programming language design efforts in the High Productivity Computing Systems (HPCS) program funded by DARPA, with a special emphasis on the Chapel language. One of the new features of this language is an object-oriented framework for user-defined distributions: rather than offering a fixed set of built-in distributions, Chapel provides a distribution class interface which allows the explicit specification of the mapping of elements in a collection to units of uniform memory access, the control of the arrangement of elements within such units, the definition of sequential and parallel iteration over collections, and the specification of allocation policies. The result is a concise high-productivity programming model that separates algorithms from data representation and enables reuse of distributions, allocation policies, and specialized data structures.

We conclude this talk with an outline of new challenges posed by heterogeneous system architectures, which are expected to play an increasingly important role in the future.

Bios for Participants of CCGSC 2006 Flat Rock, NC

Thara Angskun

Thara Angskun received his Bachelor and Master degree in Computer Engineering from Kasetsart University, Thailand. Currently, he is a PhD student and graduate research assistance at the Innovative Computing Laboratory, Department of Computer Science, University of Tennessee, Knoxville. His major research interests are in parallel and distributed environment, message passing, high performance computing, computer networking, cluster and grid computing. He was a technical leader of an open source Beowulf cluster distribution project called OpenSCE. He is also a developer of several projects including KSIX, ACI, CAMETA, Harness / FT-MPI and Open MPI.

Mark Baker

School of Systems Engineering, University of Reading
Mark Baker is a full Research Professor of Computer Science in the School of Systems Engineering at the University of Reading. Prior to this he was a Reader of Distributed Systems at the University of Portsmouth, where he set up and ran the Distributed Systems Group (DSG) for nine years. Mark's research interests are related to middleware technologies for Clusters and the Grid. Mark has published widely in journals and conferences in his research area, which includes all aspects of distributed systems. Mark has managed and been involved with a number of UK and EU funded projects. This includes, more recently, SORMA (Self-Organizing ICT Resource Management), the Sakai VRE Demonstrator project and the OGSA Testbed. Mark is a senior member of the IEEE and is involved in many of its activities.

Micah Beck

University of Tennessee
Department of Computer Science
Claxton Complex Suite 203, 1122 Volunteer Blvd
Knoxville TN 37996
Micah Beck began his research career in distributed operating systems at Bell Laboratories and received his Ph.D. in Computer Science from Cornell University (1992) in the area of parallelizing compilers. He then joined the faculty of the Computer Science Department at the University of Tennessee, where he is currently an Associate Professor working in distributed high performance computing, networking and storage.
An active participant in the Internet2 project since 1997, Dr. Beck is chair of their Network Storage Special Interest Group. Dr. Beck's other activities include co-founding a company, Lokomo Systems, in 2000; collaboration with Oak Ridge National

Laboratory's Future Technologies Group and with Vanderbilt University's Advanced Computing Center for Research and Education.

Pete Beckman

Peter Beckman has worked in systems software for parallel computing, operating systems, and Grid computing for 20 years. After receiving a Ph.D. in computer science from Indiana University, he helped create the Extreme Computing Laboratory, which focused on parallel C+ +, portable run-time systems, and collaboration technology. In 1997 Peter joined the Advanced Computing Laboratory at Los Alamos National Laboratory, where he founded the ACL's Linux cluster team and organized the Extreme Linux series of workshops and activities that helped catalyze the high-performance Linux computing cluster community. Peter has also worked in industry. For example, in 2000 he founded a research laboratory in Santa Fe (sponsored by Turbolinux Inc.), which developed the world's first dynamic provisioning system for large clusters and data centers. The following year, Peter became vice president of Turbolinux's worldwide engineering efforts, managing development offices in Japan, China, Korea, and Slovenia. Peter began working at Argonne National Laboratory in 2002. As director of engineering for the TeraGrid, a $150 million effort sponsored by the National Science Foundation to build the world's largest open Grid computing environment, he designed and deployed the world's most advanced Grid system for linking production HPC computing centers. After the TeraGrid became fully operational, Peter started a research team focusing on petascale high-performance software systems, wireless sensor networks, Linux, and the emerging field of "urgent computing" for critical, time-sensitive decision support. He has published numerous articles, served on national program committees, and presented invited papers and tutorials. Peter, his wife, and two children live in Naperville, Illinois.

Fran Berman

Dr. Francine Berman is Professor in the UCSD Department of Computer Science and Engineering, Fellow of the ACM, and first holder of the High Performance Computing Endowed Chair in the Jacobs' School of Engineering at UCSD. Dr. Berman is a pioneer in Grid Computing and an international leader in the development of Cyberinfrastructure. She has worked extensively in the areas of adaptive middleware, parallel programming environments, scheduling, and high performance computing. Since 2001, Dr. Berman has served as Director of the San Diego Supercomputer Center (SDSC) where she leads a staff of 400+ interdisciplinary scientists, engineers, and technologists in the innovation and provision of national-scale Cyberinfrastructure. SDSC is a global leader in Data Cyberinfrastructure, and works closely with the National Science Foundation, National Archives and Records Administration, Library of Congress, Department of Energy, and other federal agencies to innovate and support a national information and computational infrastructure.

Dr. Berman is one of the two founding Principle Investigators of the National Science Foundation's TeraGrid project, and also directed the National Partnership for Advanced Computational Infrastructure (NPACI), a consortium of 41 research groups, institutions, and university partners with the goal of building national infrastructure to support

research and education in science and engineering. For her accomplishments, leadership, and vision, Dr. Berman was recognized in 2004 as one of the top women in technology by BusinessWeek, as one of the top technologists by IEEE Spectrum, and most recently as a leader in science and technology by Newsweek.

Jean-Yves Berthou

Chargé de mission, Strategic Steering Manager for Simulation
EDF, Research and Development Division
Jean-Yves Berthou has been a researcher for EDF R&D since 1997. He received a Ph.D in computer science from "Pierre et Marie Curie" University (PARIS VI) in 1993. His research deals mainly with Parallelization, Parallel Programming and Software Architecture for scientific computing.

Jean-Yves Berthou was the head of the Applied Scientific Computing Group at EDF R&D from 2002 to 2006. This research group was dedicated to Simulation Platforms Development, Scientific Software Architecture, Code Optimization, High Performance Computing, Cluster and Grid Computing. He is now *Chargé de Mission* – Strategic Steering Manager for Simulation, in charge of the simulation program at EDF R&D"

George Bosilca

Research Scientist
Innovative Computing Laboratory
University of Tennessee Department of Computer Science
Dr. Bosilca is a Senior Research Associate at the Innovative Computing Laboratory (ICL). He received a Ph.D. degree in parallel architectures from the University Orsay Paris XI. His Ph.D was focused on parallel environments based on automatic data dependencies graphs, and fault tolerance. He was the main developer of the channel memory subsystem for MPICH-V. Once he joined the ICL he was continuing to work on fault tolerance related topics (such as FT-MPI), collective communications libraries as well as MPI passing libraries. Today, Dr. Bosilca is one of the main developers on the Open MPI project.

