

HP-CAST

*HP Consortium for Advanced Scientific and Technical Computing
World-Wide User Group Meeting*

ISS Hyperscale and HPC Organization

*November 11th-12th, 2011
HP-CAST 17*

Book of Abstracts

Note: For some presentations no abstract has been received. The abstracts are listed according to the sequence of the agenda version 3.1.

Friday, Nov. 11th, 2011Hewlett-Packard**Title:** HPC and Hyperscale Executive Updates**Speaker:** Paul Santeler, Vice President & General Manager, Hyperscale Industry Standard Servers and Software, HP

Bio: Paul Santeler is Vice President and General Manager of HP's Hyperscale business within the Industry-standard Servers & Software (ISS) business unit. In this role, Paul is responsible for accelerating HP's growth with the emerging independent cloud and web 2.0 service providers, and High Performing Computing segments. He is chartered with bringing innovative solutions to market and optimizing HP's supply chain and services for this new business model. Paul's career in the technology industry spans more than 30 years and includes 15 years at HP and Compaq from 1989 to 2004, where he held key executive positions within the desktops, ProLiant servers and management software businesses. Notably, he created the 8-processor Compaq ProLiant server product category, achieving 60 percent market share and \$1B in revenue within 18-months. He increased customer usage of HP Insight Manager from 34 to 70 percent and started the ProLiant Essentials software franchise which is now a multi-billion dollar business. For the past three years, Paul served as Vice President and Chief Technical Liaison at AMD where he provided technical strategy, communication, and management oversight for the overall HP account. Prior to AMD, Paul was Foxconn's Vice President of North America Operations running the region's factories and desktop ODM business. Paul has a Bachelor of Science in Electric Engineering from Rensselaer Polytechnic Institute (RPI) in Troy, NY.

Hewlett-Packard**Title:** Converged Infrastructure for HPC**Speaker:** Scott Misage, HPPurdue University**Title:** Upcoming HPC Technologies at Purdue**Speaker:** Mike Shuey

Abstract: Each year, Purdue University deploys a new compute cluster, to meet the needs of an increasingly diverse research community. As with most universities, facilities and funding are at a premium. HP, in partnership with Intel and Mellanox, helped Purdue deliver an extremely efficient HPC platform for Purdue's 2012 computing needs. Purdue's "Carter" cluster uses the latest FDR Infiniband technology, Intel's Xeon E5 processors, and HP's SL6500 platform, to provide an exceptionally capable system for campus research.

Consultant**Title:** A New Large Scale HPC System in Iceland**Speaker:** Thorvaldur Sigurdsson

Bio: International IT experience including 23 years in the US and Europe. Former CTO and Managing Director of IT operations at Skyrr which hosts critical IT infrastructure for the Icelandic government, municipalities and several companies in Iceland. Successful entrepreneur and new business development participant of several key projects and companies. Electrical Engineer graduate from the University of Iceland.

Abstract: Iceland hosts the new large scale Nordic-HPC environment based on HP blade solutions. Why? What are the advantages and limitation of Iceland as a future data center hub?

Boeing**Title:** Boeing HPC: Initial Experiences with HP and Bright**Speaker:** Greg Siekas

Bio: Greg joined The Boeing Company back in 2004 and began his career with the Boeing High Performance Computing team. This group has the responsibility for managing the overall Enterprise HPC service that is available to all engineering groups across Boeing. Prior to Boeing he worked as a Field Technical Analyst for SGI. Greg has a Masters degree in Biomedical Engineering from Tulane University. He started out as an end user of High Performance Computing resources during his undergraduate and graduate years modeling cardiac defibrillation. Greg has been busy since the birth of his son this past summer.

Abstract: In many industries, HPC-driven engineering is a crucial competitive advantage and a key success factor. Investment in HPC must address current needs, anticipate future requirements, and, at the same time, work within the resource constraints of the organization. The Enterprise HPC group of The Boeing Company needed solutions to provide the next generation HPC platforms. The

company had a wide variety of computational requirements to serve the needs of disparate and distributed engineering teams. We needed a reliable standard platform that could be applied throughout the organization. This strategic HPC platform must provide the flexibility to support every department; deliver maximum throughput to meet the demands of the business; ensure maximum security; streamline systems management; and yield improved return-on-investment.

The initial implementation involved multiple HP clusters with Bright Cluster Manager, including a 500+ node system and two smaller clusters. This rollout will expand over the next year. This presentation will focus on the Boeing HPC teams initial experiences in deploying a HP and Bright HPC platform. What we learned as we rolled out these new systems and the benefits they provided.

MIMOS Berhad

Title: Experiences and Challenges at MIMOS

Speaker: Hong Ong

University of Utah

Title: HPC Environments at the University of Utah

Speaker: Brian Haymore

Bio: I've been working for the Center for High Performance Computing at the University of Utah for nearly 13 years. I am the team lead for both our High Performance Computing and Storage resources. CHPC's mission on campus is to support and advance research computing.

Abstract: This presentation will cover the scope of research computing CHPC has provided on the University of Utah campus for about the past 15 years. We will examine the evolution of both computational needs as well as storage needs through this time. The Institute for Clean and Secure Energy will be presented as a use case to examine the ever changing and growing needs for both computational and storage resources. HP's X9000 product will be highlighted for its' role in CHPC's ability to meet and exceed the storage needs of the research community on campus.

University of Széchenyi István

Title: Setting-up Complex HPC-Clusters: Practice and Experience

Speaker: Dr. Gábor Élő, PhD

Bio: Gabor has a wide range of professional activities on the fields of information management and mobile communication. Started with Nokia Research, continued with Philis and some new ventures with partners from India. Gabor now is acting as associate professor in 'Szechenyi Istvan University' and a leader of INFCARE8 nation wide R&D&I project.

Attila Haraszti, MSc

Attila has a long history with university activities back to 32 years. Started with the Central Institute for Physics continued on Budapest University of Technology and Economics and has a long trip in Computing industry in Digital Equipment, Compaq and HP. Attila is now helping the INFCARE8 R&D activities in SZE University specifically focused on supercomputing.

Abstract: This presentation contains interesting theoretical and practical information on a project supporting new model of EEG analyzing activities in healthcare. We share some information on planning, designing phase, implementation phase and giving some experience on the first 6 months of operation. This project using the second largest HPC system in Hungary based on HP AMD blade technology and first implementation of HP's cluster file solution coming from IBRIX.

MIT Lincoln Laboratory

Title: Experiences with an Interactive, Large-Scale POD-based HPC system

Speaker: Andrew McCabe

Bio: Andrew McCabe is a System Engineer at MIT Lincoln Laboratory. He is the Lead Technical Architect for the LLGrid Program, which delivers On-Demand, Interactive High Performance Computing Resources to Scientists and Engineers in Numerous Disciplines. Mr. McCabe Earned a Bachelor's Degree in Computer Science and Spent 15 Years Developing Trading and Risk Management Systems prior to joining MIT.

Abstract: MIT Lincoln Laboratory has chosen to partner with HP on it latest Supercomputer – TX-Green Pioneer. TX-Green is an HP POD Based System that will be deployed in the City of Holyoke in Massachusetts's Pioneer Valley. We expect TX-Green and the systems to be deployed in Holyoke in the future to have a major impact on the Carbon Footprint of both MITLL's LLGrid and MIT as a whole.

Yale University

Title: Experience with the Yale Omega Cluster

Speaker: Andrew Sherman

Bio: Dr. Andrew Sherman is a Research Scientist and Lecturer in the Department of Computer Science at Yale University, where he works with users of Yale's high performance computing (HPC) clusters. He rejoined the Yale faculty in 2009 after spending more than 20 years with several small companies developing HPC software tools and applications. By training, Dr. Sherman is a numerical analyst and computer scientist, having received his Ph.D. in Computer Science from Yale in 1975. He has worked extensively on HPC algorithms and software in both academic and commercial environments, and he has authored numerous publications on numerical linear algebra and high performance scientific computation, covering both basic methodologies and applications to petroleum reservoir simulation, bioinformatics, and other scientific and engineering disciplines. In addition, Dr. Sherman holds several patents related to parallel search algorithms and distributed workflows.



