



Austin, TX | hpc transforms.

Austin
Convention Center
Austin, TX

<http://sc15.supercomputing.org/>

The International Conference for High Performance
Computing, Networking, Storage and Analysis

Conference Program

Conference Dates: Nov. 15 - 20, 2015 Exhibition Dates: Nov. 16 - 19, 2015

Sponsors:



Association for
Computing Machinery

IEEE
computer
society



The International Conference for High Performance Computing, Networking, Storage and Analysis



Sponsors:



Association for
Computing Machinery



Table of Contents

Welcome from the Chair	4	Papers.....	68
General Information	5	Posters	
Registration and Conference Store Hours	5	Research Posters.....	88
Exhibit Hall Hours.....	5	ACM Student Research Competition	114
SC15 Information Booth/Hours	5	Posters	
SC16 Preview Booth/Hours	5	Scientific Visualization/	120
Social Events.....	5	Data Analytics Showcase	
Registration Pass Access	7	Student Programs	
SCinet	8	Experiencing HPC for Undergraduates	122
Convention Center Maps	12	Mentor-Protégé Program	123
Daily Schedules	16	Student Cluster Competition Kickoff	123
Plenary/Keynote/Invited Talks		Student-Postdoc Job &	123
Plenary.....	26	Opportunity Fair	
Keynote.....	26	Tutorials.....	124
Invited Talks.....	27	Workshops	137
Awards Presentations	33	Acknowledgements	151
Birds of a Feather	36		
Doctoral Showcase.....	50		
Emerging Technologies	54		
Exhibitor Forums/	57		
HPC Impact Showcase			
Panels	63		

Welcome to SC15



Welcome to the fantastic city of Austin and to SC15, the International Conference for High Performance Computing, Networking, Storage and Analysis. With its vast technical community and the lively culture, Austin is a great place to host the 27th annual SC Conference.

HPC is transforming our everyday lives as well as our not-so-ordinary ones—from nanomaterials to jet aircrafts, from medical treatments to disaster preparedness, and even the way we wash our clothes—the HPC community has transformed the world in multifaceted ways.

The technical program is going to be fantastic this year. We have a great Keynote lined up for you with Alan Alda speaking, as well as a variety of options including invited talks, papers, panel, and so much more. We are also launching a new initiative this year, the Early Career Program, which provides an opportunity for special sessions of interest to early career researchers, including getting funding, publishing venues, establishing long-term mentor relationships, and time management. The program is intended for people in their first five years of a permanent position (assistant professors and technical staff).

Don't forget to check out the HPC Matters Plenary with Diane Bryant from Intel. She will discuss how next-generation supercomputers are transforming HPC and presenting exciting opportunities to advance scientific research and discovery to deliver far-reaching impacts on society. After this session we will be opening the show floor. Come visit the over 300 exhibitors that are here to share with you the latest in our industry. They range from industry to academia and research and have something for everyone.

Make sure you also visit the SCinet booth to view the hardware that makes our network the best and fastest in the world. This is the 25th anniversary for SCinet as a component of the SC Conference series. This team is supporting the wireless network as well as the research network and the commodity network. Each year the SCinet NRE showcases a number of interesting network-based experiments during SC. This year they also took an initiative to bring more women into the SCinet community. Another positive change this year is that we have pulled all things related to students under one umbrella. Students@SC is focused on providing the next generation of HPC Community members with resources to help them grow. Amongst other things, they will be participating in the Mentor-Protégé program, a job fair, an educational dinner, the Student Cluster Competition and volunteering in critical positions throughout the week. Plus, they will have many other opportunities to network with interesting people.

There is literally too much for me to mention here, but if you need direction or just have a few questions, please stop by one of our two information booths.

Welcome to SC15 and to Austin, and I will see you in the Convention Center!

A handwritten signature in black ink, reading "Jackie Kern". The signature is fluid and cursive.

Jackie Kern

SC General Chair

University of Illinois at Urbana-Champaign

General Information

Registration and Conference Store

The registration area and conference store are located in the Atrium of the Austin convention center.

Registration/Conference Store Hours

Saturday, November 15	1pm-6pm
Sunday, November 16	7am-6pm
Monday, November 17	7am-9pm
Tuesday, November 18	7:30am-6pm
Wednesday, November 19	7:30am-6pm
Thursday, November 20	7:30am-5pm
Friday, November 21	8am-11am

Registration Pass Access

See page 7 for access grid.

Exhibit Hall Hours

Tuesday, November 18	10am-6pm
Wednesday, November 19	10am-6pm
Thursday, November 20	10am-3pm

SC15 Information Booths

Need up-to-the-minute information about what's happening at the conference. Need to know where to find your next session? What restaurants are close by? Where to get a document printed? These questions and more can be answered by a quick stop at one of the SC Information booths. There are two booth locations for your convenience: the main booth is located in the Fourth Street Concourse; the second booth is located at the Trinity Street Concourse.

Information Booth hours are as follows (note parentheses indicate hours for secondary booth; closed if no time indicated):

Saturday, November 14	1pm-5pm
Sunday, November 15	8am-6pm (8am-5pm)
Monday November 16	8am-7pm (8am-5pm)
Tuesday November 17	8am-6pm (10am-5pm)
Wednesday November 18	8am-6pm (8am-4pm)
Thursday November 19	8am-6pm
Friday November 20	8:30am-12:30pm

SC16 Preview Booth

Members of next year's SC committee will be available in the SC16 preview booth (L12, located near the Fourth Street Entrance, outside Exhibit Hall 4) to offer information and discuss next year's SC conference in Salt Lake City, Utah. Stop by for a copy of next year's Call for Participation and pick up some free gifts!

The booth will be open during the following hours:

Tuesday, Nov. 17	10am-6pm
Wednesday, Nov. 18	10am-6pm
Thursday, Nov. 19	10am-4pm

Social Events

Exhibitors Reception

Sunday, November 16

6pm-9pm

Austin City Limits

The SC15 Exhibitors Reception will be held in the Moody Theater at Austin City Limits. The reception is open to all registered exhibitors. The reception is SC15's way of thanking exhibitors for their participation and support of the conference. The reception will feature a popular local band, The Spazmatics, access to all the memorabilia from 40 years of recording the Austin City Limits program at the theater, and a showcase of local food and drinks throughout the facility. Transportation will be provided from the bus loading area of the convention center starting at 5:45pm. An Exhibitor badge, party ticket and government-issued photo ID are required to attend this event.

Exhibits Gala Opening Reception

Monday, November 17

7pm-9pm

SC15 will host its annual Grand Opening Gala in the Exhibit Hall. This will be your first opportunity to see the latest high performance computing, networking, storage, analysis, and research products, services, and innovations. This event is open to all Technical Program, Exhibitors and Students@SC registrants.

Posters Reception**Tuesday, November 18****5:15pm-7pm****Level 4 - Concourse**

The Posters Reception will be located in the Concourse on Level 4 of the convention center. The reception is an opportunity for attendees to interact with poster presenters and includes research and ACM Student Research Competition posters. The reception is open to all attendees with Technical Program registration. Complimentary refreshments and appetizers will be available.

Technical Program Conference Reception**Thursday, November 19****6pm-9pm****The Darrell K Royal-Texas Memorial Stadium**

Hook 'em Horns! SC15 will host a conference reception for all Technical Program attendees. Join us for a sporting good time at the Darrell K Royal-Texas Memorial Stadium, home to the UT Austin Longhorns football team. The stadium is the second largest stadium in the state of Texas, the ninth largest stadium in the United States, and the twelfth largest stadium in the world. Enjoy a warm Texas welcome from the UT Spirit Team and Bevo, the school mascot, a tour of the facility, delicious food and drink, music and much more. A Tech Program badge, event ticket, and government-issued photo ID are required to attend this event. Attendees are required to wear technical program badges throughout the reception, and badges may be checked during the event. Shuttle transportation to and from the event will run starting at 5:30pm from the convention center.

Family Day

Family Day is Wednesday, November 19, 4pm-6pm. Adults and children 12 and over are permitted on the floor during these hours when accompanied by a registered conference attendee.

Lost Badge

There is a \$40 processing fee to replace a lost badge.

Facilities**Coat and Bag Check**

There are self-service locations for hanging coats within the Technical Program rooms and in the lobby. In addition, there is a coat and bag check on the premises.

Parking

There are two parking facilities at the Austin Convention Center:

- 10-story, 1,000-space garage located two blocks west of the facility, at Brazos and 210 East 2nd Street. Entrances on Brazos and San Jacinto streets.
- 5-story, 685-space garage located at the northeast corner of the facility at Red River and 4th Street. Entrances on Fifth Street.

First Aid/Emergency Medical Team

The Austin Convention Center will provide an onsite first aid room staffed with an emergency medical professional during the SC15 conference. The first aid room is located on the first level, north of the Trinity North Elevator. In the event of a medical or other emergency, attendees can dial 911 from any pay phone or dial 4111 from any house phone located in the facility. In addition, all uniformed security personnel are available to assist you in any emergency.

Wheelchair Accessibility

The Austin Convention Center complies with ADA requirements and is wheelchair accessible. Complimentary wheelchairs are available only for emergencies.

Registration Pass Access

Each registration category provides access to a different set of conference activities, as summarized below.

Type of Event	Technical Program	Technical Program + Workshops	Tutorials	Workshop Only	Exhibitor 24-hour Access	Exhibit Hall Only Tue-Thur
Awards (Thursday)	*	*		*	*	*
Birds-of-a-Feather	*	*		*	*	*
Conference Reception (Thursday)	*	*				
Doctoral Showcase	*	*				
Emerging Technologies	*	*	*	*	*	*
Exhibit Floor	*	*			*	*
Exhibitor Forum	*	*			*	*
Exhibits Gala Opening (Monday)	*	*			*	
Exhibitors Reception					*	
HPC Impact Showcase	*	*	*	*	*	*
HPC Matters Plenary (NO BADGE REQUIRED/OPEN TO PUBLIC)	*	*	*	*	*	*
Invited Talks (Non-Plenary)	*	*				
Invited Talks (Plenary)	*	*			*	
Keynote (Tuesday)	*	*	*	*	*	*
Panels (Tue-Thur)	*	*				
Panels (Friday Only)	*	*			*	*
Papers	*	*				
Posters	*	*			*	*
Poster Reception (Tuesday)	*	*				
Scientific Visualization Showcase	*	*				
Tutorial Lunch (Sun/Mon ONLY)			*			
Tutorial Sessions			*			
Students@SC	*	*	*		*	*
Student Cluster Competition	*	*			*	
Workshops		*		*		



SCinet

The Fastest Network Connecting the Fastest Computers

SC15 is once again hosting one of the most powerful and advanced networks in the world - SCinet. Created each year for the conference, SCinet brings to life a very high-capacity network that supports the revolutionary applications and experiments that are a hallmark of the SC conference. SCinet will link the convention center to research and commercial networks around the world. In doing so, SCinet serves as the platform for exhibitors to demonstrate the advanced computing resources of their home institutions and elsewhere by supporting a wide variety of bandwidth-driven applications including supercomputing and cloud computing.

Volunteers from academia, government and industry work together to design and deliver the SCinet infrastructure. Industry vendors and carriers donate millions of dollars in equipment and services needed to build the local and wide area networks. Planning begins more than a year in advance of each SC conference and culminates in a high-intensity installation in the days leading up to the conference.

For SC15, SCinet is exploring up-and-coming topics in the high-performance networking community of SC through the Network Research Exhibition (NRE) and returning from its inaugural appearance in 2014, the INDIS workshop.

In addition, SCinet has created a new Diversity Program, funded by the NSF, to fund U.S. women in their early to mid careers to participate in SCinet to get hands-on training and build their professional network.

SCinet Network Research Exhibition (NRE)

The NRE is SCinet's forum for displaying new or innovative demonstrations in network testbeds, emerging network hardware, protocols, and advanced network intensive scientific applications that will stretch the boundaries of a high-performance network.

SCinet INDIS Workshop

The second annual INDIS workshop is again part of the Technical Program. **Innovating the Network for Data-Intensive Science** (INDIS) showcases both demonstrations and technical papers highlighting important developments in high-performance networking. With participants from research, industry, and government, this workshop will discuss topics related to network testbeds, emerging network hardware, protocols, and advanced network-intensive scientific applications.

SCinet Collaborators

SCinet is the result of the hard work and significant contributions of many government, research, education and corporate collaborators. Collaborators for SC15 include:

Platinum



CenturyLink™



Gold




SCinet Collaborators

Silver

Alcatel•Lucent 

ARISTA



ixia

Reservoir Labs



SCinet Collaborators

Bronze

CABLEexpress®



COMMSCOPE®

splunk® >

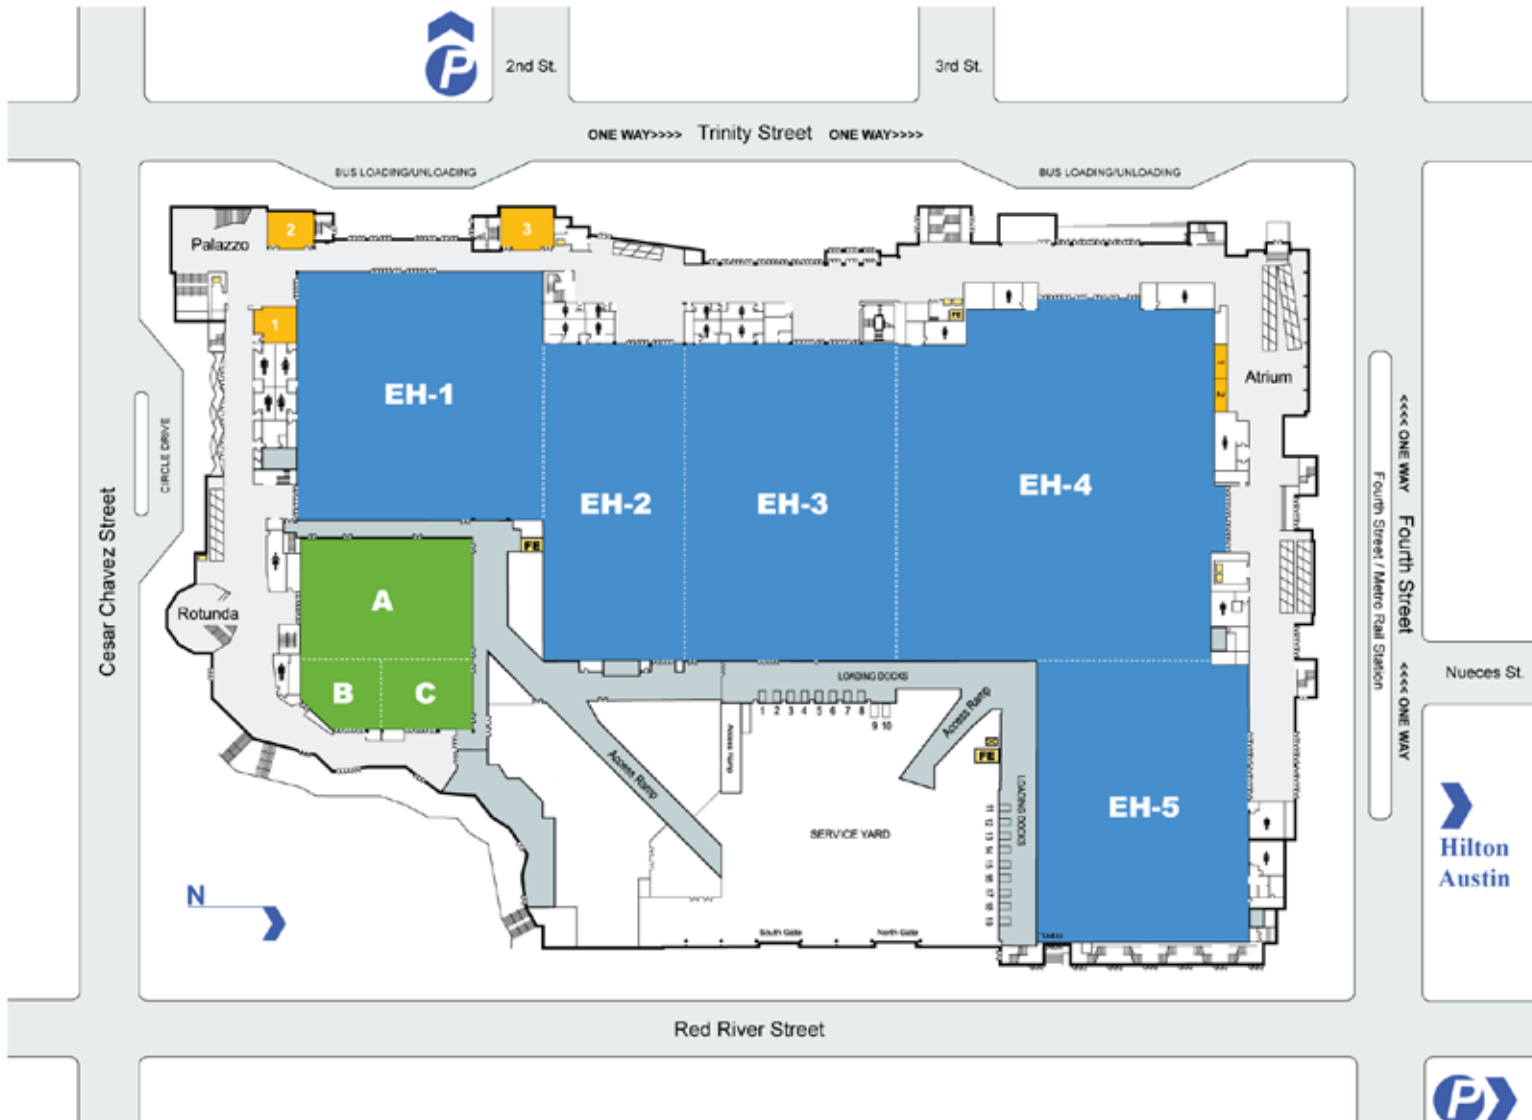


VERISIGN®

VIAVI

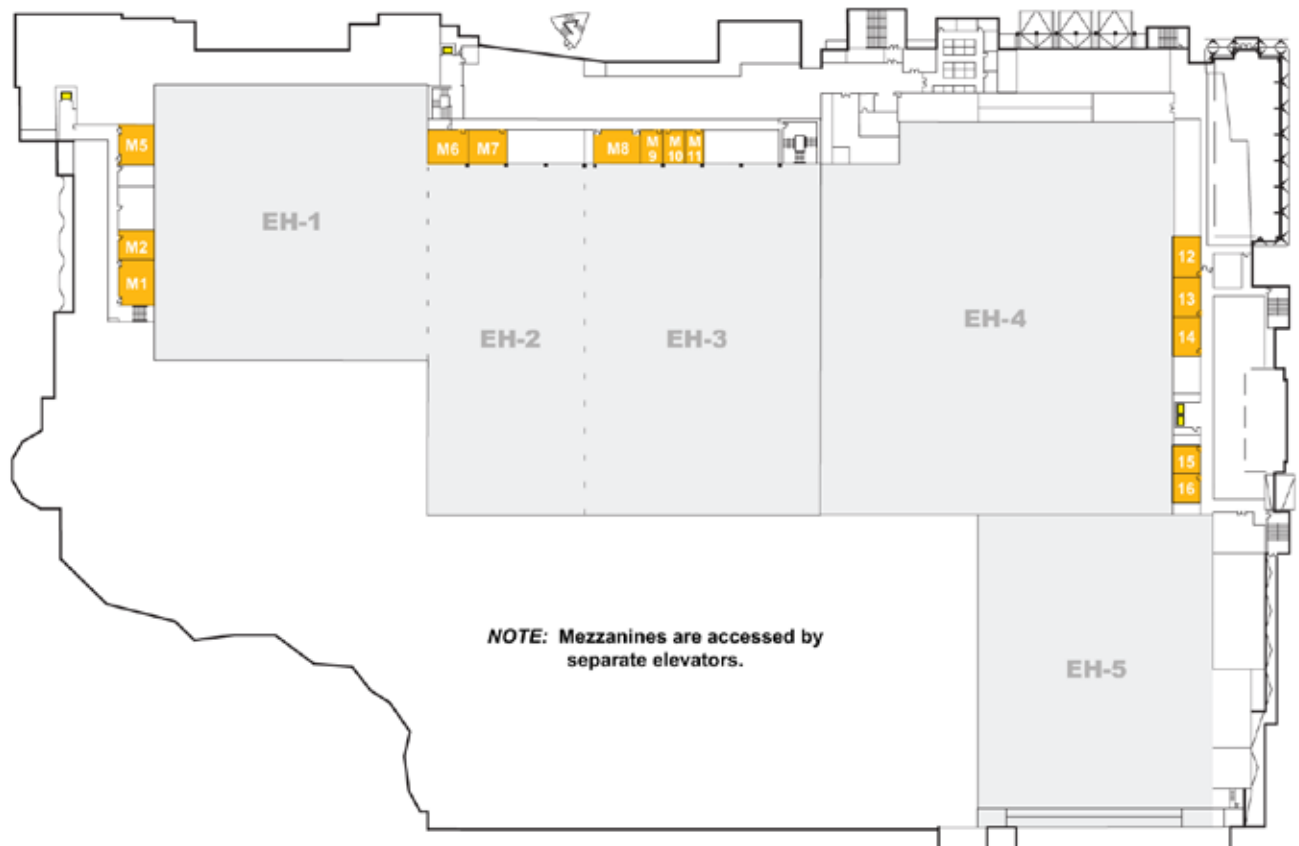
1

MULTI-LEVEL FLOOR PLAN DETAIL



2

MEZZANINES AND SHOW OFFICES



3

LEVEL 3 MEETING ROOMS



4

LEVEL 4 GRAND BALLROOM, MEETING ROOMS



SATURDAY, NOVEMBER 14

Time	Event Type	Session Title	LocationTime
5pm-6pm	Students@SC	Program Orientation	9BC

SUNDAY, NOVEMBER 15

	Time	Event Type	Session Title	LocationTime
	8:30am-10am	Students@SC	Why SC, HPC and Computer Science Need Diversity	9BC
	8:30am-12pm	Tutorial	Architecting, Implementing and Supporting Multi-Level Security Ecosystem in HPC, ISR, Big Data Analysis and Other Environments	18A
	8:30am-12pm	Tutorial	Benchmarking Platforms for Large-Scale Graph Processing and RDF Data Management: The LDBC Approach	19B
	8:30am-12pm	Tutorial	Data-Intensive Applications on HPC Using Hadoop, Spark and RADICAL Cybertools	18C
	8:30am-12pm	Tutorial	Large-Scale Visualization with ParaView	13AB
	8:30am-12pm	Tutorial	MPI+X—Hybrid Programming on Modern Compute Clusters with Multicore Processors and Accelerators	16AB
	8:30am-12pm	Tutorial	Automatic Performance and Energy Tuning with the Periscope Tuning Framework	19A
	8:30am-5pm	Tutorial	Introduction to OpenMP	17B
	8:30am-5pm	Tutorial	Efficient Parallel Debugging for MPI, Threads and Beyond	14
	8:30am-5pm	Tutorial	Fault-Tolerance for HPC: Theory and Practice	18D
	8:30am-5pm	Tutorial	Linear Algebra Libraries for HPC: Scientific Computing with Multicore and Accelerators	15
	8:30am-5pm	Tutorial	Massively Parallel Task-Based Programming with HPX	17A
	8:30am-5pm	Tutorial	Parallel Computing 101	18B
	8:30am-5pm	Tutorial	Parallel I/O in Practice	12A
	8:30am-5pm	Tutorial	Parallel Programming in Modern Fortran	12B
	9am-12:30pm	Workshop	Computational and Data Challenges in Genomic Sequencing	Hilton 406
	9am-12:30pm	Workshop	First International Workshop on Heterogeneous Computing with Reconfigurable Logic	Hilton Salon D
	9am-12:30pm	Workshop	Portability Among HPC Architectures for Scientific Applications	Hilton Salon G
	9am-5:30pm	Workshop	DISCS2015: International Workshop on Data Intensive Scalable Computing Systems	Hilton Salon J
	9am-5:30pm	Workshop	E2SC2015: Energy Efficient Super Computing	Hilton Salon F
	9am-5:30pm	Workshop	ESPM2: First International Workshop on Extreme Scale Programming Models&Middleware	Hilton 410
	9am-5:30pm	Workshop	IA ³ 2015: Fifth Workshop on Irregular Applications: Architecture and Algorithms	Hilton Salon A
	9am-5:30pm	Workshop	LLVM-HPC2015: Second Workshop on the LLVM Compiler Infrastructure in HPC	Hilton Salon E