Franck Cappello

Franck Cappello holds a Senior Researcher position at INRIA He leads the Grand-Large project at INRIA, focusing on large scale distributed systems issues. He has initiated the XtremWeb (Desktop Grid) and MPICH-V (Fault tolerant MPI) projects. He is currently the director of the Grid5000 project, a nation wide computer science platform for research in Grid and P2P. He has authored more than 60 papers in the domains of High Performance Programming, Desktop Grids, Grids and Fault tolerant MPI. He has contributed to more than 30 Program Committees. He is editorial board member of the "international Journal on GRID computing" and steering committee member of IEEE HPDC and IEEE/ACM CCGRID. He was the general chair of IEEE HPDC'2006.

Giri Chukkapalli

Chief HPC Architect
Network Systems Group
Sun Microsystems, Inc.
giri.chukkapalli@sun.com

Giri holds Ph.D from Dept. of Mechanical and Industrial Engg., Univ. of Toronto (1994-97) and ME from Dept. of Mechanical Engg., McMaster Univ. (1991-1993). As part of his multi-disciplinary Ph.D thesis, he implemented a parallel climate/weather dynamics solver with a novel mesh and advection scheme.

He worked as a programmer/Analyst, staff Scientist and Ass. Director in Scientific Computing Dept., San Diego Supercomputer Center from 1998 to 2005. During this time, he worked on several high profile HPC codes both from DoD, NSF and NIH codes including NOGAPS, COAMPS, CHARMM, Rosetta, HMMER, SPECFEM3D porting, tuning and optimizing them. He was also funded from 2000 to 2004 to work on the design and development of Bluegene with scientific applications focus. He played critical roles in selecting, configuring, running, benchmarking and supporting HPC production systems at SDSC.

Currently, He is a member of NSF HPC resources allocation review panel. He is also on the technical paper review committee of Supercomputing 04, 05 and LCI conference 04, 05 and 06.

Michel Cosnard

INRIA - FRANCE

Michel Cosnard received the Master's degree in Applied Mathematics from Cornell University in 1975 and Doctorat d'Etat degree in Computer Science from the Université de Grenoble (France) in 1983.

In 1987, he became Professor of Computer Science at the Ecole Normale Supérieure de Lyon (France) where he has founded the Laboratoire d'Informatique du Parallélisme (LIP). From September 1997 to December 2000, he served as director of the INRIA Research Unit in Lorraine and director of the LORIA laboratory. In june 2001, he has been nominated director of the INRIA Research Unit in Sophia Antipolis (Nice) and serves as Professor at the University of Nice - Sophia Antipolis from September 2002.

In May 2006, he has been appointed Chairman of the board and CEO of INRIA.

He has been Editor in Chief of Parallel Processing Letters, and member of the Editorial Board of several scientific journals (IEEE TPDS, Parallel Computing, Mathematical Systems Theory,...) and steering committees of international conferences (PACT, SPAA, IPDPS, ...).

His research interests are in the design and analysis of parallel

algorithms, in the complexity analysis of automata and neural nets. Michel Cosnard has published more than 100 papers related to parallel processing. He served as chairman of the IFIP Working Group WG 10.3 on Parallel Processing from 1989 to 1995. In 1994, Michel Cosnard was awarded a national prize from the French National Academy of Science. In 1995, he received the IFIP Silver Core and in 2003 the IPDPS Babbage award.

David Cronk

David received a BS in mathematics from Hope College in 1988, an MS in computer science from Marist College in 1992, and a Ph.D. in computer science from the College of William and Mary in 1999. He specializes in high performance computing with an emphasis on message passing libraries, parallel I/O, and performance analysis tools. David has been at the Innovative Computing Lab at the University of Tennessee working for Jack Dongarra since 1999, where he currently holds the title of Research Director. At Tennessee he leads the DoD HPCMP's PET program's computational environment functional area.

Philippe d'Anfray

GIP RENATER
151 Boulevard de l'Hôpital 75013 Paris France
Philippe.d-Anfray@renater.fr

Philippe d'Anfray began his research career in 1981 at the French Electricity Company EDF working on software development in numerical analysis. He left in 1989 for the French Aerospace Research Center ONERA where he specialized on models and tools for parallel and distributed computing and was involved in many European research Projects. In 2002 he joined the French Atomic Energy Commission CEA to start a new activity on Grid Computing. He is also a professor at Paris XIII University (post graduate courses in computer science) and manages a "Grid User's working Group" inside the Aristote Association gathering industrial and academic representatives. Since january 2006, Philippe d'Anfray is responsible, at RENATER -the French National Telecommunication Network for Technology Education and Research- for the coordination of Grid Research Projects with high network requirements.

Fredrica Darema

Frederica Darema, Ph. D., Fellow IEEE, Senior Executive Service Member Dr. Darema is the Senior Science and Technology Advisor in CNS and CISE, and Director of the Computer Systems Research (CSR) Program, and Lead of the multi-agency DDDAS Program. Dr. Darema's interests and technical contributions span the development of parallel applications, parallel algorithms, programming models, environments, and performance methods and tools for the design of applications and of software for parallel and distributed systems. Dr. Darema received her BS degree from the School of Physics

and Mathematics of the University of Athens - Greece, and MS and Ph. D. degrees in Theoretical Nuclear Physics from the Illinois Institute of Technology and the University of California at Davis Respectively, where she attended as a Fulbright Scholar and a Distinguished Scholar. After Physics Research Associate positions at the University of Pittsburgh and Brookhaven National Lab, she received an APS Industrial Fellowship and became a Technical Staff Member in the Nuclear Sciences Department at Schlumberger-Doll Research. Subsequently, in 1982, she joined the IBM T. J. Watson Research Center as a Research Staff Member in the Computer Sciences Department and later-on she established and became the manager of a research group at IBM Research on parallel applications. While at IBM she also served in the IBM Corporate Strategy Group examining and helping to set corporate-wide strategies. Dr. Darema was elected IEEE Fellow for proposing in 1984 the SPMD (Single-Program-Multiple-Data) computational model that has become the popular model for programming today's parallel and distributed computers. Dr. Darema has been at NSF since 1994, where she has developed initiatives for new systems software technologies (the Next Generation Software Program), and research at the interface of neurobiology and computing (the Biological Information Technology and Systems Program). She has led the DDDAS (Dynamic Data Driven Applications Systems) efforts including the synonymous cross-Directorate and cross-agency competition. She has also been involved in other cross-Directorate efforts such as the Information Technology Research, the Nanotechnolgy Science and Engineering, the Scalable Enterprise Systems, and the Sensors Programs. During 1996-1998 she completed a two-year assignment at DARPA where she initiated a new thrust for research on methods and technology for performance engineered systems.