Abstract: This presentation will describe the Yale Omega cluster, focusing on Yale's custom-built physical infrastructure and experience with the cluster to date.

Wellcome Trust Sanger Institute

Title: Sanger's Compute and Storage Environments: Status, Trends and Challenges

Speaker: Phil Butcher

Bio: I have been responsible for the IT at the Sanger institute for almost two decades now. The institute has approximately 850 staff and is located in the United Kingdom at the Wellcome Trust Genome Campus. In the early 1990's Sanger became the UK contingent in the Human Genome project and by the turn on the century had deposited about one third of the publicly available sequence data in public domain databases. The institute's sequencing capability has increased significantly since then and in previous years our infrastructure has grown accordingly. I now manage a team of 50 IT staff who are meeting new challenges head on.



Abstract: The Wellcome Trust Sanger Institute has been at the fore-front of very large scale DNA sequencing for many years. Sequencing has now become more widely available and over the past 3-4 years we have had discussions with many smaller labs who have taken on NGS technologies. Lots of press has been given to the very large data production rates and like many of the large sequencing groups we have had to run fast to keep up. This presentation discusses our Infrastructure, our ambitions and challenges coming our way.

Advanced Micro Devices (AMD)

Title: AMD Processor Technology & Roadmap Update

Speaker: William C. Brantley, Ph.D.

Bio: Bill Brantley completed his Ph.D. at Carnegie Mellon University in ECE after working 3 years at LANL. Next, he joined the T.J. Watson Research Center where he began a project which led to the Vector instruction set extensions for the Z-Series CPUs. Then he was one of the architects of the 64 CPU RP3 (a DARPA supported HPC system development in the mid-80s) and led the processor design including a hardware performance monitor. Later he work on a number of projects at IBM mostly dealing with system level performance of RISC 6000 and other systems, eventually joining the Linux Technology Center. In 2002, he joined Advanced Micro Devices where he is a Principal Member of Technical Staff and helped to launch the Opteron. Since 2004 he has been the manager of a performance team focused on HPC where he contributes to both current and future products.



Intel Corporation

Title: Intel in HPC: Roadmap and Processor Update

Speaker: Michael Haedrich, WW Marketing Manager for HPC

Bio: Michael Haedrich is the World Wide Marketing Manager for High Performance Computing at Intel. He is responsible for the Intel HPC Roadmap and ensuring the HPC community is positioned to take advantage of the Intel technology. In addition, Mike works closely with the HPC ecosystem to develop innovative solutions to enable user to experience faster simulation and enjoy the benefit of greater discoveries. Michael has been with Intel for over ten years with positions in Planning as well as Marketing. His passion for HPC goes well beyond Marketing to include customer advocacy and the HPC blog he regularly contributes to on the Intel Community site.

Abstract: The HPC market is very diverse. Usage models range from creating water bottles to simulating weather patterns. Intel has taken a hard look at this diverse market and will share their findings on



workloads. Intel will also provide detail on their solutions and how they are supporting High Performance Computing.

NVIDIA Corp.

Title: NVIDIA GPU Technology & Roadmap Update

Speaker: Thomas E. Reed, Solution Architect

Bio: Tom Reed is NVIDIA's solution architect for the US DOD and IC focusing on education and awareness of GPU technology and enabling customer adoption of GPU computing with those communities. Mr. Reed's research at NVIDIA centers on scalable high performance grid and cluster systems for big data problems. Prior to joining NVIDIA, Mr. Reed was with Silicon Graphics for 18 years where he held various positions including director of engineering for visualization, director of professional services, and technology center manager. Past research has included initiatives in interactive parallel computing, real-time visual simulation, large multi-display environments, the adoption of cluster and commodity technologies within SGI, and hybrid computing architectures. Mr. Reed spent as much time in customer facing roles as research and development and valued the insight gained through interaction with SGI's leading edge customers. Mr. Reed's career has always involved some combination of high performance computing (HPC) and visualization and he actively promotes the fusion of best of breed technologies in a data centric architecture as an important tool in advancing IT capabilities. In addition to SGI, Mr. Reed has worked for McDonnell Douglas, General Dynamics, and Linuxcare as well as research contracts with Kirkland Air force Base, Arnold Engineering Development Center (AEDC), and NASA. Mr. Reed attended Iowa State University for both undergraduate and graduate studies, majoring in aerospace engineering with graduate research in Computation Fluid Dynamics (CFD).

Abstract: This short presentation will discuss the following topic areas as a way of providing a high level overview of the current and future GPU computing products being fielded by HP and NVIDIA:

- Current status of GPU products going into HP system (Fermi)
- Recent developments and announcements (top 500, ISV adoption)
- Roadmap of GPU technology
- Roadmap of CUDA technology

Mellanox Technologies

Title: FDR – A New Trend in High-End Interconnects

Speaker: Gilad Shainer & Todd Wilde, Mellanox

Bio: **Gilad Shainer** is a Senior Director of HPC and Technical Computing at Mellanox Technologies focuses on high-performance computing, high-speed interconnects, leading-edge technologies and performance characterizations. Mr. Shainer holds M.Sc. degree (2001, Cum Laude) and a B.Sc. degree (1998, Cum Laude) in Electrical Engineering from the Technion Institute of Technology in Israel. He also holds several patents in the field of high-speed networking and the Chairman of the HPC Advisory Council.

Todd Wilde has over 20 years of experience in the high-tech networking industry as a silicon designer and technical support specialist. Todd is currently the Director of Technical Computing and HPC for Mellanox Technologies. During the past 6 years Todd he has managed Mellanox Technologies Field Application Engineering team where he led the support of multiple InfiniBand equipment designs using Mellanox leading edge silicon. Todd has been an evangelist and leading industry expert in InfiniBand since its inception in 2001. Mr. Wilde holds a B.S.E.E. from University of South Florida.

Abstract: Large scale HPC systems will span tens-of-thousands of nodes, all connected together via high-speed connectivity solutions. With the growing size of clusters and CPU/GPU cores per cluster node, the interconnect needs to provide not only the highest throughput and lowest latency, but to be able to offload the processing units (CPUs, GPUs) from the communications work in order to deliver the desired efficiency and scalability. Mellanox Scalable HPC solution accelerate MPI and SHMEM environments with smart offloading techniques and deliver the needed infrastructure for faster and more efficient GPU communications. The presentation will cover the latest FDR InfiniBand technology and solutions from Mellanox and the integration of it with HP platforms to deliver the most efficient, highest performance solutions for the next generation HPC systems.