SUNDAY, NOVEMBER 15

	Time	Event Type	Session Title	LocationTime
	9am-5:30pm	Workshop	PMBS2015: Sixth International Workshop on Performance Modeling, Benchmarking, and Simulation of HPC Systems	Hilton Salon K
	9am-5:30pm	Workshop	PyHPC2015: Fifth Workshop on Python for High-Performance and Scientific Computing	Hilton Salon C
	9am-5:30pm	Workshop	The Sixth International Workshop on Data-Intensive Computing in the Clouds	Hilton 412
	9am-5:30pm	Workshop	VISTech Workshop 2015: Visualization Infrastructure and Systems Technology	Hilton Salon B
	9am-5:30pm	Workshop	WORKS2015: Tenth Workshop on Workflows in Support of Large-Scale Science	Hilton 408
	10am-10:30am	Tutorial	Refreshment Break	Level 4 - Lobby
	10am-10:30am	Workshop	Refreshment Break	Hilton 4th Floor
	10:30am-12pm	Students@SC	Diversity Panel: Collective Responsibilities for Biases, Micro-aggressions, Isolation and More	9BC
	12pm-1:30pm	Tutorial	Lunch	Ballroom DEFG
	1:30pm-3pm	Students@SC	Education Panel: Making the Most of Graduate School (in HPC)	9BC
	1:30pm-5pm	Tutorial	Effective HPC Visualization and Data Analysis using VisIt	18C
	1:30pm-5pm	Tutorial	Insightful Automatic Performance Modeling	18A
	1:30pm-5pm	Tutorial	Live Programming: Bringing the HPC Development Workflow to Life	16AB
	1:30pm-5pm	Tutorial	Power Aware HPC: Challenges and Opportunities for Application Developers	19A
	1:30pm-5pm	Tutorial	Productive Programming in Chapel: A Computation-Driven Introduction	19B
	1:30pm-5pm	Tutorial	Towards Comprehensive System Comparison: Using the SPEC HPG Benchmarks HPC Systems for Better Analysis, Evaluation, and Procurement of Next-Generation	13AB
	2:00pm-5:30pm	Workshop	Computational Approaches for Cancer	Hilton 406
	2:00pm-5:30pm	Workshop	Many-Task Computing on Clouds, Grids and Supercomputers	Hilton Salon G
	2:00pm-5:30pm	Workshop	MLHPC2015: Machine Learning in HPC Environments	Hilton Salon D
	3pm-3:30pm	Tutorial	Refreshment Break	Level 4 - Lobby
	3pm-3:30pm	Workshop	Refreshment Break	Hilton 4th Floor
	3:30pm-5pm	Students@SC	Career Panel: Outlooks and Opportunities from Academia, Industry and Research Labs	9BC
	6pm-8pm	Students@SC	Dinner	Ballroom G
	6pm-9pm	Social Event	Exhibitor Reception (Exhibitor badge required for access)	Austin City Limits

MONDAY, NOVEMBER 16

	Time	Event Type	Session Title	LocationTime
	8:30am-10am	Students@SC	Research Panel: A Best Practices Guide to (HPC) Research	9BC
	8:30am-10am	Tutorial	From Description to Code Generation: Building High-Performance Tools in Python	17A
	8:30am-12pm	Tutorial	Getting Started with <i>In Situ</i> Analysis and Visualization Using ParaView Catalys	12B
	8:30am-12pm	Tutorial	InfiniBand and High-Speed Ethernet for Dummies	16AB
	8:30am-12pm	Tutorial	Managing Data Throughout the Research Lifecycle Using Globus Software-as-a-Service	15
	8:30am-12pm	Tutorial	MCDRAM (High Bandwidth Memory) on Knights Landing—Analysis Methods/Tools	18A
	8:30am-12pm	Tutorial	Practical Fault Tolerance on Today's HPC Systems	18D
	8:30am-12pm	Tutorial	Practical Hybrid Parallel Application Performance Engineering	14
	8:30am-5pm	Tutorial	Advanced MPI Programming	18B
	8:30am-5pm	Tutorial	Advanced OpenMP: Performance and 4.1 Features	13AB
	8:30am-5pm	Tutorial	How to Analyze the Performance of Parallel Codes 101	19A
	8:30am-5pm	Tutorial	Node-Level Performance Engineering	18C
	8:30am-5pm	Tutorial	OpenACC Programming For Accelerators	12A
	8:30am-5pm	Tutorial	Portable Programs for Heterogeneous Computing: A Hands-On Introduction	17B
	8:30am-5pm	Tutorial	Programming the Xeon Phi	19B
	9am-12:30pm	Workshop	INDIS-15: The Second Workshop on Innovating the Network for Data Intensive Science	Hilton 410
	9am-12:30pm	Workshop	Second SC Workshop on Best Practices for HPC Training	Hilton Salon K
	9am-12:30pm	Workshop	Ultravis '15: The Tenth Workshop on Ultrascale Visualization	Hilton Salon D
	9am-5:30pm	Workshop	ATIP Workshop on Chinese HPC Research Toward New Platforms and Real Applications	Hilton 412
	9am-5:30pm	Workshop	Co-HPC2015: Second International Workshop on Hardware-Software Co-Design for HPC	Hilton Salon C
	9am-5:30pm	Workshop	ESPT2015: Extreme-Scale Programming Tools	Hilton Salon B
	9am-5:30pm	Workshop	ExaMPI15: Workshop on Exascale MPI	Hilton Salon F
	9am-5:30pm	Workshop	PDSW10 2015: Tenth Workshop on Parallel Data Storage	Hilton Salon G
	9am-5:30pm	Workshop	Runtime Systems for Extreme Scale Programming Models & Architectures (RESPA)	Hilton 400-402
	9am-5:30pm	Workshop	SCaIA15: Sixth Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems	Hilton Salon E
	9am-5:30pm	Workshop	Sixth Annual Workshop for the Energy Efficient HPC Working Group (EE HPC WG)	Hilton Salon A
	9am-5:30pm	Workshop	Sixth SC Workshop on Big Data Analytics: Challenges and Opportunities (BDAC-15)	Hilton Salon J

MONDAY, NOVEMBER 16

	Time	Event Type	Session Title	LocationTime
	9am-5:30pm	Workshop	WACCPD: Workshop on Accelerator Programming Using Directives	Hilton 406
	9am-5:30pm	Workshop	WOLFHPC15: Fifth International Workshop on Domain-Specific Languages and High-Level Frameworks for HPC	Hilton 408
	10am-10:30am	Tutorial	Refreshment Break	Level 4 -Lobby
	10am-10:30am	Workshop	Refreshment Break	Hilton 4th Floor
	10:30am-12pm	Students@SC	What's at SC?	9BC
	12pm-1:30pm	Tutorial	Lunch	Ballroom DEFG
	1:30pm-3pm	Students@SC	Peer Speed Meetings	9ABC
	1:30pm-5pm	Tutorial	Accelerating Big Data Processing with Hadoop, Spark and Memcached on Modern Clusters	16AB
	1:30pm-5pm	Tutorial	An Introduction to the OpenFabrics Interface API	12B
	1:30pm-5pm	Tutorial	Data Management, Analysis and Visualization Tools for Data-Intensive Science	17A
	1:30pm-5pm	Tutorial	Debugging and Performance Tools for MPI and OpenMP 4.0 Applications for CPU and Accelerators/Coprocessors	14
	1:30pm-5pm	Tutorial	Getting Started with Vector Programming using AVX-512 on Multicore and Many-Core Platforms	15
	1:30pm-5pm	Tutorial	Kokkos: Enabling Manycore Performance Portability for C++ Applications and Domain Specific Libraries/Languages	18D
	1:30pm-5pm	Tutorial	Measuring the Power/Energy of Modern Hardware	18A
	2:00pm-5:30pm	Workshop	EduHPC2015: Workshop on Education for HPC	Hilton Salon K
	2:00pm-5:30pm	Workshop	ISAV2015: First Workshop on <i>In Situ</i> Infrastructures for Enabling Extreme-Scale Analysis and Visualization	Hilton Salon D
	2:00pm-5:30pm	Workshop	NDM-15: Fifth International Workshop on Network-Aware Data Management	Hilton 410
	3pm-3:30pm	Tutorial	Refreshment Break	Level 4 -Lobby
	3pm-3:30pm	Workshop	Refreshment Break	Hilton 4th Floor
	3pm-5pm	HPC for Undergraduates	Orientation	Hilton 404
	3:30pm-5pm	Students@SC	Mentor-Protégé Mixer	9ABC
	5pm-11:55pm	SCC	Student Cluster Competition Kickoff	Palazzo
	5:30pm-6:30pm	Keynote & Plenary Talks	HPC Matters Plenary (Diane Bryant)	Ballroom DEFG
	7pm-9pm	Social Event	Gala Opening Reception	Exhibit Hall

TUESDAY, NOVEMBER 17

	Time	Event Type	Session Title	Location
	6am-11:55pm	Students@SC	Student Cluster Competition	Palazzo
	8:30am-10am	Keynote & Plenary Talks	Keynote - Mr. Alan Alda - Getting Beyond a Blind Date with Science: Communicating Science for Scientists	Ballroom DEFG
	8:30am-5pm	Posters	Research Posters ACM Student Research Competition Posters	Level 4 - Lobby
	9am-5:30pm	Emerging Technologies	Emerging Technologies Exhibits	14
	10am-10:30am	Tech Program	Refreshment Break	Level 4 - Lobby
	10:30am-11:30am	Paper	Data Clustering	18CD
	10:30am-12pm	Paper	Applications: Material Science	18AB
	10:30am-12pm	Paper	Cache and Memory Subsystems	19AB
	10:30am-12pm	Birds of a Feather	HDF5: State of the Union	13A
	10:30am-12pm	Birds of a Feather	Lustre Community BOF: Enabling Big Data with Lustre	15
	10:30am-12pm	Birds of a Feather	Planning for Visualization on the Xeon Phi	13B
	10:30am-12pm	Exhibitor Forum	HPC Futures and Exascale	12AB
	10:30am-12pm	HPC for Undergraduates	Introduction to HPC Research	Hilton 404
	10:30am-12pm	Invited Talks	Invited Talks Session 1	Ballroom DEFG
	12:15pm-1:15pm	Birds of a Feather	Eleventh Graph500 List	18CD
	12:15pm-1:15pm	Birds of a Feather	Getting Scientific Software Installed: Tools and Best Practices	19AB
	12:15pm-1:15pm	Birds of a Feather	Integrating Data Commons and Other Data Infrastructure with HPC to Infrastructure with HPC to Research and Discovery	15
	12:15pm-1:15pm	Birds of a Feather	Operating System and Runtime for Exascale	17AB
	12:15pm-1:15pm	Birds of a Feather	SAGE2: Scalable Amplified Group Environment for Global Collaboration	13A
			Global Collaboration	
	12:15pm-1:15pm	Birds of a Feather	The Challenge of A Billion Billion Calculations per Second: InfiniBand	18AB
			Roadmap Shows the Future of the High Performance Standard Interconnect for Exascale Programs	
	12:15pm-1:15pm	Birds of a Feather	OpenSHMEM: User Experience, Tool Ecosystem, Version 1.2 and Beyond	13B
	12:15pm-1:15pm	Birds of a Feather	SIGHPC Annual Meeting	16AB
	1:00pm-3:30pm	HPC Impact Showcase	HPC Impact Showcase - Tuesday	12AB
	1:30pm-2:15pm	Award Presentations	SC15 Test of Time Award Special Lecture	Ballroom D
	1:30pm-3pm	Birds of a Feather	Advancing the State of the Art in Network APIs - The OpenFabrics Interface APIs	13A
	1:30pm-3pm	Birds of a Feather	Understanding User-Level Activity on Today's Supercomputers with XALT	13B
	1:30pm-3pm	Birds of a Feather	Virtualization and Clouds in HPC: Motivation, Challenges and Lessons Learned	15
	1:30pm-3pm	Panel	Post Moore's Law Computing: Digital versus Neuromorphic versus Quantum	16AB
	1:30pm-3pm	Paper	Applications: Biophysics and Genomics	18AB

TUESDAY, NOVEMBER 17

	Time	Event Type	Session Title	Location
	1:30pm-3pm	Paper	GPU Memory Management	19AB
	1:30pm-3pm	Paper	Scalable Storage Systems	18CD
	3pm-3:30pm	Tech Program	Refreshment Break	Level 4 - Lobby
	3:30pm-5pm	Paper	Applications: Folding, Imaging and Proteins	18AB
	3:30pm-5pm	Paper	Graph Analytics on HPC Systems	19AB
	3:30pm-5pm	Paper	MPI/Communication	18CD
	3:30pm-5pm	Birds of a Feather	Characterizing Extreme-Scale Computational and	13A
			Data-Intensive Workflows	
	3:30pm-5pm	Birds of a Feather	Performance Reproducibility in HPC -	13B
			Challenges and State-of-the-Art	
	3:30pm-5pm	Birds of a Feather	UCX - Communication Framework for Exascale	15
			Programming Environments	
	3:30pm-5pm	Invited Talks	Invited Talks Session 2	Ballroom D
	3:30pm-5pm	Panel	Future of Memory Technology for Exascale and Beyond III	16AB
			Beyond III	
	3:30pm-5pm	Exhibitor Forum	Hardware and Architecture	12AB
	5:15pm-7pm	Posters Reception	Research & ACM SRC Poster Reception	Level 4 - Lobby
	5:30pm-7pm	Birds of a Feather	Dynamic Liquid Cooling, Telemetry and Controls: Opportunity for Improved TCO?	18CD
	5:30pm-7pm	Birds of a Feather	Interactivity in Supercomputing	19AB
	5:30pm-7pm	Birds of a Feather	Maximizing Job Performance, Predictability and Manageability with Torque 6.0	13A
	5:30pm-7pm	Birds of a Feather	MPICH: A High-Performance Open-Source MPI Implementation	17AB
	5:30pm-7pm	Birds of a Feather	OpenMP: Where are We and What's Next?	18AB
	5:30pm-7pm	Birds of a Feather	Reconfigurable Supercomputing	16AB
	5:30pm-7pm	Birds of a Feather	Strategies for Academic HPC Centers	15
	5:30pm-7pm	Birds of a Feather	The Future of File Systems and Benchmarking, or Where Are We Going and How Do We Know We Got There?	13B
	5:30pm-7pm	Birds of a Feather	TOP500 Supercomputers	Ballroom D

WEDNESDAY, NOVEMBER 18

	Time	Event Type	Session Title	Location
	8:30am-10am	Keynote & Plenary Talks	Cray/Fernbach/Kennedy Award Recipients Talks	Ballroom D
	8:30am-5pm	Posters	Research and ACM SRC Posters	Level 4 - Lobby
	9am-5:30pm	Emerging Technologies	Emerging Technologies Exhibits	14
	10am-10:30am	Tech Program	Refreshment Break	Level 4 - Lobby
	10am-3pm	Students@SC	Student-Postdoc Job & Opportunity Fair	9ABC
	10:30am-12pm	Awards Presentations	ACM Gordon Bell Finalists I	17AB
	10:30am-12pm	Exhibitor Forum	Effective Application of HPC	12AB
	10:30am-12pm	HPC for Undergraduates	Graduate Student Perspective	Hilton 404
	10:30am-12pm	Invited Talks	Invited Talks Session 3	Ballroom D
	10:30am-12pm	Panel	Supercomputing and Big Data: From Collision to Convergence	16AB
	10:30am-12pm	Paper	Cloud Resource Management	19AB
	10:30am-12pm	Paper	Interconnection Networks	18CD
	10:30am-12pm	Paper	State of the Practice: Infrastructure Management	18AB
	10:30am-12pm	Birds of a Feather	Big Data and Exascale Computing (BDEC) Community Report	15
	10:30am-12pm	Birds of a Feather	Connecting HPC and High Performance Networks for Scientists and Researchers	13A
	10:30am-12pm	Birds of a Feather	Women in HPC: Pathways and Roadblocks	13B
	10:30am-12pm	SCi Vis & Data Analytics	Scientific Visualization & Data Analytics Showcase	Ballroom E
	12:15pm-1:15pm	Birds of a Feather	Collaborative Paradigms for Developing HPC in Constrained Environments	17AB
	12:15pm-1:15pm	Birds of a Feather	Fresco: An Open Failure Data Repository for Dependability Research and Practice	Ballroom G
	12:15pm-1:15pm	Birds of a Feather	HPC Education: Meeting of the SIGHPC Education Chapter	15
	12:15pm-1:15pm	Birds of a Feather	Migrating Legacy Applications to Emerging Hardware	16AB
	12:15pm-1:15pm	Birds of a Feather	Oil & Gas Community: Enabling FWI for Exascale	13A
	12:15pm-1:15pm	Birds of a Feather	Open MPI State of the Union	18CD
	12:15pm-1:15pm	Birds of a Feather	QuantumChemistry500	13B
	12:15pm-1:15pm	Birds of a Feather	The 2015 Ethernet Roadmap - Are We to Terabit Ethernet Yet?	19A
	12:15pm-1:15pm	Birds of a Feather	The Open Community Runtime (OCR) Framework for Extreme Scale Systems	18AB
	12:15pm-1:15pm	Birds of a Feather	The Partitioned Global Address Space (PGAS) Model	Ballroom F
	1:00pm-3:30pm	HPC Impact Showcase	HPC Impact Showcase - Wednesday	12AB
	1:30pm-3pm	Birds of a Feather	HPC 2020 in the BioPharma Industry	13A
	1:30pm-3pm	Birds of a Feather	HPCG Benchmark Update	15
	1:30pm-3pm	Birds of a Feather	Supercomputing After the End of Moore's Law	13B
	1:30pm-3pm	Invited Talks	Invited Talks Session 4	Ballroom D
	1:30pm-3pm	Panel	Mentoring Undergraduates Through Competition	16AB
	1:30pm-3pm	Paper	Applications: Climate and Weather	18CD
	1:30pm-3pm	Paper	Data Transfers and Data Intensive Applications	19AB
	1:30pm-3pm	Paper	Performance Tools and Models	18AB

WEDNESDAY, NOVEMBER 18

	Time	Event Type	Session Title	Location
	3pm-3:30pm	Tech Program	Refreshment Break	Level 4 - Lobby
	3:30pm-5pm	Exhibitor Forum	Software for HPC	12AB
	3:30pm-5pm	Invited Talks	Invited Talks Session 5	Ballroom D
	3:30pm-5pm	Panel	Programming Models for Parallel Architectures and Requirements for Pre-Exascale	16AB
	3:30pm-5pm	Paper	<i>In Situ</i> (Simulation Time) Analysis	18CD
	3:30pm-5pm	Paper	Linear Algebra	18AB
	3:30pm-5pm	Paper	Management of Graph Workloads	19AB
	3:30pm-5pm	Birds of a Feather	Fault Tolerant MPI Applications with ULFM	13A
	3:30pm-5pm	Birds of a Feather	The Message Passing Interface: MPI 3.1 Released, Next Stop MPI 4.0	15
	3:30pm-5:15pm	ACM SRC	ACM SRC Poster Presentations	17AB
	4pm-6pm	Family Day	Registered Attendee with Children ages 12+	Exhibit Hall
	5pm-6pm	SCC	Student Cluster Competition Finale	Palazzo
	5pm-6pm	Students@SC	Social Hour with Interesting People	9ABC
	5:30pm-7pm	Birds of a Feather	America's HPC Collaboration	Hilton 404
	5:30pm-7pm	Birds of a Feather	Ceph in HPC Environments	Hilton 406
	5:30pm-7pm	Birds of a Feather	Challenges in Managing Small HPC Centers	15
	5:30pm-7pm	Birds of a Feather	Flocking Together: Experience the Diverse OpenCL Ecosystem	17AB
	5:30pm-7pm	Birds of a Feather	High Performance Geometric Multigrid (HPGMG): an HPC Benchmark for Modern Architectures and Metric for Machine Ranking	18CD
	5:30pm-7pm	Birds of a Feather	Monitoring Large-Scale HPC Systems: Data Analytics and Insights	Ballroom E
	5:30pm-7pm	Birds of a Feather	OpenACC API 2.5: User Experience, Vendor Reaction, Relevance and Roadmap	Hilton Salon C
	5:30pm-7pm	Birds of a Feather	Paving the way for Performance on Intel® Knights Landing Processor and Beyond: Unleashing the Power of Next-Generation Many-Core Processors	18AB
	5:30pm-7pm	Birds of a Feather	PBS Community: Successes and Challenges in HPC Job Scheduling and Workload Management	Hilton 410
	5:30pm-7pm	Birds of a Feather	Power API for HPC: Standardizing Power Measurement and Control	13A
	5:30pm-7pm	Birds of a Feather	Software Engineering for Computational Science and Engineering on Supercomputers	Hilton 408
	5:30pm-7pm	Birds of a Feather	The Green500 List and its Continuing Evolution	16AB
	5:30pm-7pm	Birds of a Feather	Towards Standardized, Portable and Lightweight User-Level Threads and Tasks	19AB
	5:30pm-7pm	Birds of a Feather	Two Tiers Scalable Storage: Building POSIX-Like Namespaces with Object Stores	Hilton Salon A
	5:30pm-7pm	Birds of a Feather	U.S. Federal HPC Funding and Engagement Opportunities	Hilton Salon D

THURSDAY, NOVEMBER 19

	Time	Event Type	Session Title	Location
	8:30am-10am	Invited Talks	Plenary Invited Talks Session 6	Ballroom D
	8:30am-5pm	Posters	Research & ACM SRC Posters	Level 4 -Lobby
	9am-5:30pm	Emerging Technologies	Emerging Technologies Exhibits	14
	10am-10:30am	Tech Program	Refreshment Break	Level 4 -Lobby
	10:30am-11:30am	Paper	Sampling in Matrix Computations	18CD
	10:30am-11:30am	Awards Presentation	ACM Gordon Bell Finalists II	17AB
	10:30am-12pm	Paper	Programming Tools	18AB
	10:30am-12pm	Paper	Resource Management	19AB
	10:30am-12pm	Birds of a Feather	A Cohesive, Comprehensive Open Community HPC Software Stack	15
	10:30am-12pm	Birds of a Feather	Variability in Large-Scale, High-Performance Systems	13B
	10:30am-12pm	Doctoral Showcase	Doctoral Showcase	Ballroom E
	10:30am-12pm	Exhibitor Forum	Moving, Managing and Storing Data	12AB
	10:30am-12pm	HPC for Undergraduates	Careers in HPC	Hilton 404
	10:30am-12pm	Invited Talks	Invited Talks Session 7	Ballroom D
	10:30am-12pm	Panel	Asynchronous Many-Task Programming Models for Next Generation Platforms	16AB
	12:15pm-1:15pm	Awards Presentation	Award Ceremony	Ballroom D
	12:15pm-1:15pm	Birds of a Feather	Charm++ and AMPL: Adaptive and Asynchronous Parallel Programming	13B
	12:15pm-1:15pm	Birds of a Feather	Charting the PMIx Roadmap	15
	12:15pm-1:15pm	Birds of a Feather	CSinParallel.org : Update and Free Format Review	18AB
	12:15pm-1:15pm	Birds of a Feather	EMBRACE: Toward a New Community-Driven Workshop to Advance the Science of Benchmarking	18CD
	12:15pm-1:15pm	Birds of a Feather	Identifying a Few, High-Leverage Energy-Efficiency Metrics	19AB
	12:15pm-1:15pm	Birds of a Feather	Recruiting Non-Traditional HPC Users: High-Performance Communications for HPC	13A
	12:15pm-1:15pm	Birds of a Feather	Reproducibility of High Performance Codes and Simulations – Tools, Techniques, Debugging	17AB
	12:15pm-1:15pm	Birds of a Feather	Slurm User Group Meeting	16AB
	1pm-3:30pm	HPC Impact Showcase	HPC Impact Showcase - Thursday	12AB
	1:30pm-3pm	Birds of a Feather	Analyzing Parallel I/O	13B
	1:30pm-3pm	Birds of a Feather	Impacting Cancer with HPC: Opportunities and Challenges	15
	1:30pm-3pm	Birds of a Feather	Scalable <i>In Situ</i> Data Analysis and Visualization Using VisIt/Libsim	13A
	1:30pm-3pm	Doctoral Showcase	Doctoral Showcase	Ballroom E
	1:30pm-3pm	Panel	Towards an Open Software Stack for Exascale Computing	16AB
	1:30pm-3pm	Paper	Graph Algorithms and Benchmarks	18CD
	1:30pm-3pm	Paper	Resilience	19AB
	1:30pm-3pm	Paper	State of the Practice: Measuring Systems	18AB
	3pm-3:30pm	Tech Program	Refreshment Break	Level 4 - Lobby

THURSDAY, NOVEMBER 19

	Time	Event Type	Session Title	Location
	3:30pm-4:30pm	Paper	Tensor Computation	18CD
	3:30pm-5pm	Birds of a Feather	NSF Big Data Regional Innovation Hubs	15
	3:30pm-5pm	Birds of a Feather	Taking on Exascale Challenges: Key Lessons and International Collaboration Opportunities Delivered by European Cutting-Edge HPC Initiatives	13A
	5:30pm-7pm	Birds of a Feather	HPC Systems Engineering, Administration and Organization	13B
	3:30pm-5pm	Doctoral Showcase	Doctoral Showcase	Ballroom E
	3:30pm-5pm	Panel	Procuring Supercomputers: Best Practices and Lessons Learned	16AB
	3:30pm-5pm	Paper	Power-Constrained Computing	19AB
	3:30pm-5pm	Paper	Programming Systems	18AB
	6pm-9pm	Social Event	Technical Program Conference Reception	The University of Texas at Austin - Texas Memorial Football Stadium

FRIDAY, NOVEMBER 20

	Time	Event Type	Session Title	Location
	8:30am-10am	Panel	HPC Transforms DoD, DOE and Industrial Product Design, Development and Acquisition	16AB
	8:30am-10am	Panel	Return of HPC Survivor: Outwit, Outlast, Outcompute	17AB
	8:30am-12pm	Workshop	HUST2015: Second International Workshop on HPC User Support Tools	Hilton Salon B
	8:30am-12pm	Workshop	NRE2015: Numerical Reproducibility at Exascale	Hilton 400-402
	8:30am-12pm	Workshop	Producing High Performance and Sustainable Software for Molecular Simulation	Hilton 408
	8:30am-12pm	Workshop	SE-HPCCSE2015: Third International Workshop on Software Engineering for HPC in Computational Science and Engineering	Hilton Salon A
	8:30am-12pm	Workshop	Software Defined Networking (SDN) for Scientific Networking	Hilton Salon D
	8:30am-12pm	Workshop	VPA2015: Second International Workshop on Visual Performance Analysis	Hilton 410
	8:30am-12pm	Workshop	WHPCF2015: Eighth Workshop on High Performance Computational Finance	Hilton 406
	8:30am-12pm	Workshop	Women in HPC: Changing the Face of HPC	Hilton 412
	10am-10:30am	Workshop	Refreshment Break	Hilton 4th Floor
	10:30am-12pm	Panel	HPC and the Public Cloud	17AB
	10:30am-12pm	Panel	<i>In Situ</i> Methods: Hype or Necessity?	16AB

Plenary/Keynote/ Invited Talks

Plenary Talk

Monday, November 16

HPC Matters

5:30pm-6:30pm

Room: Ballroom DEFG

Fueling the Transformation

Diane Bryant (Intel Corporation)



In this lively and entertaining presentation, the audience will experience a visually powerful collection of proof points along with a thought-provoking, forward-looking glimpse of the importance of HPC on several areas of socioeconomic impact.