Michel Dayde

Frédéric Desprez

Frédéric Desprez is a director of research at INRIA and holds a position at LIP laboratory (ENS Lyon, France). He received is PhD in C.S. from the Institut National Polytechnique de Grenoble in 1994 and his MS in C.S. from the ENS Lyon in 1990. His research interests include parallel libraries for scientific computing on parallel distributed memory machines, problem solving environments, and grid computing.
See \url{http://graal.ens-lyon.fr/~desprez} for further information.

Jack Dongarra

Computer Science Department
1122 Volunteer Blvd
University of Tennessee
Knoxville, TN 37996-3450
dongarra@cs.utk.edu
Jack Dongarra received a Bachelor of Science in Mathematics from Chicago State University in 1972, a Master of Science in Computer Science from the Illinois Institute of Technology in 1973, and a Ph.D. in Applied Mathematics from the University of New Mexico in 1980. He specializes in numerical algorithms in linear algebra, parallel

computing, the use of advanced-computer architectures, programming methodology, and tools for parallel computers. His research includes the development, testing and documentation of high quality mathematical software. He has contributed to the design and implementation of the following open source software packages and systems: EISPACK, LINPACK, the BLAS, LAPACK, ScaLAPACK, Netlib, PVM, MPI, NetSolve, Top500, ATLAS, and PAPI. He was awarded the IEEE Sid Fernbach Award in 2004 for his contributions in the application of high performance computers using innovative approaches. He is a Fellow of the AAAS, ACM, and the IEEE and a member of the National Academy of Engineering.

Graham Fagg

Dan Fay

Daniel Fay is the TCI Director of North America for Microsoft's Technical Computing Group, where he works with academic research projects focused on utilizing computing technologies to aid in scientific and engineering research. Dan was previously the manager of eScience Program in Microsoft Research as well as program manager for Project 7, a unique project that brought together a number of language researchers to examine and provide feedback on the Common Language Runtime and its multi language support. Before joining Microsoft in 1992, Dan was a senior software engineer with Digital Equipment Corporation working on OSF/1 and embedded network devices. Dan is a graduate of Northeastern University in Electrical Engineering.

Wuchun Feng

Depts. of Computer Science and Electrical & Computer Engineering (0106)
Virginia Tech
Blacksburg, VA 24061
Prof. Wu-chun Feng (or more simply "Wu") is an associate professor of computer science with a joint appointment in Electrical & Computer Engineering at Virginia Tech (VT). At VT, he leads the Systems, Networking, & Renaissance Grokking (SyNeRGy) Laboratory. He received a B.S. in Electrical & Computer Engineering and Music (Honors) and an M.S. in Computer Engineering from the Pennsylvania State University in 1988 and 1990, respectively. He earned a Ph.D. in Computer Science from the University of Illinois at Urbana-Champaign in 1996. He is a senior member of the IEEE.

Geoffrey Fox

Indiana University
Community Grids Laboratory
Bloomington, IN 47401
Fox received a Ph.D. in Theoretical Physics from Cambridge University and is now professor of Computer Science, Informatics, and Physics at Indiana University. He is director of the Community Grids Laboratory of the Pervasive Technology Laboratories at Indiana University. He previously held positions at Caltech, Syracuse University and Florida State University. He has published over 550 papers in physics and computer

science and been a major author on four books. Fox has worked in a variety of applied computer science fields with his work on computational physics evolving into contributions to parallel computing and now to Grid systems. He has worked on the computing issues in several application areas – currently focusing on Earthquake Science.

Edgar Gabriel

Edgar Gabriel is an Assistant Professor in the Department of Computer Science at the University of Houston, Texas, USA. Before that, he was leader of the working group 'Clusters and Distributed Units' at the High Performance Computing Center Stuttgart (HLRS), Germany and a post-doctoral research at the Innovative Computing Laboratory (ICL) at the University of Tennessee, Knoxville, USA. His research interests are Message Passing Systems, High Performance Computing, Parallel Computing on Distributed Memory Machines, and Grid Computing.

Al Geist

Oak Ridge National Laboratory
Al Geist is a Corporate Research Fellow at Oak Ridge National Laboratory, where he leads the 35 member Computer Science Research Group. He is one of the original developers of PVM (Parallel Virtual Machine), which became a world-wide de facto standard for heterogeneous distributed computing. Al was actively involved in the design of the Message Passing Interface (MPI-1 and MPI-2) standard and more recently the development of FT-MPI, a fault tolerant MPI implementation.
Today He leads a national Scalable Systems Software effort, involving all the DOE and NSF supercomputer sites, with the goal of defining standardized interfaces between system software components. Al is co-PI of a national Genomes to Life (GTL) center. The goal of his GTL center is to develop new algorithms, and computational infrastructure for understanding protein machines and regulatory pathways in cells. He heads up a project developing self-adapting, fault tolerant algorithms for 100,000 processor systems.

Patrick Geoffray

Patrick Geoffray earned his Ph.D. at the University of Lyon, France, in 2001. His interests lie in high-performance computing. He is currently a Senior Software Architect at Myricom, in the Software Development Lab in Oak Ridge, TN. He is responsible for the design of the Myrinet Express (MX) message-passing system and its associated firmware implementations, and was earlier in charge of various middleware running on GM/Myrinet, including MPICH-GM and VIA-GM.

Olivier Gluck

Olivier Glück received both in 1999 his M.S. in computer science from Evry University and a high engineering school degree from National Institute of Telecommunications. He joined the LIP6 laboratory of University Paris 6 in 1999 where he received a Ph.D.

degree in computer science. After serving two years (2001-2003) as assistant professor in computer science in Paris 6 University, he joined the Lyon 1 University as associate professor. Since 2003, he is a permanent member of the INRIA RESO team in the LIP laboratory of Ecole Normale Superieure de Lyon, France.
His research interests include high performance, scalable communication interfaces and protocols for system area networks, operating systems for parallel processing machines, and high speed networks for clusters and grids. While at Paris 6 University, he has designed and developed a high performance implementation of the MPI Standard on a parallel computer made of clusters of workstations, providing a remote-write communication primitive; he has primarily concentrated on topics related to efficient, scalable message passing and interactions at the software/hardware boundary. Since 2003, he mainly works on an efficient distributed file system within a Myrinet Cluster and on a networks emulation tool for grids. He is also involved in two major national projects: the GridExplorer project within the ACI Grandes Masses de Données which aim is to create a large scale grid and network emulator and the GRID'5000 project within the CNRS AS enabling Grid'5000 which aims at building an experimental Grid platform gathering eight sites geographically distributed in France.