Hewlett-Packard

Title: HP ISS Platform Trends and Product and Roadmap Updates

Speaker: Ed Turkel, HP

Hewlett-Packard**Title:** From Petascale to Exascale – Challenges, Solutions and Outlook**Speaker:** Moray McLaren & Mike Tullis, HP**Bio:** Mike Tullis: 30 year computer R&D veteran, 17+ years at Hewlett-Packard companies, specializing in engineering management for compute platform R&D and advanced technologies.**Moray McLaren:** Distinguished Technologist with HP Labs, with a background in supercomputer system development and interconnect for some of the world's fastest computers.**Abstract:** Hewlett-Packard Company expects to play a key role in extending the HPC compute environment out to the exascale performance level. The timeline to exascale is being set aggressively by those who see the compelling benefits around world. We will focus research and development in key areas such as silicon photonics for interconnect technology and memristor for solid state non-volatile memory storage. HP will also make investments to solve other important challenges of exascale including resiliency, power management, system management, packaging/cooling, and fabric/switching technologies. In this presentation we will show our vision for node level architecture and discuss some important questions around software for exascale platforms. Over the next 10 years HP will co-invest with both private and public sectors in providing high end HPC compute solutions which will be deployed on critical missions benefiting all mankind and enhancing USA national security. The technologies developed in this endeavor will be leveraged across the entire compute spectrum from cell phones to super computers.Hewlett-Packard**Title:** Thermal Challenges of Future HPC System**Speaker:** Tahir Cader, HP**Bio:** Dr. Tahir Cader is a Power & Cooling Strategist within HP's Industry Standard Servers (ISS) organization, dealing with Power & Cooling issues from the product through to the data center level for HP's Hyperscale and HPC market segments. In particular, his emphasis has been on issues relating to the end-to-end management of data centers (including PODs/EcoPODs) and data center energy efficiency, encompassing existing and future technologies while working closely with HP Labs. Dr. Cader is also active with external standards bodies and public policy organizations. He has been a member of The Green Grid's BOD, is a member of The Green Grid's Technical Committee, is The Green Grid's Liaison to ASHRAE TC9.9, and is a member of ASHRAE Technical Committee 9.9. With over 18 years of experience in the thermal management and data center industries, Dr. Cader is both a sole inventor as well as a co-inventor on 20 issued and additional filed patents, and is a co-author for more than 45 peer-reviewed journal, conference, and trade journal technical articles. He was the lead editor/author for the ASHRAE/TGG joint book entitled "Real-Time Energy Consumption Measurements In Data Centers", and was also a significant contributor to several published ASHRAE Datacom Series books.**Abstract:** Server and rack power densities for HPC continue to rise rapidly. While the power levels for individual components such as CPUs, DIMMs, and HDDs are climbing relatively slowly, it is the rapid adoption of GPU-based computing that is contributing to the fast increase in server and rack power levels. Today's rack densities are now providing significant facilities thermal challenges, making it extremely difficult to provide adequate cooling without resorting to extreme measures. A major trend that is developing is a migration to containment solutions, i.e., enclosed racks, aisle containment, and even containerized data centers. For Exascale computing, preliminary analysis suggests the need to provide liquid-cooled IT in order to be able to support the projected rack power densities. The presentation will cover all of these major challenges and trends.M.D. Anderson Cancer Center**Title:** IT Challenges in supporting Personalized Cancer Medicine**Speaker:** Krishna Sankhavaram, Director, Research Information Systems & Technology Development, University of Texas M.D. Anderson Cancer Center**Bio:** Krishna Sankhavaram brings 20 years of experience building software solutions and delivering IT infrastructure services to support Basic and Translational research in cancer research environments. He has a broad background in many technologies with Solution architecture skills blended with extensive experience in project management. He spent more than a decade at St. Jude Children's Research Hospital, where he helped build the IT infrastructure for research at the Hartwell center for Bioinformatics and Biotechnology. He is the Director for Research Information Systems & Technology Development at MD Anderson Cancer Center. He helped build a large computing cluster for research, with a rapidly growing storage environment. He leads SOA based integrated software framework for researchers that integrates clinical and research areas.

Abstract: An introduction and discussion of the IT challenges we face in supporting personalized cancer medicine. The data tsunami that we have to deal with, the role of high performance cluster computing and impact will be discussed with technical details on the storage, and computing resources available at MD Anderson to support cancer research. I will also show our software framework that researchers use to work in this private cloud environment with a cluster.

Alcatel-Lucent

Title: A 100 Gbit/s WAN Research Network between Dresden and Freiberg Universities

Speaker: Rolf Sperber

Bio: End to end network design, focus on Research networks

Abstract: The main objective of the project was to demonstrate an industry first 100GE link with commercially available equipment; 1830PSS DWDM System and 7750SR Service Router. TU Dresden and TU Freiberg have linked their two high-performance data centres equipped with a cluster of 16 HP servers each with 100GE. Here we not only demonstrated rare speed but also the capability to fully utilize this bandwidth with real life applications.

- Synthetic Load tests: 100 % of bandwidth utilized
- Federated Datasets: Proof of concept, nearly 100 % utilization (98 %)
- Virtualization of Servers: Feasibility proof, latency induced limitations when going to larger distances
- QoS tests: Proof of policy implementations
- WAN acceleration: Feasibility Proof

Hewlett-Packard

Title: Implementation and Usage of Social Media Networks for HP-CAST attendees

Speaker: Jill Sweeney

Abstract: Learn about the new HP-CAST community on LinkedIN. Find out what it has to offer, how to join and how you can become a part of the HP-CAST conversation.

Saturday, Nov. 12th, 2011

Tutorial A1: Accelerators – Multi-Core Processors

NVIDIA

Title: NVIDIA Processor/GPU Update

Speaker: Thomas E. Reed, Solution Architect

Bio: Tom Reed is NVIDIA's solution architect for the US DOD and IC focusing on education and awareness of GPU technology and enabling customer adoption of GPU computing with those communities. Mr. Reed's research at NVIDIA centers on scalable high performance grid and cluster systems for big data problems. Prior to joining NVIDIA, Mr. Reed was with Silicon Graphics for 18 years where he held various positions including director of engineering for visualization, director of professional services, and technology center manager. Past research has included initiatives in interactive parallel computing, real-time visual simulation, large multi-display environments, the adoption of cluster and commodity technologies within SGI, and hybrid computing architectures. Mr. Reed spent as much time in customer facing roles as research and development and valued the insight gained through interaction with SGI's leading edge customers. Mr. Reed's career has always involved some combination of high performance computing (HPC) and visualization and he actively promotes the fusion of best of breed technologies in a data centric architecture as an important tool in advancing IT capabilities. In addition to SGI, Mr. Reed has worked for McDonnell Douglas, General Dynamics, and Linuxcare as well as research contracts with Kirkland Air force Base, Arnold Engineering Development Center (AEDC), and NASA. Mr. Reed attended Iowa State University for both undergraduate and graduate studies, majoring in aerospace engineering with graduate research in Computation Fluid Dynamics (CFD).



Abstract: This presentation will provide a technical overview of the progress in GPU computing over the last year and what to expect in the coming year. Key topics will include:

- Review of GPU computing over the last 12 months
 - o GPU HW
 - o Systems HW & SW
 - o GPU computing software
- Current and future GPU capabilities from NVIDIA
- Current and future directions for CUDA

Intel**Title:** Intel's Programming Models for Multicore and Many-core Processors**Speaker:** Michael McCool

Bio: Michael McCool has degrees in Computer Engineering (University of Waterloo, BSc) and Computer Science (University of Toronto, M.Sc. and PhD.) with specializations in mathematics (BSc) and biomedical engineering (MSc) as well as computer graphics and parallel computing (MSc, PhD). He has research and application experience in the areas of data mining, computer graphics (specifically sampling, rasterization, texture hardware, antialiasing, shading, illumination, and visualization), medical imaging, signal and image processing, financial analysis, and languages and programming platforms for high productivity parallel computing. In order to commercialize research work into many-core computing platforms done while he was a professor at the University of Waterloo, in 2004 he co-founded RapidMind, which in 2009 was acquired by Intel. Currently he is a Software Architect with Intel and an Adjunct Associate Professor with the University of Waterloo.

Abstract: Intel has taken bold steps to simplify parallel programming by providing a broad and powerful set of parallel programming models that are applicable to both multicore and many-core processors. This talk will survey the key ones: Intel® Threading Building Blocks (TBB), OpenMP, Intel® Cilk™ Plus, OpenCL, Intel® Math Kernel Library, Intel® MPI Library and Co-array Fortran. There are also a number of other models in research and development. Each of these models is appropriate in different contexts. Due to the flexibility of Intel processors, users can select the most appropriate tool for their application. Cilk Plus supports a lightweight fork-join task model with a simple syntax that supports irregular computation well, as well as explicit vectorization. TBB supports wide deployment and efficient load balancing on existing compilers, and supports a wide range of parallel constructs. The flexible support of composable programming models results in a superior solution for today's C, C++ and Fortran developers that can be deployed in nearly any application, and that works for both multicore and many-core devices.