Join Intel's Senior Vice President, Diane Bryant, as she interacts with multiple industry luminaries for this engaging stage event. During this opening plenary discussion, the audience will be taken on a journey designed to give a deep appreciation for the accomplishments attributed to HPC affecting all of us today, along with an insightful discussion on how HPC will continue to evolve—and continue to change the way we live.

Bryant draws from her experience as Intel's former CIO, her experience in driving the strategic direction of Intel's \$14 billion dollar Data Center Group which includes the company's worldwide HPC business segment, and her passion and commitment for diversity programs – a topic of priority for the SC conference and the global HPC community.

Don't miss this powerful opening plenary session that gets to the very heart of "Why HPC Matters."

Bio: Diane M. Bryant is senior vice president and general manager of the Data Center Group for Intel Corporation. Bryant leads the worldwide organization that develops the data center platforms for the digital services economy, generating more than

\$14 billion in revenue in 2014. In her current role, she manages the data center P&L, strategy and product development for enterprise, cloud service providers, telecommunications, and high-performance computing infrastructure, spanning server, storage, and network solutions. Bryant is building the foundation for continued growth by driving new products and technologies—from high-end co-processors for supercomputers to high-density systems for the cloud, to solutions for big data analytics. Previously, Bryant was corporate vice president and chief information officer of Intel. She was responsible for the corporate-wide information technology solutions and services that enabled Intel's business strategies for growth and efficiency. Bryant received her bachelor's degree in electrical engineering from U.C. Davis in 1985 and joined Intel the same year. She attended the Stanford Executive Program and holds four U.S. patents.

Keynote

Tuesday, November 17

Chair: William Kramer (University of Illinois at Urbana-Champaign)

8:30am-10am

Room: Ballroom DEFG

Getting Beyond a Blind Date with Science: Communicating Science for Scientists

Alan Alda (Science Advocate and Emmy Award Winning Actor)



Alan Alda will share his passion for science communication. Why is it so important for scientists, engineers and health professionals to communicate effectively with the public? And how can they learn to do it better? In his talk, Professor Alda will explore these questions with his characteristic warmth and wit. He will draw

on his personal experiences, including his years as host of the TV series Scientific American Frontiers

Bio: Actor, writer, science advocate, and director are just a few of Alan Alda's many job titles. Throughout his 40-year career, he has

won seven Emmys, six Golden Globes, and three DGA awards for directing. When not garnering accolades for his roles in front of and behind the camera, Alda spent 11 years hosting *Scientific American Frontiers* on PBS. One of TV Guide's 50 Greatest Television Stars of All Time, Alda is best known for portraying Hawkeye Pierce on *M*A*S*H*, which earned him five Emmys for acting, writing, and directing, the only actor in history to win in each category for a single series. 125 million people tuned in to say goodbye, making the show's finale the most watched single TV episode in US history. He wrote, directed, and starred in several films throughout the 80s and 90s. He was nominated for a British Academy Award for Crimes and Misdemeanors and received an Oscar nomination for his performance in Martin Scorsese's *The Aviator*. Alda is scheduled to appear next in Nicholas Sparks's *The Longest Ride* and Steven Spielberg's 2015 *Cold War* spy thriller. In November 2014 he returned to Broadway, starring opposite Candace Bergen in *Love Letters*. From 1993 to 2005, Alda hosted PBS's *Scientific American Frontiers*, putting the actor up close with cutting edge advancements in chemistry, technology, biology, and physics. He hosted the 2010 PBS mini-series *The Human Spark* and wrote *Radiance: The Passion of Marie Curie*, a play about the personal life of the great scientist who discovered radium. He teamed up with PBS again in 2013 for *Brains on Trial*, a neurological look at brains in the court room. A recipient of the National Science Board's Public Service Award, Alda is a visiting professor at and founding member of Stony Brook University's Alan Alda Center for Communicating Science, where he helps develop innovative programs on how scientists communicate with the public. He is also on the Board of Directors of the World Science Festival. Alda published his New York Times bestselling memoir, *Never Have Your Dog Stuffed—And Other Things I've Learned*, in 2005. His second book, 2007's *Things I Overheard While Talking to Myself*, became a New York Times bestseller, as well. His 33 Emmy nominations include recent performances for NBC's *30 Rock*, *The West Wing* (his 6th win), and *ER*.

Invited Talks

Invited Talks Session I

Chair: David Keyes (King Abdullah University of Science and Technology)

10:30am-12pm

Room: Ballroom DEFG

System Software in Post K Supercomputer

Yutaka Ishikawa (RIKEN AICS)

The next flagship supercomputer in Japan, replacement of the K supercomputer, is being designed toward general operation in 2020. Compute nodes, based on a manycore architecture, connected by a 6-D mesh/torus network is considered. A three-level hierarchical storage system is taken into account. A heterogeneous operating system, Linux and a light-weight

kernel, is designed to build suitable environments for applications. It cannot be possible without codesign of applications that the system software is designed to make maximum utilization of compute and storage resources. After a brief introduction of the post K supercomputer architecture, the design issues of the system software will be presented. Two big-data applications, genome processing and meteorological and global environmental predictions, will be sketched out as target applications in the system software design. I will then discuss how these applications' demands affect the system software.

Bio: Dr. Ishikawa is the project leader of post K supercomputer development. From 1987 to 2001, he was a member of AIST (former Electrotechnical Laboratory), METI. From 1993 to 2001, he was the chief of Parallel and Distributed System Software Laboratory at Real World Computing Partnership. He led development of cluster system software called SCore, which was used in several large PC cluster systems around 2004. From 2002 to 2014, he was a professor at the University of Tokyo. He led the project to design a commodity-based supercomputer called T2K open supercomputer. As a result, three universities, Tsukuba, Tokyo, and Kyoto, obtained each supercomputer based on the specification in 2008. He has been involved the design of the post T2K machine since 2013, which will be operated in 2020.

Trends and Challenges in Computational Modeling of Giant Hydrocarbon Reservoirs

Ali H. Dogru (Saudi Aramco)

Giant oil and gas reservoirs continue to play an important role in providing energy to the world. State of the art technologies are currently being utilized to further explore and produce these reservoirs since a slight increase in the recovery amounts to discovering a mid-size reservoir somewhere else. Mathematical modeling and numerical simulation play a major role in managing and predicting the behavior of these systems using large super computers. With the aid of evolving measurement technologies, a vast amount of geoscience, fluid and dynamic data is now being collected. Consequently, more and more high resolution, high fidelity numerical models are being constructed. However, certain challenges still remain in model construction and simulating the dynamic behavior of these reservoirs. Challenges include determination of rock property variation between the wells, accurate location of faults, effective simulation of multi component, multiphase transient flow in fractures, complex wells and rock matrix. Computational challenges include effective parallelization of the simulator algorithms, cost effective large scale sparse linear solvers, discretization, handling multi scale physics, complex well shapes, fractures, compliant software engineering with the rapidly evolving super computer architectures and effective visualization of very large data sets. This presentation will cover examples for the giant reservoir models using more

than a billion elements, model calibration to historical data, challenges, current status and future trends in computational modeling in reservoir modeling.

Bio: Ali H. Dogru is a Saudi Aramco Fellow and Chief Technologist of Computational Modeling Technology. Before joining Saudi Aramco in 1996 he worked for Core Labs Inc from 1979 to 1982 Dallas, Texas then Mobil R&D from 1982 to 1996. His academic experiences include; University of Texas at Austin, Norwegian Institute of Technology, California Institute of Technology, University of Texas, Istanbul Technical University. He is a visiting Scientist at Earth Sciences at MIT. He has 12 US patents, recipient of the SPE's John Franklin Carll, Reservoir Description and Dynamics Honorary Membership and World Oil's Innovative Thinker Awards. He holds a PhD from The University of Texas, 1974.

Invited Talks Session II

Chair: Felix Wolf (*German Research School for Simulation Science*)

3:30pm-5pm

Room: Ballroom D

Superscalar Programming Models: Making Applications Platform Agnostic

Rosa M. Badia (Barcelona Supercomputing Center)

Programming models play a key role providing abstractions of the underlying architecture and systems to the application developer and enabling the exploitation of the computing resources possibilities from a suitable programming interface. When considering complex systems with aspects such as large scale, distribution, heterogeneity, variability, etc., it is indeed more important to offer programming paradigms that simplify the life of the programmers while still providing competitive performance results.

StarSs (Star superscalar) is a task-based family of programming models that is based on the idea of writing sequential code which is executed in parallel at runtime taking into account the data dependences between tasks. The talk will describe the evolution of this programming model and the different challenges that have been addressed in order to consider different underlying platforms from heterogeneous platforms used in HPC to distributed environments, such as federated clouds and mobile systems.

Bio: Rosa M. Badia holds a PhD in Computer Science (1994) from the Technical University of Catalonia (UPC). She is a Scientific Researcher from the Consejo Superior de Investigaciones Científicas (CSIC) and team leader of the Workflows and Distributed Computing research group at the Barcelona Supercomputing Center (BSC). She was involved in teaching and research activities at the UPC from 1989 to 2008, where she was an Associated Professor since year 1997. From 1999

European Center of Parallelism of Barcelona (CEPBA). Her current research interest are programming models for complex platforms (from multicore, GPUs to Grid/Cloud). The group lead by Dr. Badia has been developing the StarSs programming model for more than 10 years, with a high success in adoption by application developers. Currently the group focuses its efforts in two instances of StarSs: OmpSs for heterogeneous platforms and COMPSs/PyCOMPSs for distributed computing including cloud. For this last case, the group has been contributing effort on interoperability through standards, for example using OCCI to enable COMPSs to interact with several cloud providers at a time. Dr Badia has published more than 150 papers in international conferences and journals in the topics of her research. She has participated in several European projects, for example BEinGRID, Brein, CoreGRID, OGF-Europe, SIENA, TEXT and VENUS-C, and currently she is participating in the project Severo Ochoa (at Spanish level), ASCETIC, Euro-server, The Human Brain Project, EU-Brazil CloudConnect, and transPLANT and is a member of HiPEAC2 NoE.

LIQUI|> and SoLi|>: Simulation and Compilation of Quantum Algorithms

Dave Wecker (Microsoft Corporation)

Languages, compilers, and computer-aided design tools will be essential for scalable quantum computing, which promises an exponential leap in our ability to execute complex tasks. LIQUI|> and SoLi|> provide a modular software architecture for the simulation of quantum algorithms and control of quantum hardware. They provide a high level interface and are independent of a specific quantum architecture. This talk will focus on simulation of quantum algorithms in quantum chemistry and materials as well as factoring, quantum error correction and compilation for hardware implementations (<http://arxiv.org/abs/1402.4467>).

Bio: Dave came to Microsoft in 1995 and helped create the "Blender" (digital video post-production facility). He designed and worked on a Broadband MSN offering when he became architect for the Handheld PC v1 & v2 as well as AutoPC v1 and Pocket PC v1. He moved to Intelligent Interface Technology and resurrected SHRDLU for Natural Language research as well as building a state of the art Neural Network based Speech Recognition system. For the Mobile Devices Division he implemented secure DRM on e-books and Pocket PCs. He created and was director of ePeriodicals before taking on the role of Architect for Emerging Technologies. This lead to starting the Machine Learning Incubation Team and then architect for Parallel Computing Technology Strategy working on Big Data and now Quantum Computing. He has over 20 patents for Microsoft and 9 Ship-It awards. He started coding professionally in 1973, worked in the AI labs at CMU while obtaining a BSEE and MSIA and was at DEC for 13 years (ask him about DIDDLE sometime).

Wednesday, November 18

Invited Talks Session III

Chair: William L. Miller (National Science Foundation)

10:30am-12pm

Room: Ballroom D

Revealing the Hidden Universe with Supercomputer Simulations of Black Hole Mergers

Manuela Campanelli (Rochester Institute of Technology)

Supermassive black holes at the centers of galaxies power some of the most energetic phenomena in the universe. Their observations have numerous exciting consequences for our understanding of galactic evolution, black hole demographics, plasma dynamics in strong-field gravity, and general relativity. When they collide, they produce intense bursts of gravitational and electromagnetic energy and launch powerful relativistic jets. Understanding these systems requires solving the highly-nonlinear and highly-coupled field equations of General Relativity and Relativistic Magnetohydrodynamics. It is only with the use of sophisticated numerical techniques for simulations, data extraction and visualization, and running on petascale supercomputers of ten to hundreds of thousands of CPUs simultaneously that this problem is tractable. This talk will review some of the new developments in the field of numerical relativity, and relativistic astrophysics that allow us to successfully simulate and visualize the innermost workings of these violent astrophysical phenomena.

Bio: Manuela Campanelli is a professor of Mathematics and Astrophysics at the Rochester Institute of Technology. She is the director of the Center for Computational Relativity and Gravitation. Dr. Campanelli was the recipient of the Marie Curie Fellowship (1998), the American Physical Society Fellowship (2009) and the RIT Trustee Award (2014). She was also the Chair of the APS Topical Group in Gravitation in 2013. Dr. Campanelli has an extensive research experience on Einstein's theory of General Relativity, astrophysics of black holes and gravitational waves. She is known for groundbreaking work on numerical simulations of binary black hole space times and for explorations of physical effects such as "super kicks" and spin-driven orbital dynamics. In 2005, she was the lead author of a work that produced a breakthrough on binary black hole simulations. In 2007, she discovered that supermassive black holes can be ejected from most galaxies at speeds of up to 4000km/s. Her more current research focuses on computer simulations of merging supermassive black holes, and on magnetohydrodynamics simulations of their accretion disk and jet dynamics, in connection with both gravitational-wave and electromagnetic observations. Dr. Campanelli's research include numerous publications and invited presentations and reviews papers. She was highlighted by the American Physical Society's Focus, New Scientist, Astronomy, and the Laser Interferometer Gravitational-Wave Observatory's LIGO Magazine. More info can be found at: <http://ccrg.rit.edu/people/campanelli>

Supercomputing, High-Dimensional Snapshots, and Low-Dimensional Models: A Game Changing Computational Technology for Design and Virtual Test

Charbel Farhat (Stanford University)

During the last two decades, giant strides have been made in many aspects of computational engineering. Higher-fidelity mathematical models and faster numerical algorithms have been developed for an ever increasing number of applications. Linux clusters are now ubiquitous, GPUs continue to shatter computing speed barriers, and exascale machines will increase computational power by at least two orders of magnitude. More importantly, the potential of high-fidelity physics-based simulations for providing deeper understanding of complex systems and enhancing their performance has been recognized in almost every field of engineering. Yet, in many applications, high-fidelity numerical simulations remain so computationally intensive that they cannot be performed as often as needed or are more often performed in special circumstances than routinely. Consequently, the impact of supercomputing on time-critical operations such as engineering design, optimization, control, and test support has not yet fully materialized. To this effect, this talk will argue for the pressing need for a game-changing computational technology that leverages the power of supercomputing with the ability of low-dimensional computational models to perform in real-time. It will also present a candidate approach for such a technology that is based on projection-based nonlinear model reduction, and demonstrate its potential for parametric engineering problems using real-life examples from the naval, automotive, and aeronautics industries.

Bio: Charbel Farhat is the Vivian Church Hoff Professor of Aircraft Structures, Chairman of the Department of Aeronautics and Astronautics, and Director of the Army High Performance Computing Research Center at Stanford University. He is a member of the National Academy of engineering, a Fellow of AIAA, ASME, IACM, SIAM, and USACM, and a designated Highly Cited Author in Engineering by the ISI Web of Knowledge. He was knighted by the Prime Minister of France in the Order of Academic Palms and awarded the Medal of Chevalier dans l'Ordre des Palmes Academiques. He is also the recipient of many other professional and academic distinctions including the Lifetime Achievement Award from ASME, the Structures, Structural Dynamics and Materials Award from AIAA, the John von Neumann Medal from USACM, the Gauss-Newton Medal from IACM, the Gordon Bell Prize and Sidney Fernbach Award from IEEE, and the Modeling and Simulation Award from DoD. Recently, he was selected by the US Navy as a Primary Key-Influencer, flown by the Blue Angels during Fleet Week 2014, and appointed to the Air Force Science Advisory Board.

Invited Talks Session IV

Chair: Padma Raghavan (Pennsylvania State University)

1:30pm-3pm

Room: Ballroom D

Reproducibility in High Performance Computing

Victoria Stodden (University of Illinois at Urbana-Champaign)

Ensuring reliability and reproducibility in computational research raises unique challenges in the supercomputing context. Specialized architectures, extensive and customized software, and complex workflows all raise barriers to transparency, while established concepts such as validation, verification, and uncertainty quantification point ways forward. The topic has attracted national attention: President Obama's July 29 2015 Executive Order "Creating a National Strategic Computing Initiative" includes accessibility and workflow capture as objectives: (1) an XSEDE14 workshop released a report "Standing Together for Reproducibility in Large-Scale Computing"; (2) on May 5, 2015, ACM Transactions in Mathematical Software released a "Replicated Computational Results Initiative"; (3) this conference is host to a new workshop "Numerical Reproducibility at Exascale"; and (4) to name but a few examples. In this context I will outline a research agenda to establish reproducibility and reliability as a cornerstone of scientific computing.

Bio: Victoria Stodden is an associate professor in the Graduate School of Library and Information Science at the University of Illinois at Urbana-Champaign. She completed both her PhD in statistics and her law degree at Stanford University. Her research centers on the multifaceted problem of enabling reproducibility in computational science. This includes studying adequacy and robustness in replicated results, designing and implementing validation systems, developing standards of openness for data and code sharing, and resolving legal and policy barriers to disseminating reproducible research. Dr. Stodden is co-chair for the Advisory Committee for the National Science Foundation's Division of Advanced Cyber-Infrastructure and is a member of the NSF CISE directorate's Advisory Committee.

Fast and Robust Communication Avoiding Algorithms: Current Status and Future Prospects

Laura Grigori (French Institute for Research in Computer Science and Automation)

In this talk Dr. Grigori address one of the main challenges in HPC which is the increased cost of communication with respect to computation, where communication refers to data transferred either between processors or between different levels of memory hierarchy, including possibly NVMs. I will overview novel communication avoiding numerical methods and algorithms that reduce the communication to a minimum for operations that are at the heart of many calculations, in particular numerical linear algebra algorithms. Those algo-

rithms range from iterative methods as used in numerical simulations to low rank matrix approximations as used in data analytics. I will also discuss the algorithm/architecture matching of those algorithms and their integration in several applications.

Bio: Laura Grigori has obtained her PhD in Computer Science in 2001 from University Henri Poincare in France. She was a postdoctoral researcher at UC Berkeley and LBNL before joining INRIA in France in 2004, where she now leads a joint research group between INRIA, University of Pierre and Marie Curie, and CNRS, called Alpines. Her field of expertise is high performance scientific computing, numerical linear algebra, and combinatorial scientific computing. She co-authored the papers introducing communication avoiding algorithms that provably minimize communication. She is leading several projects in preconditioning, communication avoiding algorithms, and associated numerical libraries for large scale parallel machines. She is currently the Program Director of the SIAM special interest group on supercomputing.

Invited Talks Session V

Chair: Bronis R. de Supinski (Lawrence Livermore National Laboratory)

3:30pm-5pm

Room: Ballroom D

The European Supercomputing Research Program Panagiotis Tsarchopoulos (Future and Emerging Technologies, European Commission)

Over the last couple of years, through a number of policy and research initiatives, the European Union has worked to put together an ambitious supercomputing research program. As part of this effort, in autumn 2015, the European Commission has launched several new supercomputing projects covering supercomputing hardware, software and applications. This launch marks an important milestone in European supercomputing research and development. The talk will provide a detailed overview of the European supercomputing research program, its current status and its future perspectives towards exascale.

Bio: Panos Tsarchopoulos is responsible for supercomputing research projects at the Future and Emerging Technologies unit of the European Commission. He holds a PhD in computer engineering from the University of Kaiserslautern, Germany and an MBA from the UBI, Brussels, Belgium.

The National Strategic Computing Initiative

Randal E. Bryant, William T. Polk (Executive Office of the President, Office of Science and Technology Policy)

U.S. President Obama signed an Executive Order creating the National Strategic Computing Initiative (NSCI) on July 31, 2015.

In the order, he directed agencies to establish and execute a coordinated Federal strategy in high-performance computing (HPC) research, development, and deployment. The NSCI is a whole-of-government effort to be executed in collaboration with industry and academia, to maximize the benefits of HPC for the United States. The Federal Government is moving forward aggressively to realize that vision. This presentation will describe the NSCI, its current status, and some of its implications for HPC in the U.S. for the coming decade.

Bio: Randal Bryant has been on the computer science faculty at Carnegie Mellon University for over 30 years, serving as Dean of the School of Computer Science from 2004 to 2014. Starting in 2014, he also has been at the White House Office of Science and Technology Policy, where he serves as Assistant Director for IT R&D.

Bio: Tim Polk joined the National Institute of Standards and Technology in 1982, where he has concentrated on Internet security since 1987. In 2013, he joined the Office of Science and Technology Policy, where high performance computing complements his duties as Assistant Director for Cybersecurity.

Thursday, November 19

Invited Talks Session VI

Chair: Bernd Mohr (Juelich Supercomputing Center)

8:30am-10am

Room: Ballroom D

2015 Quadrennial Technology Review

Franklin Orr (Department of Energy)

The United States is in the midst of an energy revolution. Over the last decade, the United States has slashed net petroleum imports, dramatically increased shale gas production, scaled up wind and solar power, and cut the growth in electricity consumption to nearly zero through widespread efficiency measures. Technology is helping to drive this revolution, enabled by years to decades of research and development that underpin these advances in the energy system.

The Department of Energy's 2015 Quadrennial Technology Review (QTR) examines the status of the science and technology that are the foundation of our energy system, together with the research, development, demonstration, and deployment opportunities to advance them. This analysis is particularly instructive in the run up to the international climate negotiations taking place later this year at the 21st Conference of Parties, as technological advancements will be crucial to achieving global greenhouse gas emissions reductions.

During his presentation, Under Secretary for Science and Energy Lynn Orr will provide an overview of the highlights of the QTR report and discuss examples of promising research and development opportunities that can help the nation achieve a low-carbon economy.

Societal Impact of Earthquake Simulations at Extreme Scale

Thomas H. Jordan (University of Southern California)

The highly nonlinear, multiscale dynamics of large earthquakes is a wicked physics problem that challenges HPC systems at extreme computational scales. This presentation will summarize how earthquake simulations at increasing levels of scale and sophistication have contributed to our understanding of seismic phenomena, focusing on the practical use of simulations to reduce seismic risk and enhance community resilience. Milestones include the terascale simulations of large San Andreas earthquakes that culminated in the landmark 2008 ShakeOut planning exercise and the recent petascale simulations that have created the first physics-based seismic hazard models. From the latter it is shown that accurate simulations can potentially reduce the total hazard uncertainty by about one-third relative to empirical models, which would lower the exceedance probabilities at high hazard levels by orders of magnitude. Realizing this gain in forecasting probability will require enhanced computational capabilities, but it could have a broad impact on risk-reduction strategies, especially for critical facilities such as large dams, nuclear power plants, and energy transportation networks.

Bio: Thomas H. Jordan is a University Professor and the W. M. Keck Foundation Professor of Earth Sciences at the University of Southern California. His current research is focused on system-level models of earthquake processes, earthquake forecasting, continental structure and dynamics, and full-3D waveform tomography. As the director of the Southern California Earthquake Center (SCEC), he coordinates an international research program in earthquake system science that involves over 1000 scientists at more than 70 universities and research organizations. He is an author of more than 230 scientific publications, including two popular textbooks. He received his Ph.D. from the California Institute of Technology in 1972 and taught at Princeton University and the Scripps Institution of Oceanography before joining the Massachusetts Institute of Technology in 1984. He was head of MIT's Department of Earth, Atmospheric and Planetary Sciences from 1988 to 1998. He has received the Macelwane and Lehmann Medals of the American Geophysical Union and the Woollard Award and President's Medal of the Geological Society of America. He is a member of the National Academy of Sciences, the American Academy of Arts and Sciences, and the American Philosophical Society.