Rich Graham

Los Alamos National Laboratory
MS B287
Los Alamos, NM 87544
Richard Graham is the Advanced Computing Laboratory acting group leader at the Los Alamos National Laboratory. He joined LANL's Advanced Computing Laboratory (ACL) as a technical staff member in 1999 and as team leader for the Resilient Technologies Team team started the LA-MPI project, and is one of the founders of the Open MPI project. Prior to joining the ACL, he spent seven years working for Cray Research and SGI.

Rich obtained his PhD in Theoretical Chemistry from Texas A&M University in 1990 and did post-doctoral work at the James Franck Institute of the University of Chicago. His BS in chemistry was from Seattle Pacific University.

Bill Gropp

William Gropp received his B.S. in Mathematics from Case Western Reserve University in 1977, a MS in Physics from the University of Washington in 1978, and a Ph.D. in Computer Science from Stanford in 1982. He held the positions of assistant (1982-1988) and associate (1988-1990) professor in the Computer Science Department at Yale University. In 1990, he joined the Numerical Analysis group at Argonne, where he is a Senior Computer Scientist in the Mathematics and Computer Science Division, a Senior Scientist in the Department of Computer Science at the University of Chicago, and a Senior Fellow in the Argonne-Chicago Computation Institute. His research interests are in parallel computing, software for scientific computing, and numerical methods for partial differential equations. He has played a major role in the development of the MPI

message-passing standard. He is co-author of the most widely used implementation of MPI, MPICH, and was involved in the MPI Forum as a chapter author for both MPI-1 and MPI-2. He has written many books and papers on MPI including "Using MPI" and "Using MPI-2". He is also one of the designers of the PETSc parallel numerical library, and has developed efficient and scalable parallel algorithms for the solution of linear and nonlinear equations.

Tony Hey

Microsoft Corp.
As corporate vice president for technical computing, Tony Hey coordinates efforts across Microsoft Corp. to collaborate with the global scientific community. He is a top researcher in the field of parallel computing, and his experience in applying computing technologies to scientific research helps Microsoft work with researchers worldwide in various fields of science and engineering.

Before joining Microsoft, Hey worked as head of the School of Electronics and Computer Science at the University of Southampton, where he helped build the department into one of the pre-eminent computer science research institutions in England. Since 2001, Hey has served as director of the U.K.'s e-Science Initiative, managing the government's efforts to provide scientists and researchers with access to key computing technologies.

Hey is a fellow of the U.K.'s Royal Academy of Engineering and has been a member of the European Union's Information Society Technology Advisory Group. He has also served on several national committees in the United Kingdom, including committees of the U.K. Department of Trade and Industry and the Office of Science and Technology. In addition, Hey has advised countries such as China, France, Ireland and Switzerland to help them advance their scientific agenda and become more competitive in the global technology economy. Hey received the award of Commander of the Order of the British Empire honor for services to science in the 2005 U.K. New Year's Honours List.

Hey is a graduate of Oxford University, with both an undergraduate degree in physics and a doctorate in theoretical physics.

Jeff Hollingsworth

Jeffrey Hollingsworth is a Professor in the Computer Science Department at the University of Maryland, College Park, and affiliated with the Department of Electrical Engineering and the University of Maryland Institute for Advanced Computer Studies. His research interests include instrumentation and measurement tools, resource aware computing, high performance distributed computing, and programmer productivity. Dr. Hollingsworth's current projects include the Dyninst runtime binary editing tool, and Harmony - a system for building adaptable, resource-aware programs. He received his PhD and MS degrees in computer science from the University of Wisconsin.

Emmanuel Jeannot

Emmanuel Jeannot is currently full-time researcher at INRIA (Institut National de Recherche en Informatique et en Automatique) and is doing its research at the LORIA laboratory. From September 1999 to September 2005 he was associate professor at the Université Henry Poincaré, Nancy 1. He received his PhD and Master degree of computer science (resp. in 1996 and 1999) both from Ecole Normale Supérieure de Lyon. His main research interests are scheduling for heterogeneous environments and grids, data redistribution, grid computing software, adaptive online compression and programming models.

Tahar Katchadi

University College Dublin
Tahar Kechadi was awarded Ph.D. and a DEA (Diplome d'Etude Approfondie) - Masters degree - in computer science, from University of Lille 1, France, in 1993 and 1990 respectively. After working as a post-doctoral fellow under TMR program at UCD, he joined UCD Computer Science Department as a permanent staff in 1999. His research interests span the areas of parallel processing, parallel architectures, scheduling, dynamic load balancing, heterogeneous distributed systems, distributed data mining, Grid computing, and optimisation techniques based on heuristics. He is a member of the communication of the ACM and IEEE computer society.

Ken Kennedy

Ken Kennedy is the John and Ann Doerr University Professor of Computational Engineering and Director of the *Center for High Performance Software Research (HiPerSoft)* at Rice University. He is a fellow of the Institute of Electrical and Electronics Engineers, the Association for Computing Machinery, and the American Association for the Advancement of Science. He was elected to the National Academy of Engineering in 1990 and to the American Academy of Arts and Sciences in 2005. From 1997 to 1999, he served as co-chair of the President's Information Technology Advisory Committee (PITAC). For his leadership in producing the PITAC report on funding of information technology research, he received the Computing Research Association Distinguished Service Award (1999) and the RCI Seymour Cray HPCC Industry Recognition Award (1999).

Prof. Kennedy has published two books and over two hundred technical articles and supervised thirty-seven Ph.D. dissertations on programming support software for high-performance computer systems. His current research focuses on programming languages tools to improve the productivity of scientists and engineers developing technical applications for complex platforms, particularly scalable parallel computers and the Grid. In recognition of his contributions to software for high performance computation, he received the 1995 W. Wallace McDowell Award, the highest research award of the IEEE Computer Society. In 1999, he was named the third recipient of the ACM SIGPLAN Programming Languages Achievement Award.

Thilo Kielmann

Vrije Universiteit
Dept. of Computer Science
De Boelelaan 1083
1081HV Amsterdam
The Netherlands

Thilo Kielmann is Associate Professor at the Department of Computer Science of Vrije Universiteit. He received his diploma in Computer Science from Darmstadt University of Technology, Germany, in 1992, a Ph.D. in Computer Engineering in 1997, and the habilitation in Computer Science in 2001, the latter both from the University of Siegen, Germany. He serves on the steering groups of the Open Grid Forum and of GridForum Nederland. His research interests include computational grids, performance analysis of networked applications, and parallel and distributed programming.

Dieter Kranzlmueller

GUP, Joh. Kepler University
Linz, Altenbergerstr. 69, A-4040
Linz, Austria/Europe

Dieter Kranzlmueller is Professor of computer science at the Johannes Kepler University Linz, Group Leader Parallel Computing and Deputy Head of the GUP - Institute of Graphics and Parallel Processing. He is a member of the Executive Commitee of the national Austrian Grid Initiative, Austrian National Representative and Vice-Chair of the European e-Infrastructure Reflection Group (e-IRG), and Area Director Applications of the Open Grid Forum (OGF).