Hewlett-Packard**Title:** GPU Acceleration of Large Scale Business Intelligence and Analytics**Speaker:** Ren Wu, HP Labs

Bio: Dr. Ren Wu is a Senior Research Scientist and the PI of the CUDA research center at HP Labs, Palo Alto. His research interests include data-intensive high-performance computing, massively parallel algorithms, machine learning and computational intelligence. In recent years he has been focusing on GPU acceleration of large scale analytics, and is well known for his work on GPU-accelerated clustering algorithms.

Abstract: In this session, we will share our research on using GPUs as accelerator for BI analytics. Using a set of the well-known BI primitives, such as very large scale clustering, predicate evaluation etc, we will discuss in details on how the GPUs can be used here for significantly better performance. Special attention will be paid to the algorithm and data structure design to better utilizing GPU's memory hierarchy, which is vital for achieving significant speedups over CPU-only approach.

[Tutorial B1: Interconnects](#)Mellanox Technologies**Title:** InfiniBand FDR Technology – Accelerations, Offloading and Environments towards Exascale**Speaker:** Gilad Shainer & Todd Wilde, Mellanox

Bios: **Gilad Shainer** is a Senior Director of HPC and Technical Computing at Mellanox Technologies focuses on high-performance computing, high-speed interconnects, leading-edge technologies and performance characterizations. Mr. Shainer holds M.Sc. degree (2001, Cum Laude) and a B.Sc. degree (1998, Cum Laude) in Electrical Engineering from the Technion Institute of Technology in Israel. He also holds several patents in the field of high-speed networking and the Chairman of the HPC Advisory Council.



Todd Wilde has over 20 years of experience in the high-tech networking industry as a silicon designer and technical support specialist. Todd is currently the Director of Technical Computing and HPC for Mellanox Technologies. During the past 6 years Todd he has managed Mellanox Technologies Field Application Engineering team where he led the support of multiple InfiniBand equipment designs using Mellanox leading edge silicon. Todd has been an evangelist and leading industry expert in InfiniBand since its inception in 2001. Mr. Wilde holds a B.S.E.E. from University of South Florida.

Abstract: Large scale HPC systems will span tens-of-thousands of nodes, all connected together via high-speed connectivity solutions. With the growing size of clusters and CPU/GPU cores per cluster node, the interconnect needs to provide not only the highest throughput and lowest latency, but to be able to offload the processing units (CPUs, GPUs) from the communications work in order to deliver the desired efficiency and scalability. Mellanox Scalable HPC solution accelerate MPI and SHMEM environments with smart offloading techniques and deliver the needed infrastructure for faster and more efficient GPU communications. The presentation will cover the latest FDR InfiniBand technology and solutions from Mellanox and the integration of it with HP platforms to deliver the most efficient, highest performance solutions for the next generation HPC systems.

QLogic Corporation

Title: QLogic TrueScale Future for High Performance Computing Clusters

Speaker: Philip Murphy, Vice President, HPC Technology, QLogic Corporation

Bio: Philip Murphy, vice president of HPC technology at QLogic Corporation, is responsible for ensuring that the company's high performance computing interconnect products remain on the forefront of the industry. Prior to his current role, Mr. Murphy was vice president of engineering within QLogic's Network Solutions Group, responsible for the design and development of all high performance computing products, as well as all storage area network switching products. Before joining QLogic, Mr. Murphy was vice president of engineering at SilverStorm Technologies, which he co-founded in 2000 and which was acquired by QLogic in 2006. SilverStorm's core focus was on providing complete network solutions for high performance computing clusters. Phil is a pioneer in developing InfiniBand infrastructure in high performance cluster environments. Prior to co-founding SilverStorm, Mr. Murphy was director of engineering at Unisys Corporation and responsible for all I/O development across the company's diverse product lines. Mr. Murphy holds an MS degree in Computer and Information Science from the University of Pennsylvania.

Abstract: Highly capable, power-efficient, and cost-effective systems are required to realize the goals of the next generation high performance computing clusters. A key enabling technology on the path to next generation computing is the compute interconnect network, which is critical to the performance of many application environments. This talk will cover the evolution of QLogic's TrueScale InfiniBand interconnect, highlighting the key architectural approaches enabling cost-effective, power-efficient scaling today and enhancements allowing scaling to extreme levels in the future.

Hewlett-Packard

Title: Scalability Investigations on an IB Cluster

Speaker: Zarka Cvetanovic, HP

Bio: Zarka has been involved in the design, architecture, and performance evaluation of a number of HP, Compaq, and DEC server products. She is currently leading efforts in software and tools development for interconnects in HP/HPC. Zarka holds a PhD degree in Computer Engineering from University of Massachusetts, Amherst.

Abstract: In this session, we will discuss results from our scalability investigations on the 704-node cluster from HP with QLogic Infiniband. This work is a result of joint collaboration between HP and QLogic. The characteristics of applications investigated include the following: one-sided programming paradigm (SHMEM/UPC), random and high-communication (all-to-all) messaging patterns, short messages, and use of large memory per node. The goal of this study goal was to determine whether InfiniBand technology can enable use of commodity clusters for such applications. Our intent was also to provide additional confirmation that this solution will ultimately provide good performance at 4K-node production scale.

Based on the results from this study, we were able to achieve the following:

- Optimal performance and scaling for 8KB messages that was achieved by enhanced routing algorithms
- Substantial improvements in message rate achieved by optimizing communication protocols
- Low memory footprint
- Comparable performance with large and small memory

We will present several software optimizations, including routing and protocol enhancements, that resulted in substantial performance improvements. Majority of these enhancements are incorporated in the current/future product releases of InfiniBand software from QLogic.

We will discuss detailed performance and scaling results across a number of different workloads that led us to conclude that InfiniBand is a viable candidate interconnect for deploying large-scale commodity clusters for one-sided programming paradigm.

[Tutorial C1a: HPC Clouds, Grids and Scheduling](#)

[Hewlett-Packard](#)

Title: The Status of HPC and Cloud
Speaker: N.N., HP

[Adaptive Computing](#)

Title: Creating a Dynamic HPC Cloud with CMU and Moab
Speaker: Scott Hurst, Director, Alliances
Abstract: Don't miss an opportunity to learn how Moab Adaptive HPC Suite™ with HP CMU provides a powerful, intelligent cluster and workload management solution for HPC workloads in a single cluster. See how together, HP and Adaptive Computing can bridge to resources from remote clusters to create a dynamic HPC Cloud to provide extended resources to users.

[Altair Engineering](#)

Title: HPC Cloud and Addressing the "Missing Middle"
Speaker: Scott J. Suchyta, Director, Partner Solutions & Integration
Bio: Scott Suchyta is the Director of Partner Solutions and Integration at Altair Engineering, Inc. He is responsible for managing the technical relationships with Altair PBS Works division's hardware and software partners. Scott supports partners' pre/post-sales engagements, helps resolve technical issues, ensures roadmap alignment, and participates in technical marketing activities. Additionally, he identifies and drives integration opportunities with partners' products that result in differentiated joint solutions. Scott is a 11 year veteran at Altair having previously worked as the PBS Works product manager responsible for PBS Professional, PBS Application Services, and PBS Portal products. Scott started his career in Altair's PBS Works division as an application engineer implementing custom solutions using PBS Professional.
Abstract: There is a significant segment of the market, the "missing middle", that is currently not well addressed by HPC solution providers. The missing middle is known to be the foundation of production, process, design and innovation across every manufacturing sector around the world. This presentation will explore the issues of the "missing middle" and present potential solutions addressing their requirements. Examples of implementations which have started to bridge the gap for the "missing middle" will be high-lighted.