Invited Talks Session VII

Chair: Taisuke Boku (University of Tsukuba)

10:30am-12pm

Room: Ballroom D

The Power of Visual Analytics: Unlocking the Value of Big Data

Daniel A. Keim (University of Konstanz)

Never before in history has data been generated and collected at such high volumes as it is today. As the volumes of multi-dimensional data available to businesses, scientists, and the public increase, their effective use becomes more challenging. Visual analytics seeks to provide people with effective ways to understand and analyze large multidimensional data sets, while also enabling them to act upon their findings immediately. It integrates the analytic capabilities of the computer and the abilities of the human analyst, allowing novel discoveries and empowering individuals to take control of the analytical process. This talk presents the potential of visual analytics and discusses the role of automated versus interactive visual techniques in dealing with big data. A variety of application examples ranging from news analysis over network security to supercomputing performance analysis illustrate the exiting potential of visual analysis techniques but also their limitations.

Bio: Daniel A. Keim is professor and head of the Information Visualization and Data Analysis Research Group in the Computer Science Department of the University of Konstanz, Germany. He has been actively involved in data analysis and information visualization research for more than 20 years and developed a number of novel visual analysis techniques for very large data sets. He has been program co-chair of the IEEE InfoVis and IEEE VAST as well as the ACM SIGKDD conference, and he is member of the IEEE VAST as well as EuroVis steering

committees. He is coordinator of the German Science Foundation funded Strategic Research Initiative “Scalable Visual Analytics” and has been scientific coordinator of the European Commission funded Coordination Action “Visual Analytics - Mastering the Information Age (VisMaster)”. Dr. Keim received his Ph.D. in computer science from the University of Munich. Before joining the University of Konstanz, Dr. Keim was associate professor at the University of Halle, Germany and Senior Technology Consultant at AT&T Shannon Research Labs, NJ, USA.

Virtual and Real Flows: Challenges for Digital Special Effects

Nils Thuerey (Technical University of Munich)

Physics simulations for virtual smoke, explosions or water are by now crucial tools for special effects in feature films. Despite their wide spread use, there are central challenges getting these simulations to be controllable, fast enough for practical use and to make them believable. In this talk Dr. Thuerey explains simulation techniques for fluids in movies, and why “art directability” is crucial in these settings. A central challenge for virtual special effects is to make them faster. Ideally, previews should be interactive. At the same time, interactive effects are highly interesting for games or training simulators. He will highlight current research in flow capture and data-driven simulation which aims at shifting the computational load from run-time into a pre-computation stage, and give an outlook on future developments in this area.

Bio: Dr. Thuerey was born in Braunschweig, Germany. He received his PhD from the University of Erlangen-Nuremberg and his Post-doc at ETH Zurich - 3 years industry: ScanlineVFX (in Los Angeles, Munich and Vancouver, 1 year each) - Tech-oscar in 2013 - since 2013, assistant professor at Technischen Universität München

Award Presentations

Tuesday, November 17

SC15 Test of Time Award Special Lecture

Chair: Jack Dongarra (University of Tennessee, Knoxville)

1:30pm-2:15pm

Room: 17AB

The NAS Parallel Benchmarks - Summary and Preliminary Results

David Bailey (Lawrence Berkeley National Laboratory), Eric Barszcz, John Barton (NASA Ames Research Center); David Browning (Computer Sciences Corporation), Russell Carter (NASA Ames Research Center), Leonardo Dagum (Assia, Inc.), Rod Fatoohi (San Jose State University), Paul Frederickson (Math Cube Associates Incorporated), Tom Lasinski (Lawrence Livermore National Laboratory), Rob Schrieber (Hewlett-Packard Development Company, L.P.), Horst Simon (Lawrence Berkeley National Laboratory), Venkat Venkatakrishnan (University of Illinois at Chicago), Sisira Weeratunga (Lawrence Livermore National Laboratory)

In 1991, a team of computer scientists from the Numerical Aerodynamic Simulation Program—predecessor to the NASA Advanced Supercomputing (NAS) facility at Ames Research Center—unveiled the NAS Parallel Benchmarks (NPB), developed in response to the U.S. space agency's increasing involvement with massively parallel architectures and the need for a more rational procedure to select supercomputers to support agency missions. Then, existing benchmarks were usually specialized for vector computers, with shortfalls including parallelism-impeding tuning restrictions and insufficient problem sizes, making them inappropriate for highly parallel systems. The NPBs mimic computation and data movement of large-scale computational fluid dynamics (CFD) applications, and provide objective evaluation of parallel HPC architectures. The original NPBs featured “pencil-and-paper” specifications, which bypassed many difficulties associated with standard benchmarking methods for sequential or vector systems. The principal focus was in computational aerophysics, although most of these benchmarks have broader relevance for many

real-world scientific computing applications. The NPBs quickly became an industry standard and have since been implemented in many modern programming paradigms. Since 1991, research areas influenced by the NPBs have broadened to include network design, programming languages, compilers, and tools. Today's version is alive and well, and continues to significantly influence NASA projects, and is used around the world by national labs, universities, and computer vendors to evaluate sustained performance of highly parallel supercomputers and the capability of parallel compilers and tools.

Wednesday, November 18

Cray/Fernbach/Kennedy Award Recipients Talks

Co-Chairs: Anand Padmanabhan (University of Illinois at Urbana-Champaign); Franck Cappello (Argonne National Laboratory)

8:30am-10am

Room: Ballroom D

ACM Gordon Bell Finalists I

Chair: Taisuke Boku (University of Tsukuba)

10:30am-12pm

Room: 17AB

Massively Parallel Models of the Human Circulatory System

Amanda Randles (Duke University), Erik W. Draeger, Tomas Oppelstrup, Liam Krauss (Lawrence Livermore National Laboratory), John Gunnels (IBM Corporation)

The potential impact of blood flow simulations on the diagnosis and treatment of patients suffering from vascular disease is tremendous. Empowering models of the full arterial tree can provide insight into diseases such as arterial hypertension and enables the study of the influence of local factors on global hemodynamics. We present a new, highly scalable implementation of the Lattice Boltzmann method which addresses

key challenges such as multiscale coupling, limited memory capacity and bandwidth, and robust load balancing in complex geometries. We demonstrate the strong scaling of a three-dimensional, high-resolution simulation of hemodynamics in the systemic arterial tree on 1,572,864 cores of Blue Gene/Q. Faster calculation of flow in full arterial networks enables unprecedented risk stratification on a per-patient basis. In pursuit of this goal, we have introduced computational advances that significantly reduce time-to-solution for biofluidic simulations.

The In-Silico Lab-on-a-Chip: Petascale and High-Throughput Simulations of Microfluidics at Cell Resolution

Diego Rossinelli (ETH Zurich), Yu-Hang Tang (Brown University), Kirill Lykov (University of Italian Switzerland), Dmitry Alexeev (ETH Zurich), Massimo Bernaschi (National Research Council of Italy), Panagiotis Hajidoukas (ETH Zurich), Mauro Bisson (NVIDIA Corporation), Wayne Joubert (Oak Ridge National Laboratory), Christian Conti (ETH Zurich), George Karniadakis (Brown University), Massimiliano Fatica (NVIDIA Corporation), Igor Pivkin (University of Italian Switzerland), Petros Koumoutsakos (ETH Zurich)

We present simulations of blood and cancer cell separation in complex microfluidic channels with subcellular resolution, demonstrating unprecedented time to solution and performing at 42% of the nominal 39.4 Peta-instructions/s on the 18'688 nodes of the Titan supercomputer.

These simulations outperform by one to three orders of magnitude the current state-of-the-art in terms of numbers of cells and computational elements. We demonstrate an improvement of up to 30X over competing state-of-the-art solvers, thus setting the frontier of particle based simulations.

The present in silico lab-on-a-chip provides submicron resolution while accessing time scales relevant to engineering designs. The simulation setup follows the realism of the conditions and the geometric complexity of microfluidic experiments, and our results confirm the experimental findings. These simulations redefine the role of computational science for the development of microfluidic devices -- a technology that is becoming as important to medicine as integrated circuits have been to computers.

Pushing Back the Limit of Ab-initio Quantum Transport Simulations on Hybrid Supercomputers

Mauro Calderara, Sascha Brueck, Andreas Pedersen, Mohammad Hossein Bani-Hashemian, Joost VandeVondele, Mathieu Luisier (ETH Zurich)

The capabilities of CP2K, a density-functional theory package and OMEN, a nano-device simulator, are combined to study transport phenomena from first-principles in unprecedentedly

large nanostructures. Based on the Hamiltonian and overlap matrices generated by CP2K for a given system, OMEN solves the Schrödinger equation with open boundary conditions (OBCs) for all possible electron momenta and energies. To accelerate this core operation a robust algorithm called Split-Solve has been developed. It allows to simultaneously treat the OBCs on CPUs and the Schrödinger equation on GPUs, taking advantage of hybrid nodes. Our key achievements on the Cray-XK7 Titan are (i) a reduction in time-to-solution by more than one order of magnitude as compared to standard methods, enabling the simulation of structures with more than 50000 atoms, (ii) a parallel efficiency of 97% when scaling from 756 up to 18564 nodes, and (iii) a sustained performance of 14.1 DP-PFlop/s.

Thursday, November 19

ACM Gordon Bell Finalists II

Chair: Subhash Saini (Chair) - NASA Ames Research Center

10:30am-12pm

Room: 17AB

Implicit Nonlinear Wave Simulation with 1.08T DOF and 0.270T Unstructured Finite Elements to Enhance Comprehensive Earthquake Simulation

Kohei Fujita (RIKEN), Pher Errol Balde Quinay (Niigata University), Lalith Maddegadara (University of Tokyo), Muneo Hori (University of Tokyo), Seizo Tanaka (University of Tsukuba), Yoshihisa Shizawa (Research Organization for Information Science and Technology), Hiroshi Kobayashi (Research Organization for Information Science and Technology), Kazuo Minami (RIKEN)

This paper presents a new heroic computing method for unstructured, low-order, finite-element, implicit nonlinear wave simulation: 1.97 PFLOPS (18.6% of peak) was attained on the full K computer when solving a 1.08T-degrees-of-freedom (DOF) and 0.270T-element problem. This is 40.1 times more DOF and elements, a 2.68-fold improvement in peak performance, and 3.67 times faster in time-to-solution compared to the SC14 Gordon Bell finalist's state-of-the-art simulation. The method scales up to the full K computer with 663,552 CPU cores with 96.6% sizeup efficiency, enabling solving of a 1.08T-DOF problem in 29.7 s per time step. Using such heroic computing, we solved a practical problem involving an area 23.7 times larger than the state-of-the-art, and conducted a comprehensive earthquake simulation by combining earthquake wave propagation analysis and evacuation analysis. Application at such scale is a groundbreaking accomplishment and is expected to change the quality of earthquake disaster estimation and contribute to society.

An Extreme-Scale Implicit Solver for Complex PDEs: Highly Heterogeneous Flow in Earth's Mantle

Johann Rudi (The University of Texas at Austin), A. Cristiano I. Malossi (IBM Corporation), Tobin Isaac (The University of Texas at Austin), Georg Stadler (Courant Institute of Mathematical Sciences), Michael Gurnis (California Institute of Technology); Peter W. J. Staar, Yves Ineichen, Costas Bekas, Alessandro Curioni (IBM Corporation)

Mantle convection is the fundamental physical process within Earth's interior responsible for the thermal and geological evolution of the planet, including plate tectonics. The mantle is modeled as a viscous, incompressible, non-Newtonian fluid. The wide range of spatial scales, extreme variability and anisotropy in material properties, and severely nonlinear rheology have made global mantle convection modeling with realistic parameters prohibitive. Here we present a new implicit solver that exhibits optimal algorithmic performance and is capable of extreme scaling for hard PDE problems, such as mantle convection. To maximize accuracy and minimize runtime, the solver incorporates a number of advances, including aggressive multi-octree adaptivity, mixed continuous-discontinuous discretization, arbitrarily-high-order accuracy, hybrid spectral/geometric/algebraic multigrid, and novel Schur-complement preconditioning. These features present enormous challenges for extreme scalability. We demonstrate that—contrary to conventional wisdom—algorithmically optimal implicit solvers can be designed that scale out to 0.5 million cores for severely nonlinear, ill-conditioned, heterogeneous and localized PDEs.

Award Ceremony

Co-Chairs: Anand Padmanabhan (University of Illinois at Urbana-Champaign); Franck Cappello (Argonne National Laboratory)

12:15pm-1:15pm

Room: Ballroom D

Birds of a Feather

Tuesday, November 17

HDF5: State of the Union

10:30am-12pm

Room: 13A

Quincey Koziol, Mohamad Chaarawi (The HDF Group)

A forum for which HDF5 developers and users can interact. HDF5 developers will describe the current status of HDF5 and discuss future plans. An open discussion will follow.

Lustre Community BOF: Enabling Big Data with Lustre

10:30am-12pm

Room: 15

Charlie Carroll (OpenSFS), Frank Baetke (European Open File System)

Lustre is the leading open source file system for HPC. Since 2011 Lustre has transitioned to a community developed file system with contributors from around the world. Lustre is now more widely used and in more mission-critical installations than ever. Lustre supports many HPC and Data Analytics infrastructures including financial services, oil and gas, manufacturing, advanced web services, and basic science. At this year's Lustre Community BoF the worldwide community of Lustre developers, administrators, and solution providers will gather to discuss challenges and opportunities emerging at the intersection of HPC and Data Analytics and how Lustre can continue to support them.

Planning for Visualization on the Xeon Phi

10:30am-12pm

Room: 13B

Hank Childs (University of Oregon and Lawrence Berkeley National Laboratory)

The most popular visualization software systems on today's supercomputers were written during an era of single core CPU. While the Xeon Phi is a compelling architecture to transition

this legacy software to the many-core era, key components are missing, namely shared-memory support for algorithms in current visualization systems, and rendering libraries that can efficiently make use of the Phi. The BoF will be an exchange, where the presenters share current plans for the Phi, and the audience provides feedback about whether new advances will meet their use cases.

Eleventh Graph500 List

12:15pm-1:15pm

Room: 18CD

Richard Murphy (Micron Technology, Inc.), David Bader (Georgia Institute of Technology), Andrew Lumsdaine (Indiana University)

Data intensive supercomputer applications are increasingly important workloads, especially for "Big Data" problems, but are ill suited for most of today's computing platforms (at any scale!). As the Graph500 list has grown to over 190 entries, it has demonstrated the challenges of even simple analytics. This BoF will unveil the eleventh Graph500 list, roll out the SSSP second kernel, and present the energy metrics from the Green Graph500. It will offer a forum for community and provide a rallying point for data intensive supercomputing problems. This is the first serious approach to complement the Top 500 with data intensive applications.

Getting Scientific Software Installed: Tools & Best Practices

12:15pm-1:15pm

Room: 19AB

Kenneth Hoste (Ghent University), Robert McLay (The University of Texas at Austin)

We intend to provide a platform for presenting and discussing tools to deal with the ubiquitous problems that come forward when building and installing scientific software, which is known to be a tedious and time-consuming task.

Several user support tools for allowing scientific software to be installed and used will briefly be presented, for example EasyBuild, Lmod, JUBE, Spack, Maali, etc.

We would like to bring various experienced members of HPC user support teams and system administrators as well as users together for an open discussion on tools and best practices.

Integrating Data Commons and Other Data Infrastructure with HPC to Accelerate Research and Discovery

12:15pm-1:15pm

Room: 15

Robert Grossman (University of Chicago and Open Cloud Consortium), Satoshi Sekiguchi (National Institute of Advanced Industrial Science and Technology); Allison Heath, Maria Patterson (University of Chicago), Beth Plale (Indiana University)

A data commons is an emerging architecture that brings together and interoperates data, storage, compute, informatics tools, and applications for a research community. There is a growing global need for data commons and more general data infrastructure to support data driven challenges, including effectively transforming data, building complex models, and supporting sophisticated visualizations. This session will discuss practical experiences developing and operating data commons and other data types of high performance data infrastructure. This is a joint BoF between the Open Commons Consortium (OCC), the Research Data Alliance (RDA), and the Center for Data Intensive Science (CDIS), University of Chicago.

OpenSHMEM: User Experience, Tool Ecosystem, Version 1.2 and Beyond

12:15pm-1:1pm

Room: 13B

Stephen Poole (Department of Defense); Manjunath Gorentla Venkata, Pavel Shamis (Oak Ridge National Laboratory)

This has been a productive year for the OpenSHMEM community with the release of OpenSHMEM 1.2 and OpenSHMEM 1.3 (expected to be released at SC15), and much of this progress resulted from discussions at the SC14 OpenSHMEM BoF. Besides specification development, the OpenSHMEM ecosystem has expanded further this year with new implementations and tools. In addition to presenting these updates, the objectives of this BoF are to gather from the audience priorities for evolving the OpenSHMEM specification and to explore ideas for improving OpenSHMEM programmer productivity.

Operating System and Run-Time for Exascale

12:15pm-1:15pm

Room: 17AB

Marc Snir (Argonne National Laboratory), Arthur Maccabe (Oak Ridge National Laboratory), John Kubiawicz (University of California, Berkeley)

DOE's Advanced Scientific Computing Research (ASCR) program funded in 2013 three major research projects on Operating System and Runtime (OS/R) for extreme-scale computing: ARGO, HOBBS, and X-ARCC. These three projects share a common vision of the OS/R architecture needed for extreme-scale and explore different components of this architecture, or different implementation mechanisms.

The goals of this BoF are to present the current status of these efforts, present opportunities for collaboration and get feedback from researchers and vendors.

SAGE2: Scalable Amplified Group Environment for Global Collaboration

12:15pm-1:15pm

Room: 13A

Jason Leigh (University of Hawaii at Manoa); Maxine Brown, Luc Renambot (University of Illinois at Chicago)

SAGE2 (Scalable Amplified Group Environment) is an innovative user-centered platform for small groups or large distributed teams to share and investigate large-scale datasets on tiled display systems in order to glean insights and discoveries. SAGE2 treats these displays as a seamless ultra-high-resolution desktop, enabling users to co-locate data from various sources and juxtapose content. SAGE2 was announced at SC14 last year; it is the next-generation SAGE (Scalable Adaptive Graphics Environment), which has been the de facto operating system for managing Big Data content on tiled display walls, providing the scientific community with persistent visualization and collaboration services for global cyberinfrastructure.

SIGHPC Annual Meeting

12:15pm-1:15pm

Room: 16AB

ACM SIGHPC (Special Interest Group on High Performance Computing) is the first international group devoted exclusively to the needs of students, faculty, and practitioners in high performance computing. Members and prospective members are encouraged to attend the annual Members Meeting. SIGHPC officers and volunteers will share what has been accomplished to date, provide tips about resources available to members, and get audience input on priorities for the future. Join us for a lively discussion of what you think is important to advance your HPC activities.

The Challenge of A Billion Billion Calculations per Second: InfiniBand Roadmap Shows the Future of the High Performance Standard Interconnect for Exascale Programs

12:15pm-1:15pm

Room: 18AB

Gilad Shainer, Bill Lee (Mellanox Technologies)

InfiniBand has long been the preferred interconnect for HPC systems and the recent Top500 list proves it with the technology powering more than 50% of the worldwide exascale systems. The InfiniBand Trade Association is continuing to invest time and resources to enable InfiniBand to deliver efficiency at a large scale.

This session will share the latest additions to the InfiniBand Roadmap with the goal of enabling data center managers and architects, system administrators and developers to prepare to deploy the next generation of the specification targeting 200 Gb/s, as well as competitive analysis to other interconnect technologies.

Advancing the State of the Art in Network APIs - The OpenFabrics Interfaces APIs

1:30pm-3pm

Room: 13A

Paul Grun (Cray Inc.), Sean Hefty (Intel Corporation), Frank Yang (NetApp)

Following a BoF held at SC13, the OpenFabrics Alliance launched an open source effort to develop a family of network APIs targeting multiple consumer models. The first result was the recent release of libfabric, an API designed to support MPI for distributed and parallel computing. Close behind is development of APIs for storage, NVM and data access. This BoF is an opportunity for consumers of network services to review the state of development of these APIs and to debate how best to provide network services into the future, particularly in light of emerging technologies such as NVM.

Understanding User-Level Activity on Today's Supercomputers with XALT

1:30pm-3pm

Room: 13B

Mark Fahey (Argonne National Laboratory), Robert McLay (The University of Texas at Austin)

Let's talk real supercomputer analytics by drilling down to the level of individual batch submissions, users, and binaries. We are after everything from which libraries and functions are in demand to preventing the problems that get in the way of successful science. This BoF will bring together those with

experience and interest in technologies that can provide this type of job-level insight. The presenters will show off their tool (XALT) with a short demo that the attendees can do on their own laptops, and we'll have a lively discussion about what can be done and what's missing.

Virtualization and Clouds in HPC: Motivation, Challenges and Lessons Learned

1:30pm-3pm

Room: 15

Jonathan Mills (NASA/Goddard), Rick Wagner (San Diego Supercomputer Center), Hans-Christian Hoppe, Marie-Christine Sawley (Intel Corporation)

Virtualization and Cloud technology, having long been flourishing in academia and the private sector, are finally entering into traditional HPC centers as means of ensuring reliable operation, improving efficiency and accommodating the needs of new communities within computational science. Apart from overcoming the prejudice about performance loss ascribed to virtualization, challenges include integration of CPU and network virtualization, ensuring fast access to storage, and management of systems.

This BoF will assemble a panel of experts who will each give a short introductory statement and discuss the state of the art, challenges, technology options and requirements with the audience.

Characterizing Extreme-Scale Computational and Data-Intensive Workflows

3:30pm-5pm

Room: 13A

Edward Seidel (National Center for Supercomputing Applications), Franck Cappello, Tom Peterka (Argonne National Laboratory)

In the future, automated workflows will need to couple tasks and transfer data across millions of cores on a variety of new architectures. This BoF will explore how to characterize workflows connecting simulations, data analytics, and physical instruments. We will examine mini workflows for this exploration. Analogous to mini apps for individual codes, mini workflows allow benchmarking, performance testing, and enable the co-design of hardware and software systems for executing full-scale workflows in the future. In particular, mini workflows can be used to explore how to bridge the gap between workflows on HPC and distributed systems.

Performance Reproducibility in HPC - Challenges and State-of-the-Art**3:30pm-5pm****Room: 13B***Helmar Burkhart (University of Basel), Gerhard Wellein (University of Erlangen-Nuremberg)*

Being able to check and reproduce research results is a major scientific principle. But computational disciplines, including high-performance computing, still find it hard to guarantee this for performance results. We explore the difficulties and general targets for reproducible research, discuss methods proposed, and demonstrate prototype tools available.

We plan to have a highly interactive session in which both the audience and invited presenters will describe their current practice and needs. The target audience is HPC users that would like to make performance results reproducible, experts and educators in HPC performance analysis, as well as tool developers and providers.

UCX - Communication Framework for Exascale Programming Environments**3:30pm-5pm****Room: 15***Pavel Shamis (Oak Ridge National Laboratory), Gilad Shainer (Mellanox Technologies)*

UCX is a new community driven, low-level communication framework that targets current and future exascale programming models. In this BoF we will discuss the UCX effort with the HPC community to further expand the UCX collaboration with help of users, application developers, and hardware/software vendors. At the BoF we will present the current state of UCX and announce the project's roadmap, including its adoption in MPI, OpenSHMEM, PGAS languages, and exascale runtime implementations.

This BoF is an excellent face-to-face opportunity to engage a broader HPC community, ask for their feedback and plan the future direction of UCX.

Dynamic Liquid Cooling, Telemetry and Controls: Opportunity for Improved TCO?**5:30pm-7pm****Room: 18CD***David Grant (Oak Ridge National Laboratory), Steven Martin (Cray Inc.), Detlef Labrenz (Leibniz Supercomputing Center)*

Today's practice for liquid cooling is to use CDUs with constant flow-rate and temperature. Tomorrow's products could be designed for variable flow-rate and temperature based on actual heat removal requirements, but there would have to be more

and finer grained telemetry and controls. How much savings can be gained in energy savings compared to the incremental capital and operational costs? Application load differences can cause rack power variation, but nodes within a rack can also vary. Where is the sweet spot for implementation? At the rack, node or even component level? This BoF will address these and other related questions.

Interactivity in Supercomputing**5:30pm-7pm****Room: 19AB***Peter Messmer (NVIDIA Corporation), Fernanda Foertter (Oak Ridge National Laboratory)*

Interactively working with big data sets, in situ visualization or application steering are compelling scenarios for exploratory science, design optimizations or signal processing. However, a range of technical and sociological challenges need to be overcome to make these workflows mainstream: What simulation scenarios or problem domains can benefit most from interactivity? How can we simplify the tool chain? What center policies are needed to support highly interactive workflows? The goal of this BoF is to bring together domain scientists, tool developers, and HPC center administrators to identify the scientific impact and technical challenges of highly interactive access to HPC resources.

Maximizing Job Performance, Predictability and Manageability with Torque 6.0**5:30pm-7pm****Room: 13A***Ken Nielson, David Beer (Adaptive Computing)*

Torque 6.0, an open-source resource manager, has added powerful new job submission features which give users and administrators more control over job resources and improves job performance and predictability. These features include the ability to select specific resources on NUMA-enabled hardware, allow multiple heterogeneous resource requests per job submission, and enforce job limits using cgroups.