Alexey Lastovetsky

Alexey Lastovetsky received a PhD degree from the Moscow Aviation Institute in 1986, and a Doctor of Science degree from the Russian Academy of Sciences in 1997. His main research interests include algorithms, models and programming tools for high performance heterogeneous computing. He is the author of mpC, the first parallel programming language for heterogeneous networks of computers. He designed HeteroMPI, an extension of MPI for heterogeneous parallel computing, and SmartNetSolve, an extension of NetSolve aimed at higher performance of scientific computing on global networks. He has made contributions into heterogeneous data distribution algorithms (with A. Kalinov, R. Reddy, etc.). He also proposed and investigated realistic performance models of processors in heterogeneous environments, including the functional model and the band model (with R. Reddy and R. Higgins). He published over 70 technical papers in refereed journals, edited books and proceedings of international conferences. He authored the monograph “Parallel computing on heterogeneous networks” published by Wiley in 2003. He is currently a senior lecturer in the School of Computer Science and Informatics at University College Dublin, National University of Ireland. At UCD, he also created and heads Heterogeneous Computing Laboratory. He is an editor of the research journals “Parallel Computing” (Elsevier) and “Programming and Computer Software” (Springer). He is a recipient of SFI Investigator Award (2004, €700,000) and SFI Basic Research Award (2004, €200,000).

Laurent Lefevre

Laurent Lefevre obtained his Ph.D in Computer Science in January 1997 at LIP Laboratory (Laboratoire Informatique du Parallelisme) in ENS-Lyon (Ecole Normale Superieure), France. From 1997 to 2001, he was assistant professor in computer science in Lyon 1 University and in the RESAM Laboratory (High Performance Networks and Multimedia Application Support Lab.). Since 2001, he is permanent researcher in computer science at INRIA (The French Institute for Research in Computer Science and Control). He is member of the RESO team (High performance networks, protocols and services) from LIP laboratory in Lyon, France. He has organized several conferences in high performance networking and computing and he is member of several program committees. He has co-authored more than 30 papers published in refereed journals and conference proceedings. He is member of IEEE and takes part in several research projects. His research interests include:
autonomic networking, high performance active networks, active services, high performance network protocols, Grid and Cluster computing network support, Active Grid, Distributed Shared Memory systems and Data consistency.

Miron Livny

Miron Livny received a B.Sc. degree in Physics and Mathematics in 1975 from the Hebrew University and M.Sc. and Ph.D. degrees in Computer Science from the Weizmann Institute of Science in 1978 and 1984, respectively. Since 1983 he has been on the Computer Sciences Department faculty at the University of Wisconsin-Madison, where he is currently a Professor of Computer Sciences and is leading the Condor project.

Dr. Livny's research focuses on distributed processing and data management systems and data visualization environments. His recent work includes the Condor high throughput computing system, the DEVise data visualization and exploration environment and the BMRB repository for data from NMR spectroscopy.

Thomas Ludwig

Ruprecht-Karls-Universität Heidelberg
Institut für Informatik
Im Neuenheimer Feld 348
69120 Heidelberg, Germany
Thomas Ludwig received his habilitation degree from Technische Universität München in Munich, Germany, where he worked for 13 years in the field of parallel computing with a focus on load balancing, development tools, and cluster and tool infrastructures. He also conducted research in the field of parallel programming, namely with computer tomography and bioinformatics. Since 2001 he is a professor for computer science at the Ruprecht-Karls-Universität Heidelberg in Heidelberg, Germany. His current research focus is in the field of high performance parallel input/output systems for cluster environments.

Andrew Lumsdaine

Andrew Lumsdaine received his Ph.D. from MIT in 1992. From 1992 through 2001, he was on the faculty in the Department of Computer Science and Engineering at the University of Notre Dame. He is presently a professor of computer science at Indiana University and director of the Open Systems Laboratory. His research interests include computational science and engineering, parallel and distributed computing, software engineering, generic programming, mathematical software, and numerical analysis.

Rusty Lusk

Ewing Lusk received his B.A. in mathematics from the University of Notre Dame in 1965 and his Ph.D. in mathematics from the University of Maryland in 1970. He is currently a senior computer scientist and acting division director in the Mathematics and Computer Science Division at Argonne National Laboratory. His current projects include implementation of the MPI Message-Passing Standard, parallel performance analysis tools, and system software for clusters. He is a leading member of the team responsible for MPICH implementation of the MPI message-passing interface standard. He is the author of five books and more than a hundred research articles in mathematics, automated deduction, and parallel computing.

Satoshi Matsuoka

Satoshi Matsuoka received his Ph. D. from the University of Tokyo in 1993. He became a full Professor at the Global Scientific Information and Computing Center (GSIC) of Tokyo Institute of Technology (Tokyo Tech / Titech) in April 2001, leading the Research Infrastructure Division Solving Environment Group of the Titech campus. He has pioneered grid computing research in Japan the mid 90s along with his collaborators, and currently serves as sub-leader of the Japanese National Research Grid Initiative (NAREGI) project, that aim to create middleware for next-generation CyberScience Infrastructure. He was also the technical leader in the construction of the TSUBAME supercomputer, which has become the fast supercomputer in Asia-Pacific in June, 2006, and also serves as the core grid resource in the Titech Campus Grid. He has been (co-) program and general chairs of several international conferences including ACM OOPSLA'2002, IEEE CCGrid 2003, HPCAsia 2004, Grid 2006, CCGrid 2006/2007, as well as countless program committee positions, in particular numerous SC technical papers committee duties including serving as the network area chair for SC2004.
He has been a Steering Group member and an Area Director of the Global Grid Forum since 1999. He has won several awards including the Sakai award for research excellence from the Information Processing Society of Japan in 1999, and recently received the JSPS Prize from the Japan Society for Promotion of Science. He serves several editorial positions of international journals including Concurrency: Practice and Experience, and the Journal of Grid Computing.

Bart Miller

Barton Miller is Professor of Computer Sciences at the University of Wisconsin, Madison. He directs the Paradyn Parallel Performance Tool project, which is investigating performance and instrumentation technologies for parallel and distributed applications and systems. He also co-directs the WiSA security project. His research interests include tools for high-performance computing systems, binary code analysis and instrumentation, computer security, and scalable distributed systems.

Miller co-chaired the SC2003 Technical Papers program, was Program co-Chair of the 1998 ACM/SIGMETRICS Symposium on Parallel and Distributed Tools, and General Chair of the 1996 ACM/SIGMETRICS Symposium on Parallel and Distributed Tools. He also twice chaired the ACM/ONR Workshop on Parallel and Distributed Debugging. Miller was on the editorial board of IEEE Transactions on Parallel and Distributed Systems, and is currently on the Boards of Concurrency and Computation Practice and Experience, Computing Systems, and the Int'l Journal of Parallel Processing. Miller has chaired numerous workshops and has been on numerous conference program committees. He is also a member of the IEEE Technical Committee on Parallel Processing.