[Intel Corporation](#)

Title: Update on an HPC Cloud Initiative
Speaker: Michael Haedrich, Intel & Ed Turkel, HP

[Tutorial C1b: Platform Computing Update – Pending Acquisition](#)

[Platform Computing](#)

Title: Platform Computing Update – Pending Acquisition
Speaker: Tripp Purvis, Dave Sherrill, Platform Computing

[Tutorial D1: Parallel File Systems for HPC](#)

[Hewlett-Packard](#)

Title: Using the HP X9000 Scalable NAS for Mixed Workloads in HPC Environments
Speaker: Sean Cochrane, X9000 CSE (Corporate Systems Engineer)
Bio: Sean Cochrane is currently a Storage Architect in the X9000 Corporate Systems Engineering team at Hewlett Packard. Sean has been working in the computer industry for over 15 years. Previous to his time at HP, Sean worked at Sun Microsystems. During much of his time at Sun, he focused on Storage Architectures in High Performance Computing and was the principal architect for the Sun Lustre Storage System.
Abstract: The task of storing large amounts of data has increasingly become a challenge in building productive HPC environments, which put unique demands on storage systems. This session will focus on the storage needs of typical HPC environments. Participants will learn when and how to use the X9000 scalable NAS systems to address the different workloads particular to HPC environments.

DDN**Title:** DDN Lustre Technology Update**Speaker:** Keith Miller, Technical Director, WW HPC & Life Sciences, DataDirect Networks**Bio:** Mr Miller is a HPC industry veteran with over 17 years of experience. He has served in several roles with DDN and currently is the Technical Director, WW HPC and Life Sciences. This role includes managing the entire field technical team for HPC as well as oversight of the HP OEM relationship. Mr Miller joined DDN over 5 years ago after serving several roles at SGI.**Abstract:** An update on the HP-DDN Lustre Solutions offering including the latest performance measurements, experience from the field and future directions. Attendees will get a clear understanding of the performance and operational advantages of these solutions, examples of successful deployments in multiple markets and information about DDN's next generation systems.**Tutorial A2: Accelerators – S/W Environments****PGI****Title:** High-level Programming for GPU Accelerators**Speaker:** Douglas Miles**Bio:** Douglas Miles is responsible for all business and technical operations of The Portland Group (PGI). He has worked in various positions over the last 25 years in HPC applications engineering, math library development and technical marketing at Floating Point Systems, Cray Research Superservers, PGI and STMicroelectronics.He can be reached by e-mail at douglas.miles@pgroup.com**Abstract:** Programming GPU accelerators involves 3 basic aspects: splitting the source code between host and GPU, managing data allocation and movement between host memory and GPU memory, and optimizing GPU kernels. This talk discusses how the first two aspects can be mostly automated using modern compiler technology, and how compiler feedback can assist with, but probably never automate, the third.**Rogue Wave Company****Title:** ReplayEngine - Deterministic Debugging for HPC**Speaker:** Nikolay Piskun**Bio:** Nikolay Piskun, Ph.D., Director of Continuing Engineering, has been with TotalView Technologies for almost ten years. He has a Doctor of Philosophy in Physics from Ohio State University and is specializing in parallel computing. Nikolay works closely with the product management team to set the direction for development of the entire family of TotalView Technologies products. He has provided tutorials, workshops, and presentations at numerous technical conferences, including SuperComputing and International Supercomputing for the past several years.**Abstract:** This talk will show you how deterministic replay technology can radically simplify debugging on multi-core HP ProLiant servers and InfiniBand accelerated BladeSystem clusters. Deterministic replay enables reverse debugging so you can stop any program and explore not just the current state of the program but also the states that lead up to the current state. This breaks down the cycle of having to restart programs over and over again during debugging. Now debugging can be done within a single session — you can let the program run to the point where the program fails, either by crashing or computing an invalid result. Then work backwards along the recorded execution sequence to quickly pin down the code defect that is the root cause of the failure. This technique works with serial, multi-threaded or MPI parallel applications. See how TotalView's ReplayEngine add-on can free you up from both the drudgery of debugging by speculative placement of print statements and the "breakpoint + restart + refine" cycle.

This talk will also provide an update on Rogue Wave's support for debugging applications that use NVIDIA GPUs and CUDA. With CUDA 4.0's new capability to define a unified virtual address block between GPU devices on a node, an increasing number of scientists and programmers will need to take advantage of TotalView's superior support for debugging multi-process, multi-node and multi-device CUDA programs.

Allinea Software**Title:** New Generation of Development tools to enable HPC from the Desktop to the Extreme**Speaker:** Jacques Philouze

Bio: Jacques, as a graduated Engineer in Electronics, started his professional career in a Sales & Marketing environment at Analog Devices where he acquired his current sense of entrepreneurial spirit, focus and being result driven. He moved to the IT industry in early 1990 with Sun Microsystems where he held leadership positions bringing both expertise and experience, particularly in the role of Managing Director for Southern Europe at Platform Computing. This successful track record in building business/organization internationally, in a fast growing environment, prepared him well in setting up and developing Sales & Marketing at Allinea Software from its very beginning in early 2004. Jacques is Global Vice President in Sales & Marketing at Allinea Software Ltd., which has become a leading supplier of tools for parallel programming and high performance computing (HPC).



Abstract: With the growing demand for performance while lowering the power consumption, Hybrid technology and large scale computers are becoming the standard. However this is drastically increasing the level of available information to the developers while debugging. Old fashioned tools now need to be revised. This talk will present the challenges developers are actually facing and the solution that Allinea is bringing to make debugging a reality at all scales, using all this technology AND ensuring usability.

Acceleware**Title:** Designing and Optimizing Software for Accelerators**Speaker:** Dan Cyca, CTO

Bio: Regarded as a leading mind in the field of parallel processing, Dan has extensive experience working with GPUs, clusters and multi-core solutions. Dan joined Acceleware in 2004 as a software developer to build the company's first product. Since then, he has served in many technical and leadership roles in the company. Most recently, as the Director of Engineering, Dan was responsible for managing the software development group. Prior to Acceleware, Dan's experience included developing 'C-to-hardware' compilers, and implementing digital signal processing and encryption algorithms on FPGAs. Dan has an M. Sc. in Electrical Engineering from the University of Calgary.



Abstract: Established in 2004, Acceleware has been developing and maintaining commercial GPU applications for 7 years. Drawing from this experience we will present simple strategies for common problems. Discussion topics will include:

- Can you design your software to minimize the time required to support different hardware and different programming languages?
- How much effort is required to support new hardware? How does this impact your software design?
- Do your algorithms evolve over time with the improved performance or additional capabilities of new hardware?

CAPS**Title:** How to write code that will survive the manycore revolution?**Speaker:** François Bodin, CAPS CTO

Bio: As CTO, François Bodin leads research and development projects to help the creation of innovative software tools. François' research focuses on code optimization and compiler technologies for high performance computers and embedded systems. François is member of HIPEAC, the European Network of Excellence on High-Performance Embedded Architecture and Compilation. François has degrees in computer science from the university of Rennes I. François Bodin is also Chairman of IRISA (Institut de Recherche en Informatique et Systèmes Aléatoires), a research unit in the forefront of information and communication science and technology.

Abstract: Pushed by the pace of innovation in the GPU architecture and more generally the manycore technology, the processor landscape is moving fast. This fast evolution makes software development more complex. Furthermore, the impact of the programming style on future performance and portability of the application is difficult to forecast. The use of directives to annotate serial languages (e.g. C/C++/Fortran) looks very promising. They abstract the programming of low-level parallelism details while preserving code assets against the evolution of processor architectures. In this presentation, we describe how to use the HMPP (Heterogeneous Manycore Parallel Programming) API, one of the directive-based approaches, to program heterogeneous compute nodes. In particular, we provide insights on how GPU / CPU can be exploited in a unified manner and how code tuning issues can be minimized. We extend the

discussion to the use of libraries that is currently one of the key elements when addressing GPU and manycores.

Reservoir Labs

Title: Automatic Optimization and Code Generation for GPU and Many Core for Seismic Kernels

Speaker: Richard A. Lethin, Ph.D.

Bio: Richard A. Lethin, Ph.D. is President a Reservoir Labs, a small business with offices in Manhattan and Portland, OR. Reservoir performs research and development services and provides technologies for commercial and government customers in the area of high performance computing, signal processing, and cyber security. Dr. Lethin also teaches electrical engineering at Yale. He received his Ph.D. from MIT in 1997 for research building the MIT J-Machine.