MPICH: A High-Performance Open-Source MPI Implementation**5:30pm-7pm****Room: 17AB***Ken Raffenetti, Pavan Balaji, Rajeev Thakur (Argonne National Laboratory)*

MPICH is a widely used, open-source implementation of the MPI message passing standard. It has been ported to many platforms and used by several vendors and research groups as the basis for their own MPI implementations. This BoF session will provide a forum for users of MPICH as well as developers

of MPI implementations derived from MPICH to discuss experiences and issues in using and porting MPICH. Future plans for MPICH will be discussed. Representatives from MPICH-derived implementations will provide brief updates on the status of their efforts. MPICH developers will also be present for an open forum discussion.

OpenMP: Where Are We and What's Next?

5:30pm-7pm

Room: 18AB

Jim Cownie (Intel Corporation (UK) Ltd.), Michael Wong (IBM Corporation)

We all know OpenMP: "Parallel loops," right? 1990s technology. Irrelevant to my current problems.

This lively interactive BoF will change your mind, showcasing modern OpenMP ("The language that lets you use all your compute resources") by presenting OpenMP 4.1 and the vision for OpenMP 5.0.

We'll have experts give short presentations (3 minutes) on key OpenMP 4.1 features, and then answer your questions on its technical aspects. Michael Wong (OpenMP's CEO) will present our vision for OpenMP 5.0, and we'll finish with an audience led discussion with a panel including members of the OpenMP Architecture Review Board.

Reconfigurable Supercomputing

5:30pm-7pm

Room: 16AB

Martin Herbordt (Boston University); Alan George, Herman Lam (University of Florida)

Reconfigurable Supercomputing (RSC) is characterized by hardware that adapts to match the needs of each application, offering unique advantages in performance per unit energy for high-end computing. 2015 is a breakout year for RSC, with datacenter deployment by Microsoft, acquisition of Altera and new devices by Intel, RSC innovations by IBM, new devices and tools by Altera and Xilinx, and success of Novo-G# (world's first large-scale RSC) in the NSF CHREC Center. This BoF introduces architectures of such systems, describes applications and tools being developed, and provides a forum for discussing emerging opportunities and issues for performance, productivity, and sustainability.

Strategies for Academic HPC Centers

5:30pm-7pm

Room: 15

Jackie Milhans (Northwestern University), Sharon Broude Geva (University of Michigan)

The purpose of this discussion is to explore some of the challenges faced by academic HPC centers. There will be three general topics of discussion: funding and sustainability; community outreach and user support; and the support of the emerging data science needs. We will begin with quick presentations from Northwestern, University of Michigan, Princeton, University of Chicago, and Columbia to introduce successes, current struggles, and future plans, followed by audience discussion and comments. The intended audience is university staff involved in research computing (primarily HPC), such as directors and managers, support staff, XSEDE campus champions, or other interested university members.

The Future of File Systems and Benchmarking, or Where Are We Going, and How Do We Know We Got There?

5:30pm-7pm

Room: 13B

Andrew Uselton (Intel Corporation), Devesh Tiwari (Oak Ridge National Laboratory), Ilene Carpenter (National Renewable Energy Laboratory)

The OpenSFS Benchmarking Working Group is organizing this BoF. The goal of this BoF is to discuss the current state of file system benchmarking practices, to identify the current gaps and to discuss how can we fill those gaps. In particular, we will discuss how to correctly, effectively, and efficiently benchmark large-scale storage/file-systems, the possible experimental methodology pitfalls, and how to avoid them. We will discuss the scalability of these practices for future emerging storage devices and beyond-petascale systems. As an outcome, we expect to release a document that will summarize our findings and areas that community needs to focus on.

TOP500 Supercomputers

5:30pm-7pm

Room: Ballroom D

Erich Strohmaier (Lawrence Berkeley National Laboratory)

The TOP500 list of supercomputers serves as a "Who's Who" in the field of High Performance Computing (HPC). It started as a list of the most powerful supercomputers in the world and has evolved to a major source of information about trends in HPC. The 46th TOP500 list will be published in November 2015 just in time for SC15.

This BoF will present detailed analyses of the TOP500 and discuss the changes in the HPC marketplace during the past years. The BoF is meant as an open forum for discussion and feedback between the TOP500 authors and the user community.

Wednesday, November 18**Big Data and Exascale Computing (BDEC) Community Report****10:30am-12pm****Room: 15**

Jack Dongarra (University of Tennessee, Knoxville), Pete Beckman (Argonne National Laboratory), Daniel Reed (University of Iowa)

The emergence of “Big Data” analytics in a wide variety of scientific fields has disrupted the entire research landscape on which emerging plans for exascale computing must develop. Participants in the international workshop series on Big Data and Extreme-scale Computing (BDEC) are endeavoring to systematically map out the ways in which the major issues associated with Big Data interact with plans for achieving exascale computing. This meeting will present an overview of the draft BDEC report and elicit SC15 community input on the development of a roadmap for the convergence on a common software infrastructure of currently disparate software ecosystems.

Connecting HPC and High Performance Networks for Scientists and Researchers**10:30am-12pm****Room: 13A**

Mary Hester (Lawrence Berkeley National Laboratory / Energy Sciences Network), Florence Hudson (Internet2), Vincenzo Capone (GEANT)

Many science communities struggle to leverage advanced network technologies in their scientific workflow. Often, scientists save data to portable media for analysis, increasing risks in reliability, scalability, and data sharing. Integration of networks and HPC to improve science workflows requires HPC user support groups and network organizations to closely collaborate, educating scientists on benefits of network-based technologies and tools to enable scientists to leverage the network.

Join this interactive BoF, led by ESnet, GÉANT, and Internet2, for networking experts, HPC user support teams, and scientists to engage in active dialogue to tackle current network challenges and begin to develop solutions.

Women in HPC: Pathways and Roadblocks**10:30am-12pm****Room: 13B**

Rebecca Hartman-Baker (National Energy Research Scientific Computing Center, Lawrence Berkeley National Laboratory), Toni Collis (Edinburgh Parallel Computing Centre), Valerie Maxville (Pawsey Supercomputing Center)

Although women comprise more than half the world’s population, they make up only a small percentage of people in the

Science, Technology, Engineering, and Mathematics (STEM) fields, including the many disciplines associated with high-performance computing (HPC). Within certain subdisciplines and occupational categories in HPC, women comprise a higher percentage of the population than the overall average. In this BoF, we seek to understand what makes these particular subfields so successful in attracting women, and how other subfields might emulate their successes.

Collaborative Paradigms for Developing HPC in Constrained Environments**12:15pm-1:15pm****Room: 17AB**

Hensley Omorodion (University of Benin), Elizabeth Leake (STEM-Trek), Linda Akli (Southeastern Universities Research Association)

Computational science and engineering provides the potential for addressing many problems facing countries on the African continent. HPC in Africa is in its embryonic stage, and growth requires significant increases in resources. Challenges include increased investment in infrastructure, accompanied by increased scope and innovation in the basic STEM disciplines. The newly chartered Special Interest Group on High Performance Computing for Resource Constrained Environments (SIGHPC-RCE) provides a forum for industry, academics, and government entities for building partnerships, collaborations, and outreach initiatives towards addressing these challenges. We invite everyone interested in learning about and solving these problems to attend this session.

Fresco: An Open Failure Data Repository for Dependability Research and Practice**12:15pm-1:15pm****Room: Ballroom G**

Saurabh Bagchi (Purdue University), Ravishankar Iyer (University of Illinois at Urbana-Champaign), Nathan Debardeleben (Los Alamos National Laboratory), Xiaohui Song (Purdue University), Zbigniew Kalbarczyk (University of Illinois at Urbana-Champaign)

This BoF will unveil a recently awarded NSF-supported effort for an open failure data repository meant to enable data-driven resiliency research for large-scale computing clusters from design, monitoring, and operational perspectives. To address the dearth of large publicly available datasets, we have started on this 3-year project to create a repository of system configuration, usage, and failure information from large computing systems. We have seeded the effort using a large Purdue computing cluster over a six-month period. Here we seek to collect requirements for a larger, multi-institution repository and demonstrate the usage and data analytics tools for the current repository.

HPC Education: Meeting of the SIGHPC Education Chapter

12:15pm-1:15pm

Room: 15

Steven Gordon (Ohio Supercomputer Center), David Halstead (National Radio Astronomy Observatory)

The SIGHPC Education chapter has as its purpose the promotion of interest in and knowledge of applications of High Performance Computing (HPC). This BOF will bring together those interested in promoting the education programs through the formal and informal activities of the chapter. The current officers will present information on the chapter organization, membership, and a review of current and proposed activities. They will then lead an open discussion from participants to solicit their ideas and feedback on chapter activities.

Migrating Legacy Applications to Emerging Hardware

12:15pm-1:15pm

Room: 16AB

Sriram Swaminarayan (Los Alamos National Laboratory), Andy Herdman (AWE), Michael Glass (Sandia National Laboratories)

Software always lives longer than expected. Hardware changes over this lifetime are hard to ignore. Current hardware presents software written in the 1990s with tremendous on-node parallelism such that an MPI-only model is insufficient. Modifying these large, complex, MPI-only applications to run well on current hardware require extensive and invasive changes. Further, new programming models for exploiting the on-node parallelism typically assume a start-from-scratch-and-application-wide approach making them difficult to use. In this BoF a panel of experts will discuss migration paths that will allow legacy applications to perform better on current and future hardware.

Oil & Gas Community BoF: Enabling FWI for Exascale

12:15pm-1:15pm

Room: 13A

Renato Miceli (SENAI CIMATEC), Gerard Gorman (Imperial College London)

While O & G revenues are huge high costs of data acquisition, drilling and production reduce profit margins to less than 10%. Strict time limits of petroleum licenses also impose a fixed timeline for exploration, from data acquisition and processing to exploratory drilling and decision-making. Only due to recent advancements in HPC new inversion algorithms and simulation approaches have become feasible, however, demanding disruptive software changes to harness the full hardware potential.

This BoF will focus on bringing together the Oil & Gas community to discuss common challenges and solutions going forward to exascale, besides facilitating international collaboration.

Open MPI State of the Union

12:15pm-1:15pm

Room: 18CD

Jeffrey Squyres (Cisco Systems, Inc.), George Bosilca (University of Tennessee, Knoxville)

It's been another great year for Open MPI. We've added new features, improved performance, completed MPI-3.0, and are continuing our journey towards a full MPI-3.1 implementation. We'll discuss what Open MPI has accomplished over the past year and present a roadmap for next year.

One of Open MPI's strength lies in its diversity: we represent many different viewpoints from across the HPC community. To that end, we'll have a Q&A session to address questions and comments from our community.

Join us to hear a state of the union for Open MPI. New contributors are welcome!

The 2015 Ethernet Roadmap - Are We to Terabit Ethernet Yet?

12:15pm-1:15pm

Room: 19AB

Scott Kipp (Brocade Communications Systems, Inc. and Ethernet Alliance), David Chalupsky (Intel Corporation)

The Ethernet Alliance has created the 2015 Ethernet Roadmap to explain the great diversification of Ethernet where 6 new Ethernet Speeds are being considered for standardization over the next 5 years. The panel discussion will discuss the path forward for Ethernet and challenges that are expected in the near future as new speeds and feeds are developed.

The Open Community Runtime (OCR) Framework for Extreme Scale Systems

12:15pm-1:15pm

Room: 18AB

Vivek Sarkar (Rice University), Barbara Chapman (University of Houston)

Exascale and extreme-scale systems impose fundamental new requirements on software to target platforms with severe energy, data movement and resiliency constraints within and across nodes, and large degrees of homogeneous and heterogeneous parallelism and locality within a node. The Open Community Runtime (OCR) was created to engage the broader community of application, compiler, runtime, and hardware experts to co-design a critical component of the future ex-

ascale hardware/software stack. We will build on the well-attended OCR BoFs in SC12, S13 and SC14, by using the recent OCR 1.0.0 release to engage with the community in setting future directions for OCR.

The Partitioned Global Address Space (PGAS) Model

12:15pm-1:15pm

Room: Ballroom F

Tarek El-Ghazawi (George Washington University), Lauren Smith (U.S. Government)

The partitioned global address space (PGAS) programming model strikes a balance between the ease of programming and potential efficiency due to locality awareness. PGAS has become even more widely acceptable as latency matters with non-uniform cache and memory access in manycore chips. There are many active efforts for PGAS paradigms such as Chapel, UPC and X10. There are also many implementations of these paradigms ranging from open source to proprietary ones. This BoF will bring researchers and practitioners from those efforts for cross-fertilization of ideas and address common issues of concern and common infrastructures software and hardware for PGAS.

HPC 2020 in the BioPharma Industry

1:30pm-3pm

Room: 13A

Brian Wong (Bristol-Myers Squibb Company), Harsha Gurukar (Merck & Co. Inc.), Michael Miller (Pfizer Inc.)

In this meeting, the Cross Pharma High Performance Computing Forum (<http://www.hpcxpharma.org>,) will present how HPC is currently incorporated in various areas in the industry, and how the Forum views HPC will evolve and play key roles by year 2020. We will then engage the audience to discuss various relevant focus areas, and exchange opinions with leaders in the audience from HPC industries and the community at large, to create a shared vision on how collectively we can leverage HPC to help enable key capabilities in the future with far-reaching impacts such as delivering cares for unmet medical needs.

HPCG Benchmark Update

1:30pm-3pm

Room: 15

Michael Heroux (Sandia National Laboratories), Jack Dongarra, Piotr Luszczek (University of Tennessee, Knoxville)

The High Performance Conjugate Gradients (HPCG) Benchmark is a new community metric for ranking high performance computing systems. The first list of results was released at ISC'14, including optimized results for systems built upon Fujitsu, Intel, NVIDIA technologies. Lists have been announced at SC14 and ISC'15, with an increase from 15 to 25 to 40 entries.

In this BoF we present an update of HPCG, highlighting version 3.0 and followwith presentations from vendors who have participated in HPCG optimization efforts. We spend the remaining time in open discussion about the future of HPCG design and implementation strategies for further improvements.

Supercomputing after the End of Moore's Law

1:30pm-3pm

Room: 13B

Erik DeBenedictis (Sandia National Laboratories), John Shalf (Lawrence Berkeley National Laboratory), Thomas Conte (Georgia Institute of Technology)

Supercomputers will inevitably reach 1 exaflops, yet the slow-down in device-level power efficiency improvements ("End of Moore's Law") makes reaching 10+ exaflops dependent on new approaches.

One scenario is the potential development of lower-power transistors and more effective use of 3D integration, preserving current HPC architectures and existing software base. However, the same devices could be used in increasingly specialized architectures like Graphics Processing Units (GPUs). This scenario would need more software recoding.

This BoF will include presentations on these and other scenarios followed by an organizational session for groups (lasting beyond the BoF) to further develop the various scenarios.

Fault Tolerant MPI Applications with ULFM

3:30pm-5pm

Room: 13A

George Bosilca (University of Tennessee), Keita Teranishi (Sandia National Laboratories)

With recent evolutions of the ULFM MPI implementation, fault tolerant MPI has become a reality. This BoF is intended as a forum between early adopting users, willing to share their experience and discuss fault tolerance techniques; new and prospective users, that want to discover the fault tolerance features now available in MPI and the current best practice; and ULFM developers presenting recent developments and sketching the future roadmap according to users' needs.

Join the meeting to see how your application could circumvent the pitfalls of exascale platforms and voice your opinion on the future of fault tolerance in MPI.

The Message Passing Interface: MPI 3.1 Released, Next Stop MPI 4.0

3:30pm-5pm

Room: 15

Martin Schulz (Lawrence Livermore National Laboratory)

The MPI forum, the standardization body for the Message Passing Interface (MPI), recently released version 3.1 of the MPI standard, which features minor improvements and corrections over MPI 3.0 released three years ago. Concurrently, the MPI forum is working on the addition of major new items for a future MPI 4.0, including mechanisms for fault tolerance, extensions for hybrid programming, improved support for RMA and new tool interfaces. We will use this BoF to introduce the changes made in MPI 3.1 and to continue an active discussion with the HPC community on features and priorities for MPI 4.0.

America's High Performance Computing Collaboration

5:30pm-7pm

Room: Hilton 404

Carlos Jaime Barrios Hernandez, Salma Jalife, Alvaro de la Ossa (SCALAC)

Since 1990 a wealth of projects has been developed in Latin America to build advanced computing platforms with the key support of different European and USA institutions. These projects constituted comprehensive initiatives of collaboration among international partners and most Latin American countries, and proved to play a key role to foster academic and industrial development. Today, the synergies that emerged from those collaborations effectively enable research, development and innovation with an important impact on the economy. This measurable growth generates new types of collaborations, following important aspects in training, education, outreach, industrial partnership, technical transference and government agreements.

Ceph in HPC Environments

5:30pm-7pm

Room: Hilton 406

James Wilgenbusch, Benjamin Lynch (University of Minnesota)

The Ceph storage platform is freely available software that can be used to deploy distributed fault-tolerant storage clusters using inexpensive commodity hardware. Ceph can be used for a wide variety of storage applications because a single system can be made to support block, file, and object interfaces. This flexibility makes Ceph an excellent storage solution for academic research computing environments, which are increasingly being asked to support complicated workflows with limited hardware and personnel budgets. This BoF will bring together experts who have deployed Ceph for a variety of applications and will illustrate how Ceph fits into research computing environments.

Challenges in Managing Small HPC Centers

5:30pm-7pm

Room: 15

Beth Anderson (Intel Corporation), Kevin Walsh (New Jersey Institute of Technology), Christopher Harrison (University of Wisconsin-Madison)

This session provides a venue for those involved in managing small/campus-scale, HPC capacity clusters to share their challenges, woes and successes. Attendees will also react to the findings of a pre-conference survey that investigated the compute, storage, submit/compile, backup and scheduler environments that are common in systems of this size. Trend data from four consecutive years will be provided to catalyze discussion.

Flocking Together: Experience the Diverse OpenCL Ecosystem

5:30pm-7pm

Room: 17AB

Simon McIntosh-Smith (University of Bristol), Tim Mattson (Intel Corporation)

Just as birds of a feather flock together, the strength of OpenCL is in how it was created and maintained by a consortium of like-minded organizations. This has resulted in its acceptance and across the ecosystem as a highly portable, non-proprietary HPC option. We will start with an overview of the newly released OpenCL 2.1 C++ kernel language and SYCL 2.1 abstraction layers. Then attendees will have the opportunity to experiment with implementations and tools from multiple vendors. We invite attendees to bring their code and their toughest questions and join the OpenCL petting zoo.

High Performance Geometric Multigrid (HPGMG): an HPC Benchmark for Modern Architectures and Metric for Machine Ranking

5:30pm-7pm

Room: 18CD

Mark Adams, Sam Williams (Lawrence Berkeley National Laboratory); Jed Brown (University of Colorado Boulder)

This BoF facilitates community participation with the HPGMG project. HPGMG is a compact benchmark to both provide architects with a tool for driving new architectures and a metric for extreme scale systems with generality comparable to High Performance Linpack (HPL). HPGMG is sensitive to multiple aspects of architectures that inhibit performance of HPC applications. We encourage community participation with contributed talks and will host an open discussion of issues relevant to extremes scale benchmarking. We present a biannual release of HPGMG metric data with analysis to provide insights into the efficacy of TOP500 machines for modern applications.

Monitoring Large-Scale HPC Systems: Data Analytics and Insights**5:30pm-7pm****Room: Ballroom E**

Jim Brandt (Sandia National Laboratories), Hans-Christian Hoppe (Intel Corporation), Mike Mason (Los Alamos National Laboratory), Mark Parsons (Edinburgh Parallel Computing Centre), Marie-Christine Sawley (Intel Corporation), Mike Showerman (National Center for Supercomputing Applications)

This BoF addresses critical issues in large-scale monitoring from the perspectives of worldwide HPC center system administrators, users, and vendors. We target methodologies, desires, and data for gaining insights from large-scale system monitoring: identifying application performance variation causes; detecting contention for shared resources and assessing impacts; discovering misbehaving users and system components; and automating these analyses. A panel of large-scale HPC stakeholders will interact with BoF attendees on these issues. We will explore creation of an open, vendor-neutral Special Interest Group. Community information, including BoF plans, archives, and plans for a follow-up webinar series will be available at: <https://sites.google.com/site/monitoring-largescalehpcsystems/>.

OpenACC API 2.5: User Experience, Vendor Reaction, Relevance, and Roadmap**5:30pm-7pm****Room: Hilton Salon C**

Duncan Poole (OpenACC), Fernanda Foertter (Oak Ridge National Laboratory), Randy Allen (Mentor Graphics)

OpenACC has had increased adoption across heterogeneous systems. Recently OpenACC compiler vendors are targeting architectures beyond GPU. Valuable user feedback has helped drive improvements to both standard and implementation. We expect OpenACC 2.5 to be complete along with the new APIs for tools added to 2.5 and we will present the features at the BoF. As in previous years, this BoF invites the user community to discuss newer features for 3.0. Additionally, we will present exciting outcomes from Hackathons that OpenACC hosted around the world; users and developers spent intense weeks with teams porting scalable applications including production codes on accelerators.

Paving the Way for Performance on Intel® Knights Landing Processors and Beyond: Unleashing the Power of Next-Generation Many-Core Processors**5:30pm-7pm****Room: 18AB**

Richard Gerber (National Energy Research Scientific Computing Center), Thomas Steinke (Zuse Institute Berlin), Kent Milfeld (Texas Advanced Computing Center)

This BoF continues the history of community building among those developing HPC applications for systems incorporating the Intel Xeon Phi many-integrated core (MIC) processor. The next-generation Intel MIC processor code-named Knights Landing introduces innovative features which expand the parameter space for optimizations. The BoF will address these challenges together with general aspects as threading, vectorization, and memory tuning.

The BoF will start with Lightning Talks that share key insights and best practices, followed by a moderated discussion among all attendees. It will close with an invitation to an ongoing discussion through the Intel Xeon Phi Users Group (IXPUG).

PBS Community BoF: Successes and Challenges in HPC Job Scheduling and Workload Management**5:30pm-7pm****Room: Hilton 410**

Bill Nitzberg (Altair Engineering, Inc.), Greg Matthews (Computer Sciences Corporation)

More than 20 years ago, NASA developed the PBS software. Now, PBS is one of the most widely used job schedulers for HPC around the world. But scheduling is hard -- every site has unique and changing processes, goals, and requirements -- so sharing successes and challenges is vital to evolving the technology to address both current and future needs.

Join fellow members of the PBS community to share challenges and solutions and learn others' tips and tricks for HPC scheduling in the modern world, including exascale, clouds, Big Data, power management, GPUs, Xeon Phi, containers, and more.

Power API for HPC: Standardizing Power Measurement and Control**5:30pm-7pm****Room: 13A**

Stephen Olivier, Ryan Grant, James Laros (Sandia National Laboratories)

The HPC community faces considerable constraints on power and energy at HPC installations going forward. A standardized, vendor-neutral API for power measurement and control is sorely needed for portable solutions to these issues at the various layers of the software stack. In this BoF, we discuss the

Power API, which was designed by Sandia National Laboratories and reviewed by representatives of Intel, AMD, IBM, Cray, and other industry, laboratory, and academic partners. The BoF will introduce the API in detail and feature an interactive panel discussion with experts from currently implementing and adopting organizations, with ample time for audience questions and comments.

Software Engineering for Computational Science and Engineering on Supercomputers

5:30pm-7pm

Room: Hilton 408

David Bernholdt (Oak Ridge National Laboratory), Neil Chue Hong (University of Edinburgh), Kengo Nakajima (University of Tokyo), Daniel Katz (University of Chicago and Argonne National Laboratory), Mike Heroux (Sandia National Laboratories), Felix Schuermann (Swiss Federal Institute of Technology in Lausanne)

Software engineering (SWE) for computational science and engineering (CSE) is challenging, with ever-more sophisticated, higher fidelity simulation of ever-larger and more complex problems involving larger data volumes, more domains and more researchers. Targeting high-end computers multiplies these challenges. We invest a great deal in creating these codes, but we rarely talk about that experience. Instead we focus on the results.

Our goal is to raise awareness of SWE for CSE on supercomputers as a major challenge, and to begin the development of an international “community of practice” to continue these important discussions outside of annual workshops and other “traditional” venues.

The Green500 List and its Continuing Evolution

5:30pm-7pm

Room: 16AB

Wu Feng (CompuGreen, LLC.), Erich Strohmaier (Lawrence Berkeley National Laboratory), Natalie Bates (Energy Efficient HPC Working Group)

The Green500 encourages sustainable supercomputing by raising awareness in the energy efficiency of such systems. This BoF will present evolving metrics, methodologies, and workloads for energy-efficient HPC, trends across the Green500 and highlights from the latest Green500 List. In addition, the BoF will review results from implementing more rigorous Green500 run rules as of Nov’15. These changes include requiring that average power be measured over the entire core run, that the network interconnect be included, and that a larger fraction of nodes be instrumented. The BoF will close with an awards presentation, recognizing the most energy-efficient supercomputers in the world.