Miller is a member of the Los Alamos National Laboratory Computing, Communications and Networking Division Review Committee, IDA Center for Computing Sciences Program Review Committee, and has been on the U.S. Secret Service Electronic Crimes Task Force (Chicago Area), the Advisory Committee for Tuskegee University's High Performance Computing Program, and the Advisory Board for the International Summer Institute on Parallel Computer Architectures, Languages, and Algorithms in Prague. Miller is an active participant in the European Union APART performance tools initiative.

Miller received his Ph.D. degree in Computer Science from the University of California, Berkeley in 1984. He is a Fellow of the ACM.

Terry Moore

Terry Moore is currently Associate Director of the Innovative Computing Laboratory at the Computer Science Department at University of Tennessee, where he is also the Managing Editor of Cyberinfrastructure Technology Watch Quarterly. He received his BA in Philosophy from Purdue University and his Ph.D. in Philosophy from the University of North Carolina, Chapel Hill. His research interests include collaboration technologies, network storage and especially logistical networking, in which he has been a leader since the inception of the field in 1997.

John Patrick Morrison

Computer Science Department
University College
Cork, Ireland.
John Patrick Morrison is an Associate Professor in the Computer Science Department of

University College Cork, Ireland. He is the founder and director of the Centre for Unified Computing, which has a staff of 20 researchers. He is a cofounder and co-director of the Boole Centre for Research in Informatics, which has a staff of 70 researchers and faculty members working at the interface between Computer Science and Mathematics. He is a cofounder and co-director of Grid-Ireland. Established in 1999, Grid-Ireland is a participant in the "Establishing Grids for e-Science in Europe" project. Currently Grid-Ireland supports four virtual organisations across 17 institutions on the island of Ireland.

Before returning to Cork in 1991, he worked for 6 years as a research scientist at the Philips' *Natuurkundig Laboratorium* in the Netherlands and received his PhD from the *Technische Universiteit Eindhoven.*

He is a Science Foundation of Ireland Investigator award holder and has attracted more than Euro 20M in research funding since 1997. He is a member of many international programme committees and editorial boards, has published widely and is a cofounder of the International Symposium on Parallel and Distributed Computing. He is a member of the ACM and a senior member of the IEEE.

Phil Papadopoulos

Dr. Papadopoulos received his PhD in 1993 from UC Santa Barbara in Electrical Engineering. He spent 5 years at Oak Ridge National Laboratory as part of the Parallel Virtual Machine (PVM) development team. He is currently the Program Director of Grid and Cluster Computing at the San Diego Supercomputer Center. Dr. Papadopoulos is deeply involved in key research projects including the Biomedical Informatics Research Network (BIRN), OptIPuter, the Geosciences Network (GEON), the NSF Middleware Initiative (NMI), The National Biomedical Computation Resource (NBCR), and the Pacific Rim Applications and Grid Middleware Assembly (PRAGMA). He is also well known for the development of the open source Rocks Cluster toolkit, which has installed base of 1000s of clusters. His research interests revolve around distributed and clustered systems and how they can be used more effectively in an expanding bandwidth-rich environment.

Jelena Pjesivac-Grbovic

Innovative Computing Laboratory, Department of Computer Science
The University of Tennessee at Knoxville
1122 Volunteer Blvd. Claxton, Suite 413
Knoxville, TN 37996
Ms. Pjesivac-Grbovic is a Graduate Research Assistant at the ICL, working toward a Ph.D. degree in computer science. She received a B.S. degree in computer science and physics from Ramapo College of New Jersey. Her research interests are parallel communication libraries and computer architectures, scientific and grid computing, and modelling of biophysical systems.

Jim Plank

Department of Computer Science
University of Tennessee
203 Claxton Complex
1122 Volunteer Blvd.
Knoxville, TN 37996-3450
plank@cs.utk.edu
http://www.cs.utk.edu/~plank

Jim Plank has been at the University of Tennessee for 13 years.
His research has encompassed many facets of fault-tolerant and dependable computing, from checkpointing of parallel and distributed applications, to scheduling on Grid systems, to scalable and shared network storage methodologies. His most recent research thrust has been to declutter the field of erasure codes for fault-tolerant storage applications, wading through the mire of overly formal coding theory, coding techniques for asymptotic content distribution systems, and the patent-laden world of large disk array systems. His mission is to blend the best of these areas to design, evaluate and implement efficient erasure codes for cluster, grid and peer-to-peer storage systems.

Yves Robert

ENS Lyon

Yves Robert received the PhD degree from Institut National Polytechnique de Grenoble in 1986. He is currently a full professor in the Computer Science Laboratory LIP at ENS Lyon. He is the author of four books, 95 papers published in international journals, and 120 papers published in international conferences. His main research interests are scheduling techniques and parallel algorithms for clusters and grids. He is a Fellow of the IEEE.

Joel Saltz

Department Chair
Biomedical Informatics
Ohio State University Medical Center
3190 Graves Hall
333 W 10th Avenue
Columbus, OH 43210
614-292-4778

Dr. Joel H. Saltz is Dorothy M. Davis Professor in Cancer Research and Chair of the Department of Biomedical Informatics, Professor in the Department of Computer Science and Engineering, Professor and Vice Chair and Director of Pathology Informatics in the Department of Pathology, and Associate Vice President for Health Sciences at The Ohio State University. He is also a Senior Fellow of the Ohio Supercomputer Center. Prior to coming to Ohio State, Dr. Saltz was Professor of Pathology and Informatics in the

Department of Pathology at Johns Hopkins Medical School and Professor in the Department of Computer Science at the University of Maryland. Dr. Saltz is trained both as a computer scientist and as a medical scientist. He received BS and MS degrees in Mathematics and Physics from the University of Michigan and his MD-PhD in Computer Science from Duke University. He completed his residency in Clinical Pathology at Johns Hopkins University with extensive training in microbiology under Dr. William Merz and Dr. James Dick, and is a board certified Clinical Pathologist.

Rob Schreiber

Rob Schreiber a Distinguished Technologist in the Advanced Architecture Program at Hewlett Packard Laboratories. He is known for important basic research in sequential and parallel algorithms for matrix computation, and compiler optimization for parallel languages. Rob was a professor of Computer Science at Stanford and at RPI. He was chief scientist and the lead architect at Saxpy Computer. He was a co-developer of the sparse matrix extension of Matlab, and a leading designer of the High Performance Fortran programming language. He was one of the developers of the NAS parallel benchmarks. He wrote the parallel matrix computation libraries at Maspar. At HP, Rob helped lead the PICO Project, a groundbreaking tool for embedded processor synthesis from high-level specifications. His current research is in algorithms for the exploration of large-scale datasets, and in programming tools for highly parallel clustered systems.