Abstract: This talk will illustrate the capabilities of the R-Stream compiler for automatically parallelizing and optimizing C to CUDA and OpenMP, through the example of some stencil kernels used in the seismic imaging application area.

Tutorial B2a: Cluster Test

Hewlett-Packard

Title: Validating Cluster Functional and Performance Health with Cluster Test

Speaker: Zarka Cvetanovic, HP

Bio: Zarka has been involved in the design, architecture, and performance evaluation of a number of HP, Compaq, and DEC server products. She is currently leading efforts in software and tools development for interconnects in HP/HPC. Zarka holds a PhD degree in Computer Engineering from University of Massachusetts, Amherst.

Abstract: Cluster Test (CT) is the tool designed to verify functional and performance health of a cluster configuration. The system components validated include: processors, memory, disks, interconnects and GPGPUs. Cluster Test is developed and supported by HP and the RPM version can be accessed via a free software download.

Cluster Test can be applied to:

- Help resolve post-installation issues
- Help diagnose functional and performance issues while cluster has been in production

We will present Cluster Test applicability and several specific use cases.

We will demonstrate the existing and new features of Cluster Test, including the following:

- Easy-to-use GUI for launching/controlling/tracking the tests
- Hardware inventory of all components, including BIOS and firmware versions
- Performance analysis to identify nodes that do not perform as expected
- InfiniBand fabric check
- Testing GPGPUs: functionality and performance
- Performance monitor to identify bottlenecks
- Troubleshooting information for failed tests
- Server health check (information on temperature, power, and fans)
- Immediate failure reporting

We will discuss 2 new versions of Cluster Test: CTlite and CTimage, and how these versions can be obtained from HP. CTlite is a light-weight version of Cluster Test with the following features:

- A simple suite of tests derived from Cluster Test
- A simple tar-ball installation
- A simple CLI interface to allow running tests from a single command line
- Ability to install and run as a non-root user
- A single report with pass/fail results for each of the tests.

CTimage is the bare-metal version of Cluster Test that contains the OS and all software needed to setup and run Cluster Test.

Tutorial B2b: Shared/ Large Memory Systems**ScaleMP**

Title: Versatile SMP – from departmental Bio-Informatics to central HPC setups; Using HP servers and CMU for any SMP challenge

Speaker: Benzi Galili

Bio: Benzi joined ScaleMP in 2005, and is responsible for ScaleMP's operations and sales worldwide. Previously VP of Services for ScaleMP, Benzi brings more than 15 years of experience in managing technology development and services, product and account management, and in product architecture and design geared towards addressing business needs of enterprises and end-users. Prior to joining ScaleMP, Benzi served as VP of Products at Participate Systems (Acquired by OutStart Inc), VP R&D followed by VP Products and Services at Kamoon Inc, and also as CIO for the ERGO Consulting Group. Throughout his career he has compiled a successful track record, managing diverse engineering and product management teams, delivering key products and services to Global 2000 companies, governments, and leading academic and research institutes.



Benzi holds a B.Sc. in industrial engineering from Tel Aviv University (major in information systems), where he graduated with honors.

Cyfronet

Title: Experiences with ScaleMP at a Production Site

Speaker: Patryk Lason, Cyfronet

Bio: M.Sc. Eng. from Wroclaw University of Technology in 2003. Since 2004 system administrator at Cyfronet. Currently main system administrator of cluster "Zeus", the fastest supercomputer in Poland (#81 on Top500 in June 2011).

Abstract: ScaleMP's product as flexible SMP alternative. Integration with Zeus' infrastructure and first experiences with applications.

University of Indiana

Title: Experiences with a Large-memory HP Cluster - Performance on Benchmarks and Genome Codes

Speaker: Craig Stewart

Bio: Craig Stewart is the Executive Director of the Pervasive Technology Institute (PTI), IU's flagship initiative for advanced information technology research, development, and delivery in support of research, scholarship, and artistic performances. Stewart is Associate Dean for Research Technologies, and leads the Research Technologies Division of University Information Technology Services. Stewart has been very active in the high performance computing for many years, including being very active in the early days of HPC applications in biology. Stewart is a past chair of the Coalition for Academic Scientific Computation (2008-2009), during which time he testified about the importance of high performance computing before the House Science and Technology Committee. Most recently Stewart served as a member of the National Science Foundation Advisory Committee on CyberInfrastructure, and Chair of the ACCI Task Force on Campus Bridging (2010-2011). Stewart has had a long career in IT at Indiana University, with extensive experience in leading and managing services to support researchers at IU, including past appointments as Director of the Center for Statistical and Mathematical Computing, Research and Academic Computing, Indiana Genomics Initiative Information Technology Core; and Special Assistant for the Life Sciences, IU Office of the Vice President for Research. Stewart is an Adjunct Professor in the School of Informatics, and also holds adjunct appointments in the Department of Medical Genetics (IU School of Medicine) and Biology (IU Bloomington). Stewart has served as a Visiting Faculty Member in Computer Science, University of Stuttgart, and as a Fulbright Senior Scholar at the Technische Universitaet Dresden (Germany). He has a Ph.D. in Ecology and Evolutionary Biology from Indiana University and a Bachelor of Arts degree in Mathematics and Biology from Wittenburg University. Most of his publications are online in the IU Scholarworks Digital Repository (scholarworks.iu.edu).



Abstract: Indiana University has implemented a large-memory HP cluster called Mason as a key resource for the National Center for Genome Analysis Support. While this cluster seems small in terms of teraFLOPS it provides a critical resource that is unique within the US in being dedicated to genome analysis and in particular to the challenges of genome assembly based on next-generation sequencer data. In this talk, we will discuss: *Challenges of genome assembly from next-generation sequencers and how this drives the design of Mason *Some early performance analysis results from Mason *Early examples of science projects using Mason *How Mason fits into XSEDE and is an example of "campus bridging" in action

[Tutorial C2: Accelerators - Ecosystems](#)

[AccelerEyes](#)

Title: ArrayFire – Simple, Fast, and Free GPU Library

Speaker: John Melonakos

Bio: John is CEO of AccelerEyes. Under his leadership, the company has grown to deliver its products to thousands of customers in more than 40 countries around the globe. He has previous experience with Lockheed Martin, Hewlett-Packard, Corning Cable Systems, and General Electric. John holds a BS in Electrical Engineering from Brigham Young University and MS/PhD degrees in Electrical Engineering from the Georgia Institute of Technology.



Abstract: GPU libraries are the key to getting fast GPU performance without time-consuming low-level hassle. ArrayFire is a new software library with a very simple array-based API that can be learned in minutes. It is available for C, C++, Fortran, and Python and runs on AMD, Intel, and NVIDIA hardware. ArrayFire includes the largest and fastest set of GPU functions in the world. It is FREE to use for most people, with some options to purchase enterprise support. Jacket is the GPU engine for MATLAB, providing the only solution capable of accelerating full M-code scripts. It includes the popular GFOR loop as well as support for all NVIDIA CUDA and ATI OpenCL GPUs.

[Multipath](#)

Title: GPU Acceleration of Matrix Algebra

Speaker: Ronald C. Young

Bio: Dr. Young received his PhD from UC Berkeley. His career has been devoted to the hardware acceleration of matrix algebra, from early array processors to current systems supporting multiple GPUs.



Abstract: Multipath's Fast Matrix Solver (FMS) is the industry standard for processing large matrices at peak performance in production application. The GPU implementation is setting performance records in the teraflops by operating multiple GPU's in parallel and overlapping computation with data staging from disk to memory to GPU memory. Attendees will learn what problems FMS solves and how it is incorporated into applications. How FMS balances terabytes of data flow with teraflops of computing to achieve this performance is also discussed. A movie showing the actual performance of the CPUs and each GPU solving a large dense matrix is also included.

[Tidepowerd](#)

Title: Everyday Supercomputing

Speaker: Jack Pappas

Bio: Jack Pappas is co-founder and CEO of TidePowerd (tidepowerd.com), a company specializing in software tools for numerical computing. Jack's interests include functional programming, computational mathematics, and college football.