Towards Standardized, Portable and Lightweight User-Level Threads and Tasks

5:30pm-7pm

Room: 19AB

Pavan Balaji, Sangmin Seo (Argonne National Laboratory)

This BoF session aims to bring together researchers, developers, vendors and other enthusiasts interested in user-level threading and tasking models to understand the current state of art and requirements of the broader community. The idea is to utilize this BoF as a mechanism to kick-off a standardization effort for lightweight user-level threads and tasks. If things go as planned, this BoF series will be continued in the future years to provide information of the standardization process to the community and to attract more participants.

Two Tiers Scalable Storage: Building POSIX-Like Namespaces with Object Stores

5:30pm-7pm

Room: Hilton Salon A

Sorin Faibish (EMC Corporation), Gary Grider (Los Alamos National Laboratory), John Bent (EMC Corporation)

Storage has seen a somewhat circular evolution over the past decade. Lack of scalability for POSIX namespaces led to flat namespace object stores. New concerns however are emerging: the lack of hierarchical namespace, to which we’re accustomed, is a hindrance to widespread object store acceptance. Emerging systems are now accessing a POSIX-like namespace layered above an object store. Have we returned to our starting point or have we made actual progress? This BoF presents talks about this trend, discussion of two exemplar systems (LANL and EMC) and their pitfalls, their relationship to burst buffers, and a panel discussion.

U.S. Federal HPC Funding and Engagement Opportunities

5:30pm-7pm

Room: Hilton Salon D

Darren Smith (National Oceanic and Atmospheric Administration), Frederica Darema (Air Force Office of Scientific Research)

In this BoF, representatives from several U.S. federal funding agencies will overview their programs and initiatives related to HPC. DoD, DOE, NASA, NIH, and NSF will briefly present opportunities spanning the areas of applications modeling and algorithms, big-data-enabled applications, programming environments and software tools, systems software, operations support, instrumentation systems, and cyber-infrastructures. The second half of the BoF will be dedicated to questions from the audience. The BoF provides a forum for the broader research community to learn about U.S. federal R&D funding opportunities in HPC and discuss program directions for new research and technology development and supporting infrastructures.

Thursday, November 19**A Cohesive, Comprehensive Open Community HPC Software Stack****10:30am-12pm****Room: 15***Karl Schulz (Intel Corporation), Susan Coghlan (Argonne National Laboratory)*

There is a growing sense within the HPC community of the need to have an open community effort to more efficiently build and test integrated HPC software components and tools. At this BoF, the speakers will discuss how such an effort could be structured and will describe the value such an effort has for the HPC community at large, and for their institutions in particular. This BoF is of interest to users, developers, and administrators of HPC machines as it proposes a more efficient vehicle to contribute to, develop on, and install an HPC stack on clusters and supercomputers.

Variability in Large-Scale, High-Performance Systems**10:30am-12pm****Room: 13B***Kirk Cameron (Virginia Polytechnic Institute and State University), Dimitris Nikolopoulos (Queen's University Belfast), Todd Gamblin (Lawrence Livermore National Laboratory)*

Increasingly, HPC systems exhibit variability. For example, multiple runs of a single application on an HPC system can result in runtimes that vary by an order of magnitude or more. Such variations may be caused by any combination of factors from the input data set or nondeterministic algorithms to the interconnect fabric, operating system, and hardware. System variability increases with scale and complexity, and it is known to limit both performance and energy efficiency in large-scale systems. The goal of this BoF is to bring together researchers across HPC communities to share and identify challenges in variability.

Charm++ and AMPI: Adaptive and Asynchronous Parallel Programming**12:15pm-1:15pm****Room: 13B***Laxmikant Kale (University of Illinois at Urbana-Champaign); Eric Bohm, Phil Miller (Charmworks, Inc.)*

This is a BoF for the burgeoning community interested in parallel programming using Charm++, Adaptive MPI, and the associated tools and domain-specific languages. This session will cover recent advances in Charm++ and the experiences of application developers with Charm++. There will also be a discussion on the future directions of Charm++ and opportunities to learn more and form collaborations.

Charm++ is one of the main programming systems on modern HPC systems. It offers features such as multicore and accelerator support, dynamic load balancing, fault tolerance, latency hiding, interoperability with MPI, and support for adaptivity and modularity.

Charting the PMIx Roadmap**12:15pm-1:15pm****Room: 15***Ralph Castain (Intel Corporation), David Solt (IBM Corporation), Artem Polyakov (Mellanox Technologies)*

The PMI Exascale (PMIx) community will be concluding its first year of existence this fall that included its first release and integration with several key resource managers. We'll discuss what PMIx has accomplished over the past year and present a proposed roadmap for next year.

The PMIx community includes viewpoints from across the HPC runtime community. To that end, we solicit feedback and suggestions on the roadmap in advance of the session, and will include time for a lively discussion at the meeting.

So please join us at the BoF to plan the roadmap. New contributors are welcome!

CSinParallel.org : Update and Free Format Review**12:15pm-1:15pm****Room: 18AB***Sharan Kalwani (Michigan State University), Richard Brown (St. Olaf College)*

Intended audience is CS curriculum educators at all levels, to share experiences and learn about the CSinParallel.org efforts (funded by NFS) on making materials and teaching modules available for Parallel and Distributed Computing. The format is brief presentation about the CSinParallel efforts to date, open discussion and finally an invitation to participate and contribute or share in this effort.

EMBRACE: Toward a New Community-Driven Workshop to Advance the Science of Benchmarking**12:15pm-1:15pm****Room: 18CD***David Bader (Georgia Institute of Technology), Jeremy Kepner (MIT Lincoln Laboratory), Richard Vuduc (Georgia Institute of Technology)*

EMBRACE is a new NSF-sponsored effort to develop a technically-focused and community-driven workshop on HPC benchmarking. The goal of this BoF is to gather feedback on EMBRACE's design and implementation. The audience consists of anyone interested in continued advances for the science of benchmarking, through benchmark development, characterization and analysis, comparison, use, and correlation to applications.

The BoF will begin with a short presentation of a strawman design for EMBRACE, which is modeled on the highly successful NSF-supported DIMACS Implementation Challenges. The rest of the time will be reserved for open discussion. The BoF's main output will be a report.

Identifying a Few, High-Leverage Energy Efficiency Metrics

12:15pm-1:15pm

Room: 19AB

Neena Imam (Oak Ridge National Laboratory), Torsten Wilde (Leibniz Supercomputing Center), Dale Sartor (Lawrence Berkeley National Laboratory)

How do we bring speed, clarity and objectivity to energy management in HPC centers? What are usage models and key metrics for current and envisioned energy information systems? What are metering requirements (what sensors and meters at which points and at what measurement frequency)? These are timely and critical questions for system administrators, operations and facilities managers and end-users. The purpose of this BoF is to review and provide feedback on a collaborative effort to identify requirements for HPC Data Center energy information systems. The BoF will host a panel of seasoned administrators and managers, followed by lively audience discussion.

Recruiting Non-Traditional HPC Users: High-Performance Communications for High-Performance Computing

12:15pm-1:15pm

Room: 13A

Jack Wells (Oak Ridge National Laboratory), Michele De Lorenzi (Swiss National Supercomputing Center)

This BoF, "Recruiting Non-Traditional HPC Users: High-Performance Communications for High-Performance Computing," will bring together leaders from supercomputing centers, universities, industry, and laboratories to discuss the challenges in attracting non-traditional users to supercomputing centers. We will focus our discussion on the sharing of outreach strategies and support mechanisms to engage high-potential user communities currently under represented within HPC centers, e.g., biotechnology, social science, digital humanities, experimental science, etc. The BoF will focus mainly on a lively discussion in which attendees will be encouraged to share their ideas, experiences, and best practices. We will also focus on networking and community building.

Reproducibility of High Performance Codes and Simulations – Tools, Techniques, Debugging

12:15pm-1:15pm

Room: 17AB

Miriam Leeser (Northeastern University), Dong Ahn (Lawrence Livermore National Laboratory), Michela Taufer (University of Delaware)

As applications are run on increasingly parallel and heterogeneous platforms, reproducibility of numerical results or code behaviors is becoming less and less obtainable. The same code can produce different results or occasional failures such as a crash on different hardware or even across different runs on the same hardware. This is due to many factors including non-determinism and the lack of numerical reproducibility with floating-point implementations. The target audience is application and tool developers interested in discussing case studies, tools, and techniques needed to address reproducibility on exascale systems. The format will be a panel of experts plus audience participation.

Slurm User Group Meeting

12:15pm-1:15pm

Room: 16AB

Morris Jette (SchedMD LLC), Jason Coverston (Cray Inc.)

Slurm is an open source workload manager used by many on TOP500 systems and provides a rich set of features including topology aware optimized resource allocation, the ability to expand and shrink jobs on demand, the ability to power down idle nodes and restart them as needed, hierarchical bank accounts with fair-share job prioritization and many resource limits. The meeting will consist of three parts: the Slurm development team will present details about changes in the new version 15.08, describe the Slurm roadmap, and solicit user feedback. Everyone interested in Slurm use and/or development is encouraged to attend.

Analyzing Parallel I/O

1:30pm-3pm

Room: 13B

Julian Kunkel (German Climate Computing Center), Philip Carns (Argonne National Laboratory)

Parallel application I/O performance often does not meet user expectations. Additionally, slight access pattern modifications may lead to significant changes in performance due to complex interactions between hardware and software. These challenges call for sophisticated tools to capture, analyze, understand, and tune application I/O.

In this BoF, we will highlight recent advances in monitoring tools to help address this problem. We will also encourage community discussion to compare best practices, identify gaps in measurement and analysis, and find ways to translate parallel I/O analysis into actionable outcomes for users, facility operators, and researchers.

Impacting Cancer with HPC: Opportunities and Challenges**1:30pm-3pm****Room: 15**

Eric Stahlberg (Frederick National Laboratory for Cancer Research), Patricia Kovatch (Icahn School of Medicine at Mount Sinai), Thomas Barr (Research Institute at Nationwide Children's Hospital)

HPC technologies have long been employed in multiple roles in cancer research and clinical applications. The importance of HPC in medical applications has been recently highlighted with the announcement of the National Strategic Computing Initiative. Despite the broad impact of HPC on cancer, no venue has been present to specifically bring together the SC community around this topic. This session will focus on bringing together those with an interest in furthering the impact of HPC on cancer, providing overviews of application areas and technologies, opportunities and challenges, while providing for in-depth interactive discussion across interest areas.

Scalable In Situ Data Analysis and Visualization Using VisIt/Libsim**1:30pm-3pm****Room: 13A**

Brad Whitlock (Intelligent Light), Cyrus Harrison (Lawrence Livermore National Laboratory)

In situ data analysis and visualization enables scientific insight to be extracted from large simulations, without the need for costly I/O. Libsim is a popular in situ infrastructure that enables tight-coupling of simulations with VisIt, a powerful data analysis and visualization tool. This session will present successful uses of Libsim and build a community to share in situ experiences and lessons learned. The format will consist of lightning talks that speak to Libsim features and usage, followed by community discussions of user experiences, applications, and best practices. The intended audience consists of simulation code developers, analysts, and visualization practitioners.

HPC Systems Engineering, Administration, and Organization**3:30pm-5pm****Room: 13B**

William R. Scullin (Argonne National Laboratory), Timothy Wickberg (George Washington University)

Systems chasing exascale often leave their administrators chasing yottascale problems. This BOF is a forum for the administrators, systems programmers, and support staff behind some of the largest machines in the world to share solutions and approaches to some of their most vexing issues and meet other members of the community. This year we are focusing

on bylaws for a new SIGHPC virtual chapter; how to best share with others in the community; training users, new administrators, and peers; management of expectations; and discussing new tools, tricks, and troubles.

NSF Big Data Regional Innovation Hubs**3:30pm-5pm****Room: 15**

Srinivas Aluru (Georgia Institute of Technology), Chaitanya Baru, Fen Zhao (National Science Foundation)

The National Science Foundation is establishing four Big Data Regional Innovation Hubs to foster multi-stakeholder partnerships in pursuit of regional and national big data challenges. Each hub serves the broad interests of its entire region, and they collectively give rise to a national Big Data innovation ecosystem. This session will introduce the leadership of the newly established hubs, and present initial focus areas and opportunities for engagement. Much of the session will be devoted to audience participation and interaction with cognizant NSF officers and hub leadership. Expected audience includes anyone interested in big data research, application, policy, or national strategy.

Taking on Exascale Challenges: Key Lessons and International Collaboration Opportunities Delivered by European Cutting-Edge HPC Initiatives**3:30pm-5pm****Room: 13A**

Mark Parsons, Jean-François Lavignon (Edinburgh Parallel Computing Centre); Jean-Philippe Nominé (European HPC Technology Platform and French Alternative Energies and Atomic Energy Commission), Estela Suarez (Juelich Research Center), Jesus Carretero (Charles III University of Madrid)

The established EU-funded exascale projects and initiatives (CRESTA, DEEP/DEEP-ER, MontBlanc, NUMEXAS, EXA2CT, EPIGRAM and NESUS) will present their status, lessons learned and potential cross-region synergies.

In addition, the long-term EU effort to develop, produce and exploit exascale platforms continues with 19 projects within the first part of the Horizon 2020 program addressing: HPC core technologies and architectures, programming methodologies, languages and tools, APIs and system software, new mathematical and algorithmic approaches.

A European status update will be presented with an emphasis on international collaboration opportunities and mechanisms needed to integrate different approaches, both in hardware and software.

Thursday, November 19

Doctoral Showcase

10:30am-12pm

Room: Ballroom E

The Relentless Execution Model for Task-Uncoordinated Parallel Computation in Distributed Memory Environments

Lucas A. Wilson (The University of Texas at San Antonio)

This work demonstrates the feasibility of executing tightly-coupled parallel algorithms in a task-uncoordinated fashion, where the tasks do not use any explicit interprocess communication (e.g., messages, shared memory, semaphores, etc.). The proposed model achieves this through the use of dataflow program representation, and the use of a distributed, eventually-consistent key/value store to memorize intermediate values and their associated state.

In addition, this work also details a new domain-specific language called StenSAL which allows for simple description of stencil-based scientific applications. Experiments performed on the Stampede supercomputer in Austin, Texas, demonstrate the ability of task-uncoordinated parallel execution models to scale efficiently in both shared memory and distributed memory environments.

Contech: Parallel Program Representation and High Performance Instrumentation

Brian P. Railing (Georgia Institute of Technology)

This summary of my dissertation work explores a pair of problems: how can a parallel program's execution be comprehensively represented? How would this representation be efficiently generated from the program's execution?

I demonstrated that the behavior and structure of a shared-memory parallel program can be characterized by a task graph that encodes the instructions, memory accesses, and dependen-

cies of each component of parallel work. The task graph representation can encode the actions of any threading library and is agnostic to the target architecture. Subsequently, I developed an open source, LLVM-based instrumentation framework, Contech, for generating dynamic task graphs from arbitrary parallel programs. The Contech framework supports a variety of languages, parallelization libraries, and architectures, with an average instrumentation overhead of less than 3x. Various analyses have been applied to Contech task graphs, including modeling a parallel, reconfigurable architecture.

Active Global Address Space (AGAS): Global Virtual Memory for Dynamic Adaptive Many-Tasking Runtimes

Abhishek Kulkarni (Indiana University)

While the communicating sequential processes model, as realized by the Message Passing Interface, is presently the dominant scalable computing paradigm, such programs neither excel at the irregular computation patterns present in big data and adaptive execution, nor obviously scale to exascale. Dynamic and adaptive many-tasking execution models are suggested as alternatives to MPI in these realms. In such models, programs are described as lightweight tasks concurrently operating on data residing in a global address space. By their nature, these models need load balancing to be effective. When work follows data this naturally takes the form of dynamically balancing data. We present High Performance ParallelX (HPX), a dynamic adaptive many-tasking runtime that maintains a scalable high-performance virtual global address space using distributed memory hardware. We describe the design and implementation of active global address space (AGAS) in HPX and demonstrate its benefit for global load balancing.

Distributed NoSQL Storage for Extreme-Scale System Services

Tonglin Li (Illinois Institute of Technology)

On both HPC systems and clouds the continuously widening performance gap between storage and computing resource prevents us from building scalable data-intensive systems. Distributed NoSQL storage systems are known for their ease of use

and attractive performance and are increasingly used as building blocks of large scale applications on cloud or data centers. However there are not many works on bridging the performance gap on supercomputers with NoSQL data stores. This work presents a convergence of distributed NoSQL storage systems in clouds and supercomputers. It firstly presents ZHT, a dynamic scalable zero-hop distributed key-value store, that aims to be a building block of large-scale systems on clouds and supercomputers. This work also presents several distributed systems that have adopted ZHT as well as other NoSQL systems, namely FREIDA-State, WaggleDB, and Graph/Z. These systems have been simplified due to NoSQL storage systems and shown good scalable performance.

Energy-Efficient Hybrid DRAM/NVM Main Memory

Ahmad Hassan (Queen's University Belfast)

DRAM, the de-facto technology for main memory, consumes significant static energy due to continuous leakage and refresh power. Various byte-addressable non-volatile memory (NVM) technologies promise near-zero static energy and persistence, however they suffer from increased latency and increased dynamic energy. A hybrid main memory, containing both DRAM and NVM components, can provide both low energy and high performance although such organizations require that data is placed in the appropriate component. This paper proposes a user-level software management methodology for a hybrid DRAM/NVM main memory system with an aim to reduce energy. We propose an operating system (OS) and programming interface to place data from within the application. We present a tool to help programmers decide on the placement of application data in hybrid memory. Cycle-accurate simulation of modified applications confirms that our approach is more energy-efficient than state-of-the-art hardware or OS approaches at equivalent performance.

Facilitating Irregular Applications on Many-Core Processors

Da Li (University of Missouri)

Despite their widespread use in application acceleration, many-core processors are still considered relatively difficult to program, in that they require the programmer to be familiar with both parallel programming and the hardware features of these devices. Irregular applications are characterized by irregular and unpredictable memory access patterns, frequent control flow divergence, and runtime (rather than compile time) parallelism. If the effective deployment of regular applications on many-core processors has been extensively investigated, the one of irregular applications is still far from understood. Yet, many established and emerging applications based on graphs and trees are irregular in nature. My research focuses on addressing important issues related to the deployment of irregular computations on many-core processors. Specifically, my contributions are in three directions: (1) unifying programming interfaces for many-core processors; (2) runtime support for efficient execution of applications on irregular datasets; and (3) compiler support for efficient mapping of applications onto hardware.

1:30pm-3pm

Room: Ballroom E

Efficient Multiscale Platelets Modeling Using Supercomputers

Na Zhang (Stony Brook University)

My dissertation focuses on developing multiscale models and efficient numerical algorithms for simulating platelets on supercomputers. More specifically, the development of multiple time-stepping algorithm can be applied to optimally use computing resources to model platelet structures at multiple scales, enabling the study of flow-induced platelet-mediated thrombogenicity. In order to achieve this, sophisticated parallel computing algorithms are developed and detailed performance analysis has been conducted. The performance results manifest the possibility of simulating the millisecond-scale hematology at resolutions of nanoscale platelets and mesoscale bio-flows using millions of particles. The computational methodology using multiscale models and algorithms on supercomputers will enable efficient predictive simulations for initial thrombogenicity study and may provide a useful guide for exploring mechanisms of other complex biomedical problems at disparate spatiotemporal scales.

High Performance Earthquake Simulations

Alexander N. Breuer (Technical University of Munich)

The understanding of earthquake dynamics is greatly supported by highly resolved coupled simulations of the rupture process and seismic wave propagation. This grand challenge of seismic modeling requires an immense amount of supercomputing resources. The increasing demand for parallelism and data locality often requires replacing major software parts.

I present a new computational core for the seismic simulation package SeisSol. The new core is designed to maximize value and throughput of the FLOPs performed in the underlying ADER-DG method. Included are auto-tuned matrix kernels, hybrid parallelization up to machine-size and a novel high performance clustered LTS scheme.

The presented computational core reduces time-to-solution of SeisSol by several factors and scales beyond 1 million cores. At machine-size the new core enabled a landmark simulation of the Landers 1992 earthquake. For the first time this simulation allowed the analysis of the complex rupture behavior and seismic wave propagation at high geometric complexity.

High-Performance Algorithms For Large-Scale Singular Value Problems and Big Data Applications

Lingfei Wu (William & Mary)

As “big data” has increasing influence on our daily life and research activities, it poses significant challenges on various research areas. Some applications often demand a fast solution of large-scale singular value problems; In other applications, extracting knowledge from large-scale data requires many techniques such as data mining, and high performance computing. We firstly present a preconditioned hybrid, two-stage method to effectively and accurately compute a small number of extreme singular triplets. More importantly, we have implemented a high-performance preconditioned SVD solver software, PRIMME_SVDS. PRIMME_SVDS fills a gap in production level software for computing the partial SVD, especially with preconditioning. In addition, we propose a real-time outlier detection algorithm to efficiently find blob-filaments in fusion experiments and numerical simulations. We have implemented this algorithm with hybrid MPI/OpenMP and show that we can achieve linear time speedup and complete blob detection in two or three milliseconds using a HPC cluster.

Irregular Graph Algorithms on Parallel Processing Systems

George M. Slota (Pennsylvania State University)

My dissertation research lies at the interface of high performance computing and big data analytics. Specifically, my work focuses on designing new parallel approaches for analyzing large real-world networks. The high complexity, scale, and variation of graph-structured data poses an immense challenge in the design of techniques to study and derive insight from such data. Optimizing graph analysis software for complex and heterogeneous modern HPC systems poses an additional set of challenges for algorithm designers to overcome. My primary research goals focus on tackling these problems through the optimization of graph analytics at all levels of hardware architecture, from thread to core to processor to single-node to multi-node to system-level scale. The fact that graph-structured data is so universal means that this research is useful to a large collection of data-intensive problems within both the social and physical sciences.

Framework for Lifecycle Enrichment of HPC Applications on Exascale Heterogeneous Architecture

Karan Sapra (Clemson University)

Recently, HPC has experienced advancements with the introduction of multi-core, many-core, and massively parallel architectures such as Intel’s Xeon-Phi, Graphical Processing Units (GPUs), and Field Programmable Gate Arrays (FPGAs). Although the aforementioned architectures can provide significant speedup, the performance greatly depends on the

choice of architecture for the given application. Furthermore, stepping into the exascale-computing era, heterogeneous supercomputers are gaining popularity due to wide-array applications hosted on the supercomputers. In this research, we perform enrichment and acceleration of HPC application on heterogeneous architectures. The enrichment of an HPC application is achieved in two phases: Development Phase using our Application-to-Architecture (A2A) followed by utilizing Functional Partitioning (FP) framework during runtime. We evaluate the accuracy of our framework using diverse micro-benchmark and applications that belong to different regions in the application space. These benchmark and applications encompass a broad spectrum of application behavior, making them highly suitable for our studies.

Mitigation of Failures in High Performance Computing via Runtime Techniques

Xiang Ni (University of Illinois at Urbana-Champaign)

The number of components assembled to create a supercomputer keeps increasing in pursuit of more computational power required to enable breakthroughs in science and engineering. However, the reliability and the capacity of each individual component has not increased as fast as the increase in the total number of components. As a result, the machines fail frequently and hamper smooth execution of high performance applications. This thesis strives to answer the following questions with regard to this challenge: how can a runtime system provide fault tolerance support more efficiently with minimal application intervention? What are the effective ways to detect and correct silent data corruption? Given the limited memory resource, how do we enable the execution and checkpointing of data intensive applications?

Thursday, November 19

Doctoral Showcase

3:30pm-4pm

Room: Ballroom E

Kanor: an EDSL for Declarative Communication

Nilesh N. Mahajan (Indiana University)

High performance programs, using explicit communication calls to specify communication structure, need considerable programming expertise. It is difficult to guarantee properties like deadlock freedom while avoiding significant performance degradation. We have developed a domain-specific language embedded in C++ called Kanor that allows programmers to specify the communication declaratively. Kanor programs are written in the Bulk Synchronous Style (BSP) and deadlock freedom is guaranteed by construction. In this work, we present the design and implementation of Kanor. We start with a

description of Kanor syntax and explain how the syntax makes it easier to write communication patterns. Next we describe Kanor semantics and explain how Kanor programs guarantee certain properties. Finally, we describe how the declarative nature of Kanor communication allows us to optimize Kanor programs.

Designing High Performance and Energy-Efficient MPI Collectives for Next Generation Clusters

Akshay Venkatesh (Ohio State University)

The past decade has witnessed a notable change in HPC landscape with the proliferation of accelerators/co-processors. The presence of processing cores and memory of a different nature has had an adverse impact on the MPI and PGAS programming models that are predominantly used for developing scientific applications. As accelerators/co-processors are primarily available as external PCIe devices, the external cores and memory units render systems heterogeneous. From the perspective of collectives communication routines, where hundreds to thousands of cores are likely to be involved, traditional assump-

tions such as uniform processing and communication costs within the system are rendered invalid. As the exascale report identifies communication and energy optimization among the foremost challenges in reaching the exaflop mark, this work addresses: (1) the performance of collective algorithms by distinguishing heterogeneous communication path costs and (2) the energy of both collective and point-to-point operations by generating rules that apply energy saving strategies at the fundamental underlying communication algorithms.

Doctoral Showcase Poster Session

4pm-5pm

Room: Ballroom E

All Doctoral Showcase presenters will be present, with their posters, for an open discussion.