David Scott

Dr. David Scott is the Petascale Product Line Architect for Intel Corporation. He is responsible for formulating and driving Intel strategy for developing the technology needed for high end computing, including processors, platforms, interconnect, and software. He is Intel's interface to the programs around the world that are developing systems that scale to a Petaflop and beyond.

Dr. Scott joined Intel in 1985, developing algorithms and math libraries for a sequence of distributed memory message passing computers, culminating in the first teraflop on the Linpack benchmark in 1996. He worked for five years in the Floating Point Center of Expertise, helping to improve floating-point design and validation for all of Intel's 32 and 64-bit processor projects. He spent three years in Singapore supporting High Performance Computing in the Asia-Pacific region, before returning to Oregon to take his current position in July 2006.

David received his PhD in Numerical Analysis from UC Berkeley in 1978. He was a research scientist at the Oak Ridge National Lab for 3 years and taught in the Computer Sciences department of the University of Texas for four years before joining Intel. While working for Intel, he received his MBA from the University of Oregon in 1990.

Stephen Scott

Stephen L. Scott is a Senior Research Scientist in the Network and Cluster Computing Group of the Computer Science and Mathematics Division at the Oak Ridge National Laboratory (ORNL), Oak Ridge, USA. Dr. Scott's research interest is in experimental systems with a focus on high performance distributed, heterogeneous, and parallel computing. He is a founding member of the Open Cluster Group (OCG) and Open Source Cluster Application Resources (OSCAR). Stephen is presently the OCG steering committee chair and has served as the OSCAR release manager and working group chair. Dr. Scott is the lead principal investigator for the Modular Linux and Adaptive Runtime support for HEC OS/R research (MOLAR) research team. This multi-lab and multi-education institution research effort concentrates on adaptive, reliable, and efficient operating and runtime system solutions for ultra-scale scientific high-end computing (HEC) as part of the Forum to Address Scalable Technology for Runtime and Operating Systems (FAST-OS). Stephen has published numerous papers on cluster and distributed computing and has both a Ph.D. and M.S. in computer science. He is also a member of ACM, IEEE Computer, and the IEEE Task Force on Cluster Computing.

Chuck Seitz

Chuck Seitz, CEO of Myricom, earned BS, MS, and PhD degrees from MIT, and during the 17 years prior to founding Myricom in 1994 was a Professor of Computer Science at Caltech. Myricom products are based in part on the communication, switching, and software technologies that Chuck's Caltech research group developed under DARPA sponsorship for advanced multicomputers. Among Chuck's professional honors, he was elected to the US National Academy of Engineering in 1992 "for pioneering contributions to the design of asynchronous and concurrent computing systems."

Keith Seymour

Computer Science Department
1122 Volunteer Blvd
University of Tennessee
Knoxville, TN 37996-3450
seymour@cs.utk.edu

Keith Seymour is a Senior Research Associate at the Innovative Computing Laboratory (ICL) in the Computer Science Department at the University of Tennessee, Knoxville. He received his M.S. in Computer Science from the University of Tennessee in 1997. His research interests include grid computing and empirical code optimization.

Thomas Sterling

Dr. Thomas Sterling is a Professor of Computer Science at Louisiana State University, a Faculty Associate at California Institute of Technology, and a Distinguished Visiting Scientist at Oak Ridge National Laboratory. He received his PhD as a Hertz Fellow from MIT in 1984. Dr. Sterling is probably best known as the "father" of Beowulf clusters and

for his research on Petaflops computing architecture. He was one of several researchers to receive the Gordon Bell Prize for this work on Beowulf 1997. In 1996, he started the inter-disciplinary HTMT project to conduct a detailed point design study of an innovative Petaflops architecture. He currently leads the MIND memory accelerator architecture project for scalable data-intensive computing and is an investigator on the DOE sponsored Fast-OS Project to develop a new generation of configurable light-weight parallel runtime software system. Thomas is co-author of five books and holds six patents.

Vaidy Sunderam

Vaidy Sunderam is Professor of Computer Science and Chair of the Department of Mathematics and Computer Science at Emory University. His research interests are in parallel and distributed computing systems, software tools and architectures for metacomputing, high-performance message passing environments and infrastructures for collaborative computing. His prior and current research efforts focus on system architectures and implementations for heterogeneous computing middleware, collaboration tools, and fault tolerant distributed systems. He has been a principal architect of several software systems including PVM and other frameworks such as IceT, H2O, and Harness. Professor Sunderam teaches computer science at the beginning, advanced, and graduate levels, and advises graduate theses in the area of computer systems.

Bernard Tourancheau

Prof. Bernard Tourancheau on leave from University of Lyon - France
After his studies of applied mathematics and computer science Bernard Tourancheau got a PhD from Institut National Polytechnique of Grenoble. He then obtained a CNRS Researcher position at Ecole Normale Superieur de Lyon in 1989 and spent two years on leave at the University of Tennessee. In 1995, he got a Professor position at the University of Lyon where he created a research laboratory associated within INRIA, specialized in high speed networking and clusters. He joined SUN Microsystems Laboratories in 2000 as a Principal Investigator in the DARPA HPCS project. He is actually ending his sabbatical period, learning about research in renewable energy technologies.

Jeff Vetter

Jeffrey Vetter is a computer scientist in the Computer Science and Mathematics Division (CSM) of Oak Ridge National Laboratory (ORNL), where he leads the Future Technologies Group. His research interests lay largely in the areas of experimental software systems and architectures for high-end computing. Jeff earned his Ph.D. in Computer Science from the Georgia Institute of Technology; he joined CSM in 2003. Jeff is also an Adjunct Professor in the College of Computing at Georgia Tech.

David Walker

Prof. David W. Walker is head of the Distributed and Collaborative Computing group in the department of Computer Science at Cardiff University, and the Director of the Welsh e-Science Centre. Professor Walker's research interests focus on software, algorithms, and environments for computational science on high performance computers. He has been closely involved in the development of the ScaLAPACK parallel software library, and the MPI message passing standard. He has also contributed to the design of a parallel version of the Community Climate Model, and has published a number of papers on the parallel implementation of particle-in-cell algorithms for plasma simulations. He has also been involved in the benchmarking of science and engineering applications codes on parallel computers. Prof. Walker has published over 70 papers in the area of parallel computing and has co-authored three books on the subject.

Rich Wolski

Rich Wolski is an Associate Professor in Computer Science at the University of California, Santa Barbara (UCSB). Having received his M.S. and Ph.D. degrees from the University of California at Davis (while he held a full-time research position at Lawrence Livermore National Laboratory) he has also held positions at the University of California, San Diego, and the University of Tennessee. He is currently also a strategic advisor to the San Diego Supercomputer Center and an adjunct faculty member at the Lawrence Berkeley National Laboratory. Dr. Wolski heads the Middleware and Applications Yielding Heterogeneous Environments for Metacomputing (MAYHEM) Laboratory which is responsible for several national scale research efforts in the area of high-performance distributed computing and grid computing.