Abstract: Current HPC applications are heavily oriented towards industries which rely on numerical heavy-lifting for their day-to-day work. However, there are also many unexplored areas of computing which could provide for novel applications of HPC. This talk will cover software architecture and development tools for HPC and include demonstrations of novel applications built with GPU.NET.

[Wolfram Research, Inc.](#)

Title: Mathematica for Parallel and GPU Computation

Speaker: Ulises Cervantes-Pimentel

Bio: Ulises Cervantes-Pimentel is a senior kernel developer at Wolfram Research with more than 10 years of experience in developing and implementing algorithms and scientific visualization in Mathematica. He is currently leading the development of GPU computing in Mathematica as well as the implementation of state of the art computational geometry algorithms. He is a Master and PhD graduate in Applied Mathematics from the University of Illinois at Urbana-Champaign, and worked for several years as a Research Assistant at the NCSA and Beckman Institute.

Abstract: Mathematica is a computational tool used by the leading commercial and academic institutions. It has a built-in toolset for visualization, image processing, statistics, and finance providing a cohesive compute, develop, and deploy environment. In this talk we present Mathematica 8 and its Parallel and GPU capabilities. We will show how to use Mathematica to solve large problems, how to use it in teaching environments. We will then discuss how the symbolic and functional

nature of its language, as well its large library of built-in functions, makes it easy to not only get started, but to also to get speedups from your program easily.

Hewlett-Packard

Title: HP and GPU Affinity - Discussion

Speaker: Glenn Lupton, HP Accelerator Team Technical Lead

Bio: Glenn is the technical team leader for accelerators in HPC R&D. His background includes visualization projects and software development tools.

Abstract: For best application performance, the effects of memory and GPU bandwidth must be understood. This presentation will discuss how to use the server most effectively to get the most out of your system when bandwidth is critical.

Tutorial D2: Insight CMU

Hewlett-Packard

Title: CMU: Features and Functions Reviewed in Detail

Speaker: Chris Holmes, HP & Sébastien Cabaniols, HP

Bio: **Christopher Holmes** has degrees in Computer Engineering and Computer Science, and has been an HP HPC Linux cluster software engineer for 9 years. Prior to that he worked on software debugging tools, computer vision systems, and was an aircraft mechanic and flight-deck troubleshooter for the US Navy.

Sébastien Cabaniols has a MSc in Advanced Computer Science from the University of Manchester and also graduated from ENS Lyon option parallel computation in 2000. His area of expertise is the linux kernel, linux distributions and parallel computing in general. Sébastien joined the HP HPC team in 2001 working in many different areas from high speed networking to high performance storage and cluster management.

Abstract: This talk will present an overview and live demonstration of the existing features in HPs latest Cluster Management Utility (CMU) for Linux clusters. CMU provides scalable and efficient tools to provision, monitor, and manage a Linux cluster of any size (up to 4k nodes). CMU includes industry standard cluster tools such as pdsh, and unique differentiators such as its highly scalable image cloning process, its live graphic monitoring display, and its one-of-a-kind data filter, cmudiff, for identifying outliers in a cluster configuration. This talk will begin with an overview of CMU and then go into detail on many of the latest new features.

Accelerator SIG Session

Karlsruhe Institute of Technology (KIT), Engineering Mathematics and Computing Lab (EMCL) & Institute for Applied and Numerical Mathematics 4

Title: Library design and new parallel methods for accelerators

Speaker: Jan-Philipp Weiss, Head of Shared Research Group

Bio: Jan-Philipp Weiss studied mathematics at the universities of Stuttgart and Freiburg and worked as a researcher at the universities of Kaiserslautern and Karlsruhe. He received a Ph.D. from University Karlsruhe in 2006 and was appointed junior professor in 2008. He is leading a Shared Research Group at the Engineering Mathematics and Computing Labs (EMCL) at Karlsruhe Institute of Technology (KIT) in joint collaboration with the company Hewlett-Packard. Research of his group addresses parallel programming techniques and numerical methods for emerging multi- and manycore technologies in numerical simulation and scientific computing.



Abstract: Modern numerical simulation methods rely on both efficient parallel solution schemes and platform-optimized parallel implementations. The single code base approach of EMCL's Multi-platform Linear Algebra Toolbox is a flexible and efficient code framework that is portable across various computing platforms including x86 multicore CPUs, GPUs and OpenCL-enabled accelerators. In this talk we show how to adapt parallel numerical methods like preconditioners and smoothers to fine grained-parallelism by means of improved algorithms. Moreover, we detail the configuration of our parallel library and show performance results on various platforms.

KIT & Hewlett-Packard

Discussion: Multicore, Manycore and Accelerators - the User Perspective

Speaker: Jan-Philipp Weiss, KIT & Jim Bovay, HP

Abstract: The landscape of accelerators is still highly dynamic -- we have or will see new devices like NVIDIA's Kepler, AMD's Fusion or Intel's MIC architecture. The tutorial A1 has given an update on the latest products and roadmaps. In this session we would like to discuss the perspective of these technologies from the audience's point of view. What is your experience and expectation with respect to performance and productivity? How do you deal the multitude of different programming approaches in your ongoing/planned projects? Where do you see the future?

Visualization SIG Session

Nice Software

Title: **The Evolution of Remote 3D Visualization for Technical Computing**

Speaker: **Antonio Arena**

Bio: Antonio Arena has been working in HPC world for the last 10 years. Currently he's employed at Nice Software. The company relocated him to Houston to as business developer.

Abstract: 3D visualization is pivotal for innovation and efficiency in many fields, including Energy, Health Care, Engineering, Life Sciences, Research. A rich and fluid remote, interactive 3D visualization experience enables enterprises and institutions to make timely decisions that involve large data volumes that would be impractical to deliver to a user's desktop or workstation. In this talk we will review the present 3D technology landscape, their requirements and performance implications in real world scenarios. We'll also provide insight of multi-user collaboration and of the upcoming revolutionary synergy between GPUs and Virtualization technologies that allows cost optimizations and GPU sharing across multiple users and effectively opens up the on-demand 3D visualization in the Cloud era.



Axceleon

Title: **Opportunities and Challenges for Image Processing on the Cloud**

Speaker: **Mike Duffy, President and Chief Operating Officer**

Bio: Mr. Duffy brings more than 30 years of global executive experience to Axceleon with a special focus on engineering and operations. As president and COO of Axceleon, Mr. Duffy runs all aspects of the company's global software business. Mr. Duffy has spent the last 10 years of his career working the in the High Performance Computing (HPC) arena.

Prior to Axceleon Mr. Duffy served as senior vice president of sales and global enterprise services for TurboLinux, a Linux company, where Mr. Duffy ran the US, Latin America and European sales and professional services teams. With his strong revenue focus he doubled the company's revenues on a quarter by quarter basis and closed large deals with IBM, HP, Compaq and Intel. Mr. Duffy has held positions with Solectron Inc. as vice president of international business development and Sun Microsystems as Director of Enterprise Services for Latin America and Asia. Mr. Duffy holds degrees in electrical engineering and financial management from University College Dublin & College of Technology, UCLA and UC Berkeley.

Abstract: Challenges and Opportunities for Image processing on the Cloud (The Democratization of Digital Content Creation).

The democratization of digital content creation may be getting closer especially with the onset of image rendering services in the cloud. Cloud rendering is expanding from just providing raw horsepower to it providing services and features that being connected to the cloud can deliver. With the advent of on-demand pay as you go rendering services the massive power of HPC, using hundreds of cores for image rendering, can be accessed easily and cheaply by the smaller studios and the individual artist. However, along with these opportunities come the challenges. Data Transfer times to and from these "render clouds" are directly related to file sizes and the available bandwidth. What does the workflow look like from the Artist's perspective; do they have to become "cloud technology experts"? How is the cloud started up, shut down and managed in general. There are ways to mitigate and minimize these problems that could impact adoption and usage of these cloud based services for rendering and we will explore these in more detail and what services are out there today.