Emerging Technologies

**Tuesday, November 17 –
Thursday, November 19**

Emerging Technologies Exhibits

Chair: *Simon McIntosh-Smith (University of Bristol);
Sadaf R. Alam (Swiss National Supercomputing Center)*
9am-5:30pm
Room: 14

Liquid Cooling for Exascale Computing

Peter Hopton (Iceotope)

Iceotope has been selected as the infrastructure and cooling provider for the largest EU Exascale research H2020 Project - ExaNeSt. Within the infrastructure group of 4, Iceotope is involved in an exciting project which has successfully secured EU funds totaling over \$20m, and is aimed at the rapid development of high efficiency, ultra-dense compute that could one day be deployed to exascale. Peter Hopton, CVO of Iceotope, will look at the heat density problem produced by Exascale computing and how Iceotope's unique technology is being adapted to power and cool high density systems. Peter will discuss key issues that need to be incorporated into infrastructure to allow for scalability. He will also provide insight into the direction of liquid cooling and its future potential as a wide scale disruptor of TCO rather than just an added cost and complexity to increase density.

Emulating Future HPC SoC Architectures

Farzad Fatollahi-Fard, David Donofrio, John Shalf (Lawrence Berkeley National Laboratory)

Current HPC ecosystems rely upon Commercial Off-the-Shelf (COTS) building blocks to enable cost-effective design by sharing costs across a larger ecosystem. Modern HPC nodes use commodity chipsets and processor chips integrated together on custom motherboards. Commodity HPC is heading into a new era where the chip acts as the "silicon motherboard"

that interconnects commodity Intellectual Property (IP) circuit building blocks to create a complete integrated System-on-a-Chip (SoC). These SoC designs have the potential for higher performance at better power efficiency than current COTS solutions. We will showcase how the advanced tools provided by the Co-Design for Exascale (CoDEX) project running on a cloud-based FPGA system can provide powerful insights into future SoC architectures tailored to the needs of HPC. The large-scale emulation environment shown here will demonstrate how we are building the tools needed to evaluate new and novel architectures at speeds fast enough to evaluate whole application performance.

LiquidCool Solutions Project Clamshell: Novel Cooling Technology

Herb Zien, Jay Ford, Kim McMahon (LiquidCool Solutions, Inc.)

LiquidCool Solutions presents an innovation in extreme energy efficiency, significant space conservation, and zero water use supercomputing. Customers are no longer bound by complicated and costly mechanical systems that consume capital, natural resources, and time. LiquidCool un-tethers compute systems from the brick and mortar data center.

Imagine what one would do if no longer bound by brick walls, complicated air handling, budget pressure, and scarcity of water. Others have attempted but produced mixed results. LiquidCool offers a maintenance-free total solution that not only cools, but protects and extends the life of the platform, while offering unmatched flexibility in system choice: any combination of motherboard, processor, memory cluster, solid state drive, and even power supply.

Project Clamshell takes existing cooling technologies to the next level with 100% of components cooled, no water, and 100% reliability. It has been selected for the Wells Fargo Innovation Incubator program for its innovative environmental technologies.

GoblinCore-64: An Emerging, Open Architecture for Data Intensive High Performance Computing*John Leidel, Xi Wang, Yong Chen (Texas Tech University)*

The current class of mainstream microprocessor architectures relies upon multi-level data caches and relatively low degrees of concurrency to solve a wide range of applications and algorithmic constructs. These mainstream architectures are well suited to efficiently executing applications that are generally considered to be cache friendly. These may include applications that operate on dense linear data structures or applications that make heavy reuse of data in cache. However, applications that are generally considered to be data intensive in nature may access memory with irregular memory request patterns or access such large data structures that they cannot reside entirely in an on-chip data cache. The goal of GoblinCore-64 (GC64) is to provide a scalable, flexible and open architecture for efficiently executing data intensive, high performance computing applications and algorithms. Further, we seek to build an architectural infrastructure with the aforementioned properties without the need for new and exotic programming models.

Emerging Technology for Persistent Challenges in Optical Interconnects and Photonics Packaging*R. Ryan Vallance, Rand Dannenberg, Shuhe Li, Yang Chen, Matt Gean (nanoPrecision Products Inc.)*

Ultra-high precision forming is an emerging technology capable of manufacturing low volumes (thousands per year) and high volumes (hundreds of millions per year) of metallic components with sub-micrometer tolerances. It is enabling the next-generation of fiber-optic packaging and connectors for photonics in supercomputing and networks. Ultra-high precision forming relies on advances in material characterization, simulation and optimization methods used to design the forming process, the manufacture of hardened tools with sub-micrometer precision, and dimensional metrology. This exhibit highlights these technical advances and illustrates how they enable the forming of : (1) meso-scale components with sub-micrometer tolerances, (2) micro-structured grooves that hold and accurately locate optical fibers, and (3) micro aspherical mirror arrays that focus and bend light from optical fibers into photonic integrated circuits, laser arrays, and detector arrays. These elements are integrated together into stamped components that simplify and reduce the cost of photonic packaging and single-mode fiber-optic cabling.

Automata Processing: A Massively Parallel Computing Solution*Susan Platt, Indranil Roy, Matt Tanner, Dan Skinner (Micron Technology, Inc.)*

Many of today's most challenging computer science problems---such as those involving very large data structures, unstructured data, random access or real-time performance requirements - require highly parallel solutions. The current implementation of parallelism can be cumbersome and complex, challenging the capabilities of traditional CPU and memory system architectures and often requiring significant effort on the part of programmers and system designers. For the past seven years, Micron Technology has been developing a hardware co-processor technology that can directly implement large-scale Non-deterministic Finite Automata (NFA) for efficient parallel execution. This new non-Von Neumann processor, currently in fabrication, borrows from the architecture of memory systems to achieve massive data parallelism, addressing complex problems in an efficient, manageable method.

openHMC - A Configurable Open-Source Hybrid Memory Cube Controller*Juri Schmidt (Heidelberg University)*

The link between the processor and the memory is one of the last remaining parallel buses and a major performance bottleneck in HPC systems. The Hybrid Memory Cube (HMC) was developed with the goal of helping overcome this memory wall. The presented project is an open-source implementation of an HMC host controller that can be configured to different datapath widths and HMC link variations. Due to its modular design, the controller can be integrated in many different system environments. Many HPC applications will benefit from using openHMC, which is a verified, no-cost alternative to commercially available HMC host controller IP.

Mineral Oil-Based Direct Liquid Cooling System*Gangwon Jo (Seoul National University), Jungho Park (Many-CoreSoft Co., Ltd.), Jaejin Lee (Seoul National University)*

Direct liquid cooling is one of the promising cooling solutions for data center servers. It reduces operating expenses and energy consumption. Moreover, it minimizes floor spaces for servers. Metal plates are attached to server components and the coolant, typically water, flows through the plates. The goal of this project is to develop a new direct liquid cooling solution that uses electrically non-conductive mineral oil instead of water. To overcome the low thermal conductivity of mineral oil, we designed special cooling plates for CPUs and GPUs that maximizes the liquid-contacting area and circulation systems that maximize the flow rate. As a result, our system keeps the temperature of CPUs and GPUs below 70°C at full load. This is

the first mineral-oil-based direct liquid cooling solution that is safe, cheap, high density, and easy to maintain. ManyCoreSoft Co., Ltd. has released the MCS-4240 cluster system that uses our liquid cooling solution.

ExaNoDe: European Exascale Processor and Memory Node Design

Denis Dutoit (French Alternative Energies and Atomic Energy Commission)

ExaNoDe is a collaborative European project within the “Horizon 2020 Framework Programme”, that investigates and develops a highly integrated, high-performance, heterogeneous System-on-a-Chip (SoC) aimed towards exascale computing. It is addressing these important challenges through the coordinated application of several innovative solutions recently deployed in HPC: ARM-v8 low-power processors for energy efficiency, 3D interposer integration for compute density and an advanced memory scheme for exabyte level capacities. The ExaNoDe SoC will embed multiple silicon “chiplets”, stacked on an active silicon interposer in order to build an innovative 3D-Integrated-Circuit (3D-IC). A full software stack allowing for multi-node capability will be developed within the project. The project will deliver a reference hardware to enable the deployment of multiple 3D-IC System-on-Chips and the evaluation, tuning and analysis of HPC mini-apps along with the associated software stack.

DOME Hot-Water Cooled Multinode 64 Bit Microserver Cluster

Ronald P. Luijten (IBM Corporation)

We provide a live demonstration of the IBM / ASTRON “DOME” 64-bit high-performance microserver, consisting on an 8-way hot water cooled cluster running a parallel application. This cluster is based on the Freescale P5040 2.2GHz 64bit ppc SoC. In addition, we will demonstrate a single node Freescale T4240 air cooled node. These systems run Fedora 20 with the XFCE desktop, CPMD, STREAM, IBM DB2 and many other applications. The demonstration system fits on a single tabletop, and only a main supply is needed for a self-contained demo (the water cooling is done with a small, portable PC-modder cooling unit).

Exhibitor Forums/ HPC Impact Showcase

Tuesday, November 17

HPC Futures and Exascale

Chair: Lucas A. Wilson (Texas Advanced Computing Center)

10:30am-12pm

Room: 12AB

Fujitsu HPC: the Next Step Toward Exascale

Toshiyuki Shimizu (Fujitsu)

Fujitsu has been leading the HPC market for well over 30 years, and today offers a comprehensive portfolio of computing products: Fujitsu's SPARC64-based supercomputer PRIMEHPC series, as well as our x86-based PRIMERGY clusters, software, and solutions, to meet wide-ranging HPC requirements. The latest addition to the PRIMEHPC series, PRIMEHPC FX100, is now available and is operational at several customer sites. Fujitsu also has been working together with RIKEN on Japan's next-generation supercomputer, and has completed the basic design phase. The national project, which is called the "FLAGSHIP2020 Project," aims to develop the successor to the K computer with 100 times more application performance, and begin its operation in fiscal year 2020. At SC15, Fujitsu will mention what we have learned using the PRIMEHPC FX100 systems with our customers, and provide updates on the Post K supercomputer development process.

NEC SX-ACE Vector Supercomputer and Its Successor Product Plan

Shintaro Momose (NEC Corporation)

NEC gives an overview of the latest model vector supercomputer SX-ACE with several performance evaluations. NEC also provides the plan for a successor product of SX-ACE. NEC has launched SX-ACE by aiming at a much higher sustained performance particularly for memory-intensive applications. It is designed based on the big core strategy targeting higher sustained performance. It provides both the world top-level single core performance of 64 GFlops and the world highest memory bandwidth per core of 64 GBytes/s. This high capabil-

ity core has demonstrated the highest peak performance ratio with competitive power consumption for the HPCG benchmark program that is designed to create a more relevant metric for ranking HPC systems than the High Performance LINPACK. NEC will also talk about the vision and concept of its future vector architectural product, which is aimed at not only conventional high-performance computing but also emerging big data analytics applications.

The Evolution Continues - HPC, Data Centric Architecture and CORAL, and Cognitive Computing for Modeling & Simulation

Dave Turek (IBM Corporation). Arthur S. (Buddy) Bland (Oak Ridge Leadership Computing Facility)

The demand for HPC solutions is continuously evolving. Data Centric architecture, moving HPC to data, addresses the explosive growth in compute and data requirements due to new workloads in big data, analytics, machine learning, and cognitive computing. Dave Turek, IBM VP HPC, and Buddy Bland, Director Oak Ridge Leadership Computing Facility, will provide examples of how data centric computing is propelling customers towards open, innovative solutions solving the most complex HPC problems. Requirements for these new workloads are converging and why conventional datacenter architectures have not kept pace. The next-generation computing architectures require adoption of innovative, open models that include tighter integration between partners using accelerators, servers, storage, fabric, workload, resource, scheduling and orchestration.

These initiatives are spawning a new frontier that requires HPC to link modeling-simulation with cognitive computing and Deep Learning. New IBM research projects in Petroleum and Manufacturing push the boundaries for analytics for next-generation HPC.

Hardware and Architecture

Chair: Paul Domagala (Argonne National Laboratory)

3:30pm-5pm

Room: 12AB

Extremely Dense 2-Phase Immersion Cooled HPC System

Mondrian Nuessle (EXTOLL GmbH)

EXTOLL network technology is a direct network and is dedicated to HPC through its high bandwidth, low latency and high hardware message rate. This technology has been brought to market by EXTOLL's TOURMALET chip which is available as PCIe Express gen3 x16 board or as a single-chip solution for integration into customers' hardware. Besides performance, EXTOLL's network technology allows for hostless nodes, i.e. one node comprises an EXTOLL NIC and an accelerator like Intel's KNC. In order to create an extremely dense system with high compute performance, EXTOLL constructed its novel 2-phase immersion cooling system GreenICE. It is capable of running 32 nodes within a 19" 10U chassis yielding 38.4 TFLOPS. 10KW heat is dissipated by hot water cooling. System design is outlined and measurements from first installations are reported.

High Performance Compute Cluster Constructed with High Density Rigid-Flex Based Modules

How Lin (i3 Electronics, Inc.)

A high performance compute cluster constructed using high density modular subsystems is described in this paper. Distance between a large number of processing nodes are minimized by using 3-D packaging and complex flex circuit technologies. The base processing module (a leaf), contains local memory and power supply, is fabricated using Direct Chip Attach (DCA) technology onto a rigid-flex substrate. A four processing node is formed by attaching four leafs to a node block (a circuit board with a centrally located connector) via flex circuits. The processing nodes are interconnected via a cluster board (motherboard). The leafs of each processing node are folded 90° with respect to the surface of the node block and intimately contracting a liquid cooled thermal management unit for heat removal. Up to 64 processing nodes can be housed on a 22" x 19" motherboard. Minimum distance between processing nodes on the same motherboard is thus achieved.

Reinstating Moore's Law with Transistor-Level 3D Using Foundry Wafers and Depreciated Process Nodes

David Chapman (Tezzaron)

More than a decade of 3D packaging efforts have yielded marginal, but sorely needed performance improvements. Yet those efforts have done nothing directly to address the demise of transistor density improvements, the undoing of Moore's Law. Even using mature process nodes, transistor-scale 3D introduced

by Tezzaron at SC15, allows transistor densities that exceed any anticipated at even 7 nm, the broadly acknowledged "end of silicon." Interconnect parasitics, which have become the primary source of delays and power consumption, can easily be reduced by factors of 10 to 1000. Transistor-scale 3D applied to memory produces dramatic reductions in latency and power and allows vastly improved transaction rates, key to many HPC applications; and can be expected to produce similar results when applied to the design and fabrication of processors and other complex ICs, allowing cost effective extension of Moore's law in conventional CMOS for another one or two decades.

Wednesday, November 18

Effective Application of HPC

Chair: Dane Skow (University of Wyoming)

10:30am-12pm

Room: 12AB

Bringing Multiple Organizations Together to Share Common HPC Resources Dynamically, Securely and Equitably

Nick Ihli (Adaptive Computing)

The necessary technology didn't exist when Hospital for Sick Kids and Princess Margaret Cancer Center dreamed up the idea of creating a central HPC environment where hospitals across Canada could share. Adaptive will take you from the initial vision through the deployment. HPC4Health converges HPC, Cloud and Big Data Solutions into a single pool of resources with each organization having dedicated resources plus a common communal pool of resources. Each organization manages their dedicated resources just as if it were a private environment. As workloads increase, automation manages each organization's growth requirements and dynamically obtains additional resources from the communal pool to handle the peak loads and then relinquishes those resources back to the communal pool, fully in-line with strict privacy policies, for the next peak workload requirement from any organization. All workloads are tracked per user/organization and accounted for with extensive reporting capabilities.

Champions - Made Possible by Cray

Tim White (Cray Inc.)

Competition is part of human nature. Sport marries our competitive instinct with another key human pursuit — entertainment. As a global market, sport generates \$600+ billion annually with an estimated additional \$700 billion coming from derivative markets such as fantasy sports, gambling, and enhanced media experiences. Pointing to the cultural importance of sport, the movie Moneyball highlights an ongoing transformation in decision making in Major League Baseball from one of tacit knowledge to actuarial modeling. In this presentation we will look at sports

analytics, discuss how HPC is furthering this transformation, and highlight complex problems such as discovering competitive advantages, injury risk, player selection, and the March Madness bracket. The analytics developed for sport have wide application outside of sport as well. Overall, the development of complex analytics in sport is a cultural validation point for the importance of HPC.

High Performance Data Analytics: Constructing Complex Analytic Workflows in an HPC Enclave

William Leinberger (General Dynamics)

The objectives of High Performance Data Analytic (HPDA) systems are to maintain the productivity of Open Source Big Data frameworks (e.g. Hadoop, Accumulo, Spark) while improving their performance through the use of High Performance Computing (HPC) systems technology. We describe our efforts to define an HPDA Workflow Architecture. This architecture is enabled by three key components. First, a visual workflow tool is required, to provide a parameterizeable and repeatable interface for multi-step analytics. Second, an extensible tool is needed to dynamically deploy a heterogeneous suite of open source frameworks to HPC resources in a just-in-time fashion. Finally, a memory-centric data grid is needed to maximize the performance of data transfers between analytic steps in the workflow. We conclude with a description of a prototype implementation of the workflow architecture, and show how it enables the construction of a complex multi-step graph analytic solution to a real-world Social Network Analysis problem.

Software for HPC

Chair: Dane Skow (University of Wyoming)

3:30pm-5pm

Room: 12AB

SYCL for OpenCL: Accelerate High-Performance Software on GPUs, CPUs, FPGAs and Other Processors Using Open Standards in C++

Andrew Richards (Codeplay)

SYCL is the royalty-free, cross-platform abstraction layer that enables the development of applications and frameworks that build on the underlying concepts, portability and efficiency of OpenCL™, while adding the ease-of-use and flexibility of C++. For example, SYCL can provide single source development where C++ template functions can contain both host and device code to construct complex algorithms that use OpenCL acceleration – and then enable re-use of those templates throughout the source code of an application to operate on different types of data.

This talk will present SYCL as an open standard as well as demonstrate implementations of SYCL and show how using C++ and open standards can deliver huge performance, quickly and easily across a range of hardware platforms.

Strangely Parallel ... New, Old Parallel Hardware, Architecture and Languages

Michael C. Leventhal (Xilinx)

Massive, efficient, and manageable parallelism is the core of HPC and the next stage of computing evolution. Computing is increasingly dominated by intrinsically parallel interaction with the physical world using approaches such as deep learning. Incremental improvement to conventional compute architectures have dramatically slowed - some to say we are on the cusp of a silicon “apocalypse”. We are rethinking computing; disruptive hardware and architecture are increasingly appearing at SC. GPUs started the trend but are challenged in power efficiency, failing to address the communication cost of Von Neumann-style computing. The Micron Automata Processor used two old/new concepts, NFA and memory topology to propose a new computing approach. FPGAs are the old/new architecture now, with their intrinsic parallelism, reconfigurability and heterogeneous fabric offering tantalizing possibilities for energy-efficient high performance computation. With new architectures come new ways to express parallel computation, e.g., CUDA, ANML, OpenCL and domain-specific languages for neural nets.

Docker Containers - Future of Packaging Applications for Heterogeneous Platforms in HPC

Subbu Rama (Bitfusion.io, Inc.)

Docker containers wrap up a piece of software in a complete filesystem that contains everything it needs to run: code, runtime, system tools, system libraries – anything you can install on a server. This guarantees that it will always run the same, regardless of the environment it is running in. This talk will provide some background on linux containers, and its applicability to HPC, heterogeneous platforms in specific, and challenges in adoption. The talk will give an overview of packaging HPC applications as containers versus existing options and shed light on performance versus portability.

Thursday, November 19

Moving, Managing, and Storing Data

Chair: John Cazes (*The University of Texas at Austin*)

10:30am-12pm

Room: 12AB

The Role of Converged Architectures in Delivering IO for Exascale

James Coomer (*DataDirect Networks*)

In this talk Dr. Coomer will give a brief survey of emerging high performance data storage media, and then deliver a detailed view of emerging architectures which take advantage of them to deliver Petascale and even Exascale IO. Specific attention will be paid to next generation converged architectures and the how role of flash in the server, fabric and storage is likely to evolve over time.

Virtualized LNET Routing for Multi-Fabric Lustre Deployments

Daniel Rodwell (*NCI Australia*)

NCI has installed 30PB of high-performance storage, presented as multiple Lustre filesystems across the facility, which includes the Fujitsu Primergy 57,000 core (Intel Sandy-Bridge) HPC system known as Raijin.

To provide world-class research storage capabilities that are both scalable and reliable, NCI has architected an InfiniBand fabric design that is highly performing and resilient. As part of this architecture, a cost-to-performance effective method of scaling LNET routing was required to link the storage and HPC systems.

This presentation will discuss the architecture of NCI's InfiniBand fabrics, the implementation of Virtualized LNET routing for presenting NCI's Lustre filesystems across the facility, and a performance comparison with traditional bare-metal LNET routing approaches.

Building a High IOPS Flash Array: A Software-Defined Approach

Weafon Tsao (*AcceleStor, Inc.*)

Dr. Weafon Tsao is the Vice President of Research and Development Division, AcceleStor, Inc., a startup company devoted to deliver the true performance of flash-based storage solutions with a software-defined approach. Dr. Tsao's primary responsibility is to drive the cutting-edge developments of software-defined storage. He has more than 17 years of work and research experiences in the telecommunication, network

and storage territory. The presentation will introduce AcceleStor's exclusive software-defined approach and technology for performance enhancement of all-flash arrays. Dr. Tsao will talk about the current myth of flash arrays, the root of challenge and the solution to significantly improve the IO performance of all-flash arrays.

HPC Impact Showcase

Tuesday, November 17

Chair: David Martin (*Argonne National Laboratory*)

1pm-3:30pm

Room: 12AB

System Level Large Eddy Simulation of a Large Gas Turbine Combustor

Dheeraj Kapilavai (*General Electric*)

Combustion in large industry gas turbine consists of complicated physics such as turbulence, multi-specie mixing, chemical reaction, and heat transfer. The combustion process is highly turbulent and unsteady due to the interaction between the turbulence and the combustion. The better understanding of the unsteady combustion will help to reduce the emission and enhance the durability of the gas turbine. At the exit of the combustor, the first stage nozzle is critical for the durability of the turbine. The cooling and hence the heat transfer in the first stage nozzle is of great importance. In order to predict the heat transfer correctly, the upstream conditions form the combustor is needed. In the traditional approach, the 2D exit temperature profile is derived from combustor model and used as inlet condition for the nozzle simulation. However, the combustor exit flow is unsteady and 3 dimensional. It is desired to model the combustor and first stage nozzle together. In this study, a complete combustor can and nozzles of GE 9HA gas turbine are modeled using large eddy simulation (LES) method. The model reveals the intricate details of the turbulent flame structures. The unsteady flow field and the upstream conditions of the nozzles are studied too.

Realizing Actionable Computer Security Results Through HPC-Powered Analysis

Eric Dull (*Deloitte and Touche, LLP*)

Deloitte is realizing significant, actionable results in cyber analysis for clients by utilizing a cyber analytics infrastructure built on High-Performance and commodity computing architectures. This infrastructure is allowing Deloitte to rapidly develop "big-data" analytic workflows by combining many large data sets that address clients' analytic challenges in the areas of cyber reconnaissance, cyber hunt, and supply chain analysis. These workflows can then be executed efficiently enough to delivery actionable results to clients.

This presentation will include Deloitte's cyber analytic architecture, outline data flows, highlight useful algorithms and data transforms, and discuss two case studies, highlighting analytic successes and client outcomes.

Use of HPC Throughout the Product Cycle

Mohamad S. El-Zein (John Deere)

HPC has the reputation of solving problems related to climate prediction, aerospace and defense, oil discovery, etc. However, the use of HPC in everyday product design from inception through the supply chain is rarely highlighted due to its scarcity in terms of "HPC". A scan through areas of use of HPC through the supply chain of product design will be highlighted and discussed.

Wednesday, November 18

HPC Impact Showcase

Chair: Christy Adkinson (Cray Inc.)

1pm-3:30pm

Room: 12AB

Establishing an Enterprise HPC Service at Boeing

Joerg Gablonsky (Boeing),

Joerg Gablonsky, Chair of the Enterprise HPC Council, will present on establishing an enterprise HPC service at the Boeing company.

Discrete Element Method (DEM) Modeling on HPC Machines at Procter & Gamble

Chris Stoltz (Procter & Gamble)

This talk will focus on the use of Discrete Element Method (DEM) simulations at P&G for the simulation of granular materials, and how we have moved towards the use of open-source software in an HPC environment. We will present a broad overview of DEM, including a discussion of some of the capabilities and limitations inherent in conducting DEM simulations, and then discuss how we have taken advantage of our HPC environment to dramatically improve our capability for simulation in this area. We also include some benchmark studies we have carried out to evaluate the performance and scalability of a variety of simulation types.