Hans P. Zima

Principal Scientist	Professor
Jet Propulsion Laboratory	University of Vienna
California Institute of Technology	Inst.of Scientific Computing
4800 Oak Grove Drive MS 602-149	Nordbergstrasse 15/C/4
Pasadena, California 91109-8099	A-1090 Vienna, Austria
Phone (818) 354-8980	Phone +43-1-4277-39401
E-Mail zima@jpl.nasa.gov	E-Mail zima@par.univie.ac.at

Hans P. Zima is a Professor of Applied Computer Science at the University of Vienna, Austria, and a Principal Scientist at the Jet Propulsion Laboratory, California Institute of Technology, in Pasadena, California. His major research interests have been in the fields of programming languages, compilers, and software tools. He led the design of of SUPERB, the first Fortran-based compilation system for distributed-memory architectures (1989) and was the chief designer of the Vienna Fortran language (1992) that became a major input for High Performance Fortran. His current research focuses on the design of an advanced programming language in the framework of the DARPA-sponsored HPCS project "Cascade" led by Cray Inc. Dr. Zima is the author or co-author of more than 160 publications, including 4 books.

Highland Lake Inn

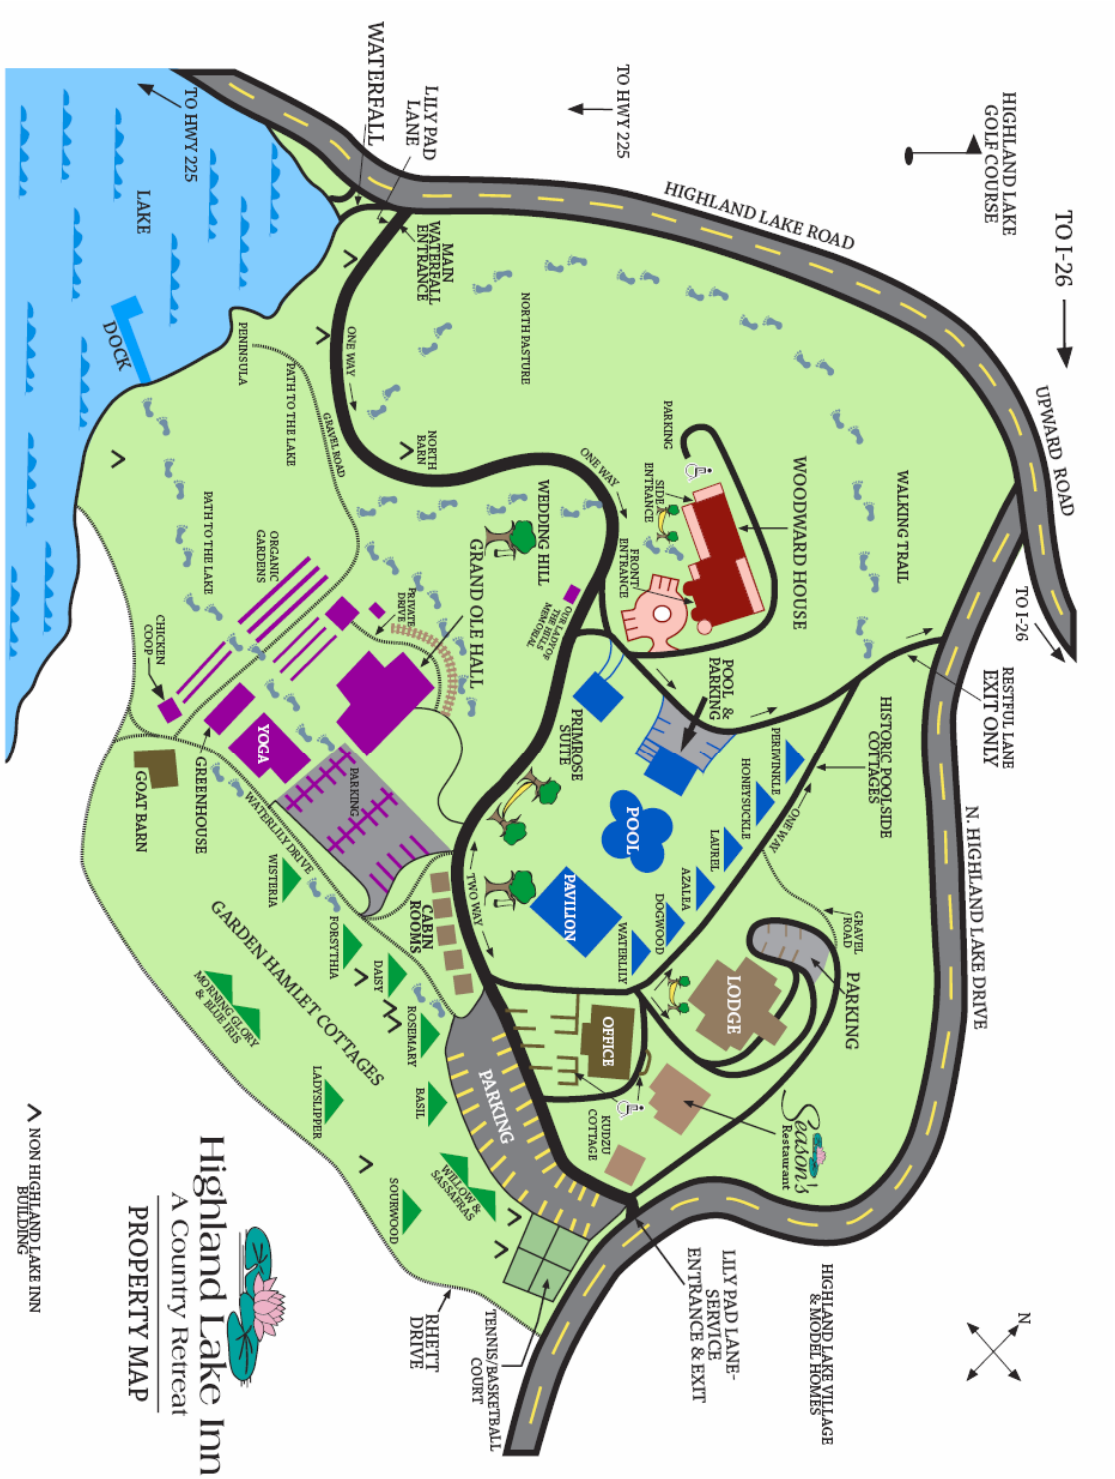

CCGSC 1998 Participants, Blackberry Tennessee

CCGSC 2000 Participants, Faverges, France

CCGSC 2002 Participants, Faverges, France

CCGCS 2004 Participants, Faverges, France