NVIDIA

Title: Remote Visualization on Quadro and Tesla M2070Q
Speaker: Bob Crovella
Bio: HPC Solutions Engineer supporting NVIDIA Tesla GPU computing products
Abstract: "Remote Visualization on Quadro and Tesla M2070Q" - A discussion of remote/virtualization technologies including RemoteFX, VirtualGL, VMware, Citrix, RGS, and others, which make use of the NVIDIA Tesla M2070Q and Quadro products for the visualization engine.

Oxalva

Title: Boosting HPC Project Reviews
Speaker: Alban Schmutz, founder and CEO of Oxalva
Bio: Alban Schmutz is founder and CEO of Oxalva, a software techno-provider dedicated to HPC and scientific visualization. Started in 2005, the company has a wide range of customers in Energy, Manufacturing, Medical and Electronic industries with prestigious customers like CEA, EDF, Thales, Areva, Arcelor-Mittal, Renault, Plastic Omnium, Valeo, and many others. Since 2008, he is elected to the board of the Systematic competitiveness cluster, which involves 500 research organizations in the Paris Region Area (big companies, research labs and SMEs). Deeply involved in R&D partnerships in the field of numerical simulation, he has coordinated several R&D platform projects up to 28 partners co-funded by the French National Research Agency or the French Ministry of Finance. Research areas are focused on HPC datacenter automation with VirtualNodes and Scientific Visualization with VisuPortal. Previously, he created successively in 1998 and 2000, two open source services companies. The second one became the leading open source company in France. He wrote many articles on Open Source business models. Alban Schmutz is 32 yo, holds three master degrees in Political Sciences, Telecommunication Management and Entrepreneurship Engineering. He is namely a graduate of the French National Institute for Telecommunications and Institut d'Etudes Politiques de Paris.



Abstract: "Powerwalls, CADwalls, CAVEs and/or standard desktops are used to perform Project Reviews for design. The presentation will show how high end management tools can boost Project Reviews through high performance visualization systems automation, large data management and remote collaboration."

Cloud & Grid SIG SessionUniversity of Oregon

Title: The University of Oregon ACISS Science Cloud
Speaker: Allen Malony
Abstract: An interdisciplinary group of computer scientists, psychologists, biologists, chemists, neuroscientists, and physicists at the University of Oregon (UO) is creating a large-scale computational and storage resource to support next-generation scientific research. Funded by an award from the NSF Major Research Instrumentation program and UO's commitment to research computing, the "Applied Computational Instrument for Scientific Synthesis" (ACISS) emphasizes both the productive shared use of the instrument and the potential synergies it will afford for integrative scientific research. The ACISS hardware will consist of general purpose multicore computing nodes, high-performance computing nodes augmented with GPGPU acceleration, a high-volume storage system, and high-bandwidth networking infrastructure, all to be housed in an updated UO Computing Center. A unique feature of ACISS will be its management as a cloud system for computational science, informatics, and data science.

Energy, Cooling and Infrastructure SIG SessionConsultant

Title: Energy Issues of a Large Center in Iceland
Speaker: Thorvaldur Sigurdsson
Bio: International IT experience including 23 years in the US and Europe. Former CTO and Managing Director of IT operations at Skyrr which hosts critical IT infrastructure for the Icelandic government, municipalities and several companies in Iceland. Successful entrepreneur and new business development participant of several key projects and companies. Electrical Engineer graduate from the University of Iceland.
Abstract: Iceland hosts the new large scale Nordic-HPC environment based on HP blade solutions. Why? What are the advantages and limitation of Iceland as a future data center hub?

CMU SIG Session

Discussion: New features. Best practices for CMU. Customer requirements.

Audience: HP, customers and resellers. Others by Invitation ONLY.

Abstract: This new SIG is being added to provide customers and our HP team an opportunity to share best practices for CMU deployment and usage, and discuss new and upcoming features, integration with other HPC tools, and how to leverage the online CMU Forum to tap expertise from the CMU community. This is an opportunity for users to discuss issues and requirements for cluster management.

Language and Programming SIG Session

Tokyo Institute of Technology

Title: A 2-Petaflops Stencil Application on SL390/ GPU-based TSUBAME 2.0

Speaker: Takayuki Aoki

Bio: Takayuki Aoki received a BSc in Applied Physics (1983), an MSc in Energy Science and Dr.Sci (1989) from Tokyo Institute of Technology, was a Visiting Fellow in Cornell University and the Max-Planck Institute in Germany for one year, has been a professor in Tokyo Institute of Technology since 2001 and a deputy director of the Global Scientific Information and Computing Center since 2009. He has received the Computational Mechanics Achievement Award from Japan Society of Mechanical Engineers and many awards and honors in visualization, and others. He has authored the first book in the Japanese language on the CUDA programming and applications.

Abstract: Several large-scale stencil applications have been successfully developed on GPU-rich supercomputer TSUBAME 2.0, which consists of 1420-SL390 nodes and is equipped with 4224 NVIDIA Tesla M2050 GPUs, and has started the operation since November 2010 at Tokyo Tech. Stencil computing on a regular structured grid is suitable for GPU computing since high performance of the memory access can be achieved for the on-board memory. In an explicit time integration, we introduce a technique overlapping the GPU-to-GPU communication with the computation into large-scale applications for the purpose to hide the communication overhead. A phase-field simulation runs for a dendritic solidification of the Al-Si binary alloy. In our largest configuration of $4096 \times 6500 \times 10400$, TSUBAME 2.0 achieved 2.0 Petaflops in single precision, using 4,000 GPUs along with 16,000 CPU cores. The sustained performance has reached 45 % of the peak performance. We also demonstrate gas-liquid two-phase flows and results of Lattice Boltzmann Method.

Workshop for Programmers 1: OpenCL

Advanced Micro Devices (AMD)

Title: OpenCL: A Tutorial for Programmers

Speaker: Udeпта Bordoloi

Bio: Udeпта Bordoloi is the lead engineer for HPC/GPGPU applications at AMD. He obtained his MS degree in Electrical Engineering from Washington University in St. Louis and his PhD in Computer Science from The Ohio State University. Since then he has worked in the field of graphics, visualization and cloud-based rendering with industry leading companies and also with bootstrapping startups. He has authored academic publications and patents, tutorials and lectures, and continues research and development on the state-of-the-art in GPUs. Besides GPGPU and HPC, he has interests in Image Processing, Graphics and Visualization.

Abstract: We will introduce the OpenCL language to programmers using introductory examples. The goal is that after attending the tutorial the audience will be able to create a simple OpenCL program and run it on both a CPU and a GPU. The talk will touch upon fundamental OpenCL concepts of memory spaces, queues, synchronization issues etc. We will also go over the mapping of OpenCL concepts to the underlying hardware both for a CPU and a GPU.

[Workshop for Programmers 2: CUDA](#)[NVIDIA](#)**Title:****CUDA: A Tutorial for Programmers – All you need to start accelerating applications****Speaker:****Tom Reed, HPC & Visualization Solution Architect****Bio:**

Tom Reed currently works with NVIDIA's key customers and OEMs to optimize the benefits of GPU technology. Prior to joining NVIDIA, Mr. Reed spent 20 years with Silicon Graphics where his tenure included roles in systems engineering, benchmarking & performance engineering, software development, and professional services. Mr. Reed's career has always involved some combination of high performance computing (HPC) and visualization and he's long been an advocate for the fusion of best-of-breed technologies to drive sustainable and efficient advances in IT capability. Past research work has included initiatives in interactive parallel computing, data-centric systems design, and hybrid computing architectures. In addition to SGI, Mr. Reed has worked for McDonnell Douglas, General Dynamics, and Linuxcare as well as research contracts with Kirkland Air force Base, Arnold Engineering Development Center (AEDC), and NASA. Mr. Reed attended Iowa State University for both undergraduate and graduate studies, majoring in aerospace engineering with graduate research in computation fluid dynamics (CFD).