An HPC-Based Computational Epidemiology Infrastructure for Ebola Response

Keith Bisset (Virginia Bioinformatics Institute)

The 2014-2015 Ebola outbreak in Western Africa has illuminated the significant threat posed by infectious diseases to human lives and society. As of September 2015, over 12,000 individuals have been confirmed infected, with another 5000 suspected cases, and almost 6500 deaths. Applying interventions in a resource-poor area during an ongoing epidemic is not easily done and the level of success is uncertain. Modeling disease outbreaks can therefore be helpful by providing epidemic forecasts that explain the complex dynamics of infectious diseases. Simulation and modeling can predict the likely impact of possible interventions before they are implemented. As a result, policy makers and health care workers are provided with guidance and support. High performance computing, big-data analytics, algorithms, visualizations and distributed systems are all important while developing scalable tools and decision support systems. Real-time computational epidemiology incorporating currently available ground information is critical to support decisions in the field. A state-of-the-art environment can deliver detailed situational awareness and decision support for an epidemic such as Ebola. NDSSL is supporting several federal agencies, including the DoD and NIH, by supplying them analytic results, enabling government decision makers to make more informed scientifically backed decisions and policies. The Infectious Disease Modeling Environment was quickly retargeted to Ebola, enabling us to: infer disease parameters in the very early phase of the epidemic, produce weekly forecasts of disease incidence and mortality, understand the chances of Ebola spreading in continental US, provide input on the location of Ebola.

Thursday, November 19

HPC Impact Showcase

Chair: David Martin (Argonne National Laboratory)

1pm-2:40pm

Room: 12AB

HPC Transforming Approaches to Alzheimer's and Parkinson's

Joseph Lombardo (University of Nevada, Las Vegas)

The University of Nevada Las Vegas and the Cleveland Clinic Lou Ruvo Center for Brain Health have been awarded an \$11.1 million federal grant from the National Institutes of Health and National Institute of General Medical Sciences to advance the understanding of Alzheimer's and Parkinson's diseases. In this session, we will present how UNLV's National Supercomputing Center will be a critical component to this research which will combine brain imaging with neuropsychological and behavioral studies to increase our knowledge of dementia-related and age-associated degenerative disorders.

HPC at BP: Enabling Seismic Imaging and Rock Physics Research

Keith Gray (BP)

Seismic Imaging and rock physics are key research focus areas for BP's Upstream Business and high-performance computing (HPC) is critical to enable the research breakthroughs required for our successful exploration and development programs.

Since 2013, our team has completed construction and moved into a new computing facility. We have delivered a 3.8-petaflop system. We have been able to strengthen our Software Development and Computational Science Team. The presentation will review the strategy for our HPC team, our history, current capabilities, and near-term plans. Also discussed will be the business value delivered by our seismic imaging research.

Panels

Tuesday, November 17

Post Moore's Law Computing: Digital versus Neuromorphic versus Quantum

1:30pm-3pm

Room: 16AB

Moderator: George Michalogiannakis (Lawrence Berkeley National Laboratory)

Panelists: John Shalf (Lawrence Berkeley National Laboratory), Bob Lucas (University of Southern California), Jun Sawada (IBM Corporation), Mattias Troyer (ETH Zurich), David Donofrio (Lawrence Berkeley National Laboratory), Shekhar Bokhar (Intel Corporation)

The end of Moore's Law scaling has sparked research into preserving performance scaling through alternative computational models. This has sparked a debate for the future of computing. Currently, the future of computing is expected to include a mix of quantum, neuromorphic, and digital computing. However, a range of questions remain unanswered for each option. For example, what problems each approach is most efficient for remains to be determined, and so are issues such as manufacturability, long-term potential, inherent drawbacks, programming, and many others. Can neuromorphic or quantum ever replace digital computing? Can we find alternative CMOS technologies and clever architectures to preserve digital computing performance scaling? What is the upper limit of CMOS? This is a critical debate for a wide audience, because solving many of tomorrow's problems requires a reasonable expectation of what tomorrow looks like.

Future of Memory Technology for Exascale and Beyond III

3:30pm-5pm

Room: 16AB

Moderator: Richard Murphy (Micron Technology, Inc.)

Panelists: Shekhar Borkar (Intel Corporation), Bill Dally (NVIDIA Corporation), Wendy Elasser (ARM Ltd.), Mike Ignatowski (Advanced Micro Devices, Inc.), Doug Joseph (IBM Corporation), Peter Kogge (University of Notre Dame), Steve Wallach (Micron Technology, Inc.)

Memory technology is in the midst of profound change as we move into the exascale era. Early analysis, including the DARPA UHPC Exascale Report correctly identified the fundamental technology problem as one of enabling low-energy data movement throughout the system. However, the end of Dennard Scaling and the corresponding impact on Moore's Law has begun a fundamental transition in the relationship between the processor and memory system. The lag in the increase in the number of cores compared to what Moore's Law would provide has proven a harbinger of the trend towards memory systems performance dominating compute capability.

Wednesday, November 18

Supercomputing and Big Data: From Collision to Convergence

10:30am-12pm

Room: 16AB

Moderator: George O. Strawn (*Networking and Information Technology Research and Development National Coordination Office*)

Panelists: David Bader (*Georgia Institute of Technology*), Ian Foster (*University of Chicago*), Bruce Hendrickson (*Sandia National Laboratories*), Randy Bryant (*Executive Office of the President, Office of Science and Technology Policy*), George Biros (*The University of Texas at Austin*), Andrew W. Moore (*Carnegie Mellon University*)

As data intensive science emerges, the need for HPC to converge capacity and capabilities with Big Data becomes more apparent and urgent. Capacity requirements have stemmed from science data processing and the creation of large scale data products (e.g., earth observations, Large Hadron Collider, square-kilometer array antenna) and simulation model output (e.g., flight mission plans, weather and climate models). Capacity growth is further amplified by the need for more rapidly ingesting, analyzing, and visualizing voluminous data to improve understanding of known physical processes, discover new phenomena, and compare results. (How does HPC need to change in order to meet these Big Data needs? What can HPC and Big Data communities learn from each other? What impact will this have on conventional workflows, architectures, and tools?)

An invited international panel of experts will examine these disruptive technologies and consider their long-term impacts and research directions.

Mentoring Undergraduates Through Competition

1:30pm-3pm

Room: 16AB

Moderator: Brent Gorda (*Intel Corporation*)

Panelists: Jerry Chou (*Tsinghua University*), Rebecca Hartman-Baker (*Lawrence Berkeley National Laboratory*), Doug Smith (*University of Colorado Boulder*), Xuanhua Shi (*Huazhong University of Science and Technology*), Stephen Lien Harrell (*Purdue University*)

The next generation of HPC talent will face significant challenges to create software ecosystems and optimally use the next generation of HPC systems. The rapid advances in HPC make it difficult for academic institutions to keep pace. The Student Cluster Competition, now in its ninth year, was created

to address this issue by immersing students into all aspects of HPC. This panel will examine the impact of the Student Cluster Competition on the students and schools that have participated. Representatives from five institutions from around the world will talk about their experiences with the Student Cluster Competition with respect to their students' career paths, integration with curriculum and academic HPC computing centers. The panel will further discuss whether "extracurricular" activities, such as the Student Cluster Competition, provide sufficient return on investment and what activities could change or replace the competition to meet these goals more effectively.

Programming Models for Parallel Architectures and Requirements for Pre-Exascale

3:30pm-5pm

Room: 16AB

Moderator: Fernanda Foertter (*Oak Ridge National Laboratory*)

Panelists: Barbara Chapman (*University of Houston*), Steve Oberlin (*NVIDIA Corporation*), Satoshi Matsuoka (*Tokyo Institute of Technology*), Jack Wells (*Oak Ridge National Laboratory*), Si Hammond (*Sandia National Laboratories*)

Relying on domain scientists to provide programmer intervention to develop applications to emerging exascale platforms is a real challenge. A scientist prefers to express mathematics of the science, not describe the parallelism of the implementing algorithms. Do we expect too much of the scientist to code for high parallel performance given the immense capabilities of the platform? This ignores that the scientist may have a mandate to code for a new architecture and yet preserve portability in their code. This panel will bring together user experience, programming models, and architecture experts to discuss the pressing needs in finding the path forward to port scientific codes to such a platform. We hope to discuss the evolving programming stack and application-level requirements and address the hierarchical nature of large systems in terms of different cores, memory levels, power consumption and the pragmatic advances of near term technology.

Thursday, November 19

Asynchronous Many-Task Programming Models for Next Generation Platforms**10:30am-12pm****Room: 16AB****Moderator:** Robert Clay (Sandia National Laboratories)**Panelists:** Alex Aiken (Stanford University), Martin Berzins (University of Utah), Matthew Bettencourt (Sandia National Laboratories), Laxmikant Kale (University of Illinois at Urbana-Champaign), Timothy Mattson (Intel Corporation), Lawrence Rauchwerger (Texas A&M University), Vivek Sarkar (Rice University), Thomas Sterling (Indiana University), Jeremiah Wilke (Sandia National Laboratories)

Next generation platform architectures will require us to fundamentally rethink our programming models and environments due to a combination of factors including extreme parallelism, data locality issues, and resilience. As seen in the computational sciences community, asynchronous many-task (AMT) programming models and runtime systems are emerging as a leading new paradigm. While there are some overarching similarities between existing AMT systems, the community lacks consistent 1) terminology to describe runtime components, 2) application- and component-level interfaces, and 3) requirements for the lower level runtime and system software stacks. This panel will engage a group of community experts in a lively discussion on status and ideas to establish best practices in light of requirements such as performance portability, scalability, resilience, and interoperability. Additionally, we will consider the challenges of user-adoption, with a focus on the issue of productivity, which is critical given the application code rewrite required to adopt this approach.

Toward an Open Software Stack for Exascale Computing**1:30pm-3pm****Room: 16AB****Moderator:** Nicolás Erdödy (Open Parallel Ltd.)**Panelists:** Pete Beckman (Argonne National Laboratory), Chris Broekema (Netherlands Institute for Radio Astronomy), Jack Dongarra (University of Tennessee), John Gustafson (Ceranovo, Inc.), Thomas Sterling (Indiana University), Robert Wisniewski (Intel Corporation)

The panel will discuss what an open software stack should contain, what would make it feasible, and what is not looking possible at the moment. The discussion is inspired by the fact that “this time, we have time” before the hardware actually reaches the market after 2020, so we can work on a software stack accordingly.

We will cover questions such as: Which would be the software development costs? What industries will migrate first? Would a killer app accelerate this process? Do we focus on algorithms to save power? How heterogeneous would/should “your” exascale system be? Is there a role for co-design toward exascale? Is the Square Kilometre Array (SKA) project an example to follow? Would cloud computing be possible for exascale? Who will “own” the exascale era?

Procuring Supercomputers: Best Practices and Lessons Learned**3:30pm-5pm****Room: 16AB****Moderator:** Bilel Hadri (King Abdullah University of Science and Technology)**Panelists:** Katie Antypas (National Energy Research Scientific Computing Center), Bill Kramer (University of Illinois at Urbana-Champaign), Satoshi Matsuoka (Tokyo Institute of Technology), Greg Newby (Compute Canada), Owen Thomas (Red Oak Consulting)

Procuring HPC systems is a challenging process of acquiring the most suitable machine under technical and financial constraints aiming at maximizing the benefits to the users’ applications and minimizing the risks during its lifetime. In this panel, HPC leaders will discuss and debate on key requirements and lessons learned for successful procurements of supercomputers. How do we define the requirements of the system? Is it to acquire a system for maximizing the capacity and capability, assessing new/future technologies, deliver a system designed for specific applications or provide an all-purpose solution to a broad range of applications? Is the system just a status symbol or must it do useful work? This panel will give the audience an opportunity to ask questions of panelists who are involved in the procurement of leadership-class supercomputers and capturing lessons learned and turning that hindsight into best practices to procure and the most suitable HPC system.

Friday, November 20

Return of HPC Survivor: Outwit, Outlast, Outcompute

8:30am-10am

Room: 17AB

Moderator: Cherri Pancake (Oregon State University)

Panelists: Robin Goldstone (Lawrence Livermore National Laboratory), Steve Hammond (National Renewable Energy Laboratory), Jennifer M. Schopf (Indiana University), John E. West (The University of Texas at Austin)

Back by popular demand, this panel brings together HPC experts to compete for the honor of “HPC Survivor 2015.” Following up on the popular Xtreme Architectures (2004), Xtreme Programming (2005), Xtreme Storage (2007)), Build Me an Exascale (2010), and Does HPC Really Matter? (2014) competitions, the theme for this year is “HPC Transformed: How to Reduce/Recycle/Reuse Your Outdated HPC System.”

The contest is a series of “rounds,” each posing a specific question about system characteristics and how that affects its transformation to new and exciting uses. After contestants answer, a distinguished commentator furnishes additional wisdom to help guide the audience. At the end of each round, the audience votes (applause, boos, etc.) to eliminate a contestant. The last contestant left wins.

While delivered in a light-hearted fashion, the panel pushes the boundaries of how HPC can/should affect society in terms of impact, relevancy, and ROI.

HPC Transforms DoD, DOE, and Industrial Product Design, Development and Acquisition

8:30am-10am

Room: 16AB

Moderator: Loren Miller (DataMetric Innovations, LLC)

Panelists: Christopher Atwood (U.S. Department of Defense High Performance Computing Modernization Program and CREATE Program), Col Keith Bearden (United States Air Force), Douglas Kothe (Oak Ridge National Laboratory), Edward Kraft (United States Air Force), Lt Col Andrew Lofthouse (United States Air Force Academy), Douglass Post (U.S. Department of Defense High Performance Computing Modernization Program and CREATE Program)

Supercomputing has been shown to enable massive reductions in product development time, significant improvements in product capability, greater design innovation in new products, and effective systems engineering implementations. Our panelists will share their intimate knowledge of the various methods and practices by which these results have been

achieved in the U.S. Departments of Defense & Energy, and in industry. Topics will include the Digital Thread & Twin of Air Force acquisition; the development and deployment of physics-based engineering analysis and design software for military aircraft, ships, ground vehicles, and antennas; high fidelity predictive simulation of challenging nuclear reactor conditions; accessibility in the era of hacking and exfiltration; STEM education using HPC; and cultural barriers to organizational adoption of HPC-based product development. Audience questions and contributions to the list of key enablers and pitfalls for the implementation of HPC-based product development within both government and industry will be encouraged and discussed.

HPC and the Public Cloud

10:30am-12pm

Room: 17AB

Moderator: Kevin D. Kissell (Google)

Panelists: Jeff Baxter (Microsoft Corporation), Shane Canon (Lawrence Berkeley National Laboratory), Brian Cartwright (MetLife Insurance Company), Steve Feldman (CD-adapco), Bill Kramer (University of Illinois at Urbana-Champaign), Kanai Pathak (Schlumberger Limited), Steve Scott (Cray Inc.)

Where high-performance computing collides with cloud computing, just about the only point where most interested and informed parties agree is that the overlap is incomplete, complex, and dynamic. We are bringing together stakeholders on all sides of the issue to express and debate their points of view on questions such as:

Which HPC workloads should be running in the public Cloud? Which should not? How will Cloud economics affect the choices of algorithms and tools? How does Cloud computing impact computational science? Is there a line to be drawn between “Big Data” and “HPC”? If so, where? Will Cloud HPC encourage or discourage innovation in HPC hardware and software? What is it about HPC that the Cloud providers “don’t get”?

In Situ Methods: Hype or Necessity?**10:30am-12pm****Room: 16AB**

Moderator: *Wes Bethel (Lawrence Berkeley National Laboratory)*

Panelists: *Patrick O’Leary (Kitware, Inc.), John Clyne (National Center for Atmospheric Research), Venkat Vishwanath (Argonne National Laboratory), Jacqueline Chen (Sandia National Laboratories)*

Due to the widening gap between the FLOP and I/O capacity of HPC platforms, it is increasingly impractical for computer simulations to save full-resolution computations to disk for subsequent analysis. “In situ” methods offer hope for managing this increasingly acute problem by performing as much analysis, visualization, and related processing while data is still resident in memory. While in situ methods are not new, they are presently the subject of much active R&D, though as yet are not widespread in deployment or use.

This panel examines different aspects of in situ methods, with an eye toward increasing awareness of the current state of this technology, how it is used in practice, and challenges facing widespread deployment and use. The panel will also explore the issue of whether in situ methods are really needed or useful in the first place, and invites discussion and viewpoints from the SC community.

Tuesday, November 17

Applications - Material Science

10:30am-12pm

Room: 18AB

Massively Parallel Phase-Field Simulations for Ternary Eutectic Directional Solidification

Authors: Martin Bauer (FAU Erlangen Nuremberg); Johannes Hötzer, Marcus Jainta (Karlsruhe University of Applied Sciences); Philipp Steinmetz, Marco Berghoff (Karlsruhe Institute of Technology); Florian Schornbaum, Christian Godenschwager, Harald Köstler (FAU Erlangen Nuremberg); Britta Nestler (Karlsruhe Institute of Technology), Ulrich Rude (FAU Erlangen Nuremberg)

Microstructures forming during ternary eutectic directional solidification processes have significant influence on the macroscopic mechanical properties of metal alloys. For a realistic simulation, we use the well-established thermodynamically consistent phase-field method and improve it with a new grand potential formulation to couple the concentration evolution. This extension is very compute intensive due to a temperature dependent diffusive concentration. We significantly extend previous simulations that have used simpler phase-field models or were performed on smaller domain sizes. The new method has been implemented within the massively parallel HPC framework waLberla that is designed to exploit current supercomputers efficiently. We apply various optimization techniques, including buffering techniques, explicit SIMD kernel vectorization, and communication hiding. Simulations utilizing up to 262,144 cores have been run on three different supercomputing architectures and weak scalability results are shown. Additionally, a hierarchical, mesh-based data reduction strategy is developed to keep the I/O problem manageable at scale.

Award: Best Paper Finalist

Parallel Implementation and Performance Optimization of the Configuration-Interaction Method

Authors: Hongzhang Shan, Samuel Williams (Lawrence Berkeley National Laboratory); Calvin Johnson (San Diego State University), Kenneth McElvain (University of California, Berkeley), W. Erich Ormand (Lawrence Livermore National Laboratory)

The configuration-interaction (CI) method, long a popular approach to describe quantum many-body systems, is often cast as a very large sparse matrix eigenpair problem with matrices whose dimension can exceed one billion. Such formulations place high demands on memory capacity and memory bandwidth. In this paper, we describe an efficient, scalable implementation, BIGSTICK, which, by factorizing both the basis and the interaction into two levels, can reconstruct the nonzero matrix elements on the fly and reduce the memory requirements by one or two orders of magnitude allowing researchers to trade reduced resources for increased computational time. We optimize BIGSTICK on two leading HPC platforms. Specifically, we not only develop an empirically-driven load balancing strategy that can evenly distribute the matrix-vector multiplication across 256K threads, we also developed techniques that improve the performance of the Lanczos reorthogonalization. Combined, these optimizations improved performance by 1.3-8x depending on platform and configuration.

Efficient Implementation of Quantum Materials Simulations on Distributed CPU-GPU Systems

Authors: Raffaele Solcà (ETH Zurich), Anton Kozhevnikov (Swiss National Supercomputing Center); Azzam Haidar, Stanimire Tomov, Jack Dongarra (University of Tennessee, Knoxville); Thomas C. Schulthess (Swiss National Supercomputing Center)

We present a scalable implementation of the Linearized Augmented Plane Wave method for distributed memory systems, which relies on an efficient distributed, block-cyclic setup of the Hamiltonian and overlap matrices and allows us to turn around highly accurate 1000+ atom all-electron quantum materials simulations on clusters with a few hundred nodes.

The implementation runs efficiently on standard multi-core CPU nodes, as well as hybrid CPU-GPU nodes. Key for the latter is a novel algorithm to solve the generalized eigenvalue problem for dense, complex Hermitian matrices on distributed hybrid CPU-GPU systems. Performance test for Li-intercalated CoO₂ supercells containing 1501 atoms demonstrate that high accuracy, transferable quantum simulations can now be used in throughput materials search problems. A systematic comparison between our new hybrid solver and the ELPA2 library shows that the hybrid CPU-GPU architecture is considerably more energy efficient than traditional multi-core CPU-only systems for such complex applications.

Award: Best Paper Finalist

Cache and Memory Subsystems

10:30am-12pm

Room: 19AB

Runtime-Driven Shared Last-Level Cache Management for Task-Parallel Programs

Authors: Abhisek Pan, Vijay S. Pai (Purdue University)

Task-parallel programming models with input annotation-based concurrency extraction at runtime present a promising paradigm for programming multicore processors. Through management of dependencies, task assignments, and orchestration, these models markedly simplify the programming effort for parallelization while exposing higher levels of concurrency.

In this paper we show that for multicores with a shared last-level cache (LLC), the concurrency extraction framework can be used to improve the shared LLC performance. Based on the input annotations for future tasks, the runtime instructs the hardware to prioritize data blocks with future reuse while evicting blocks with no future reuse. These instructions allow the hardware to preserve all the blocks for at least some of the future tasks and evict dead blocks. This leads to a considerable improvement in cache efficiency over what is achieved by hardware-only replacement policies, which can replace blocks for all future tasks resulting in poor hit-rates for all future tasks.

Frugal ECC: Efficient and Versatile Memory Error Protection through Fine-Grained Compression

Authors: Jung-rae Kim, Michael Sullivan, Seong-Lyong Gong, Mattan Erez (The University of Texas at Austin)

Because main memory is vulnerable to errors and failures, large-scale systems and critical servers utilize error checking and correcting (ECC) mechanisms to meet their reliability requirements. We propose a novel mechanism, Frugal ECC (FECC), that combines ECC with fine-grained compression to

provide versatile protection that can be both stronger and lower overhead than current schemes, without sacrificing performance. FECC compresses main memory at cache-block granularity, using any leftover space to store ECC information. Compressed data and its ECC information are then frequently read with a single access even without redundant memory chips; insufficiently compressed blocks require additional storage and accesses. As examples, we present chipkill-correct ECCs on a non-ECC DIMM with $\times 4$ chips and the first true chipkill-correct ECC for $\times 8$ chips using a ECC DIMM. FECC relies on a new Coverage-oriented-Compression that we developed specifically for the modest compression needs of ECC and for floating-point data.

Automatic Sharing Classification and Timely Push for Cache-Coherent Systems

Authors: Malek Musleh, Vijay S. Pai (Purdue University)

This paper proposes and evaluates Sharing/Timing Adaptive Push (STAP), a dynamic scheme for preemptively sending data from producers to consumers to minimize critical-path communication latency. STAP uses small hardware buffers to dynamically detect sharing patterns and timing requirements. The scheme applies to both intra-node and inter-socket directory-based shared memory networks. We integrate STAP into a MOESI cache-coherence (prefetching-enabled) protocol using heuristics to detect different data sharing patterns, including broadcasts, producer/consumer, and migratory-data sharing. Using 15 benchmarks from the PARSEC and SPLASH-2 suites, we show that our scheme significantly reduces communication latency in NUMA systems and achieves an average of 9% performance improvement, with at most 3% on-chip storage overhead.

Data Clustering

10:30am-11:30am

Room: 18CD

BD-CATS: Big Data Clustering at Trillion Particle Scale

Authors: Md. Mostafa Ali Patwary (Intel Corporation), Suren Byna (Lawrence Berkeley National Laboratory), Nadathur Rajagopalan Satish, Narayanan Sundaram (Intel Corporation), Zarija Lukic (Lawrence Berkeley National Laboratory), Vadim Roytershteyn (Space Science Institute), Michael J. Anderson (Intel Corporation), Yushu Yao, Mr Prabhat (Lawrence Berkeley National Laboratory), Pradeep Dubey (Intel Corporation)

Modern cosmology and plasma-physics codes are now capable of simulating trillions of particles on petascale-systems. Each timestep output from such simulations is on the order of 10s of TBs. Summarizing and analyzing raw particle data is challenging, and scientists often focus on density structures,

whether in the real 3D space, or a high-dimensional phase space. In this work, we develop a highly scalable version of the clustering algorithm DBSCAN and apply it to the largest datasets produced by state-of-the-art codes. Our system, called BD-CATS, is the first one capable of performing end-to-end analysis at trillion particle scale. We show analysis of 1.4 trillion particles from a plasma-physics simulation, and a $10,240^3$ particle cosmological simulation, utilizing $\sim 100,000$ cores in 30 minutes. BD-CATS is helping infer mechanisms behind particle acceleration in plasma-physics and holds promise for qualitatively superior clustering in cosmology. Both of these results were previously intractable at the trillion-particle scale.

Performance Optimization for the K-Nearest Neighbors Kernel on x86 Architectures

Authors: Chenhan D. Yu, Jianyu Huang, Woody Austin, Bo Xiao, George Biros (The University of Texas at Austin)

Nearest neighbor search is a cornerstone problem in computational geometry, non-parametric statistics, and machine learning. For N points, exhaustive search requires quadratic work, but many fast algorithms reduce the complexity for exact and approximate searches. The common kernel (kNN kernel) in all these algorithms solves many small-size problems exactly using exhaustive search. We propose an efficient implementation and performance analysis for the kNN kernel on x86 architectures. By fusing the distance calculation with the neighbor selection, we are able to utilize memory throughput. We present an analysis of the algorithm and explain parameter selection. We perform an experimental study varying the size of the problem, the dimension of the dataset, and the number of nearest neighbors. Overall we observe significant speedups. For example, when searching for 16 neighbors in a point dataset with 1.6 million points in 64 dimensions, our kernel is over 4 times faster than existing methods.

Applications - Biophysics and Genomics

1:30pm-3pm

Room: 18AB

HipMer: An Extreme-Scale De Novo Genome Assembler

Authors: Evangelos Georganas (University of California, Berkeley), Aydin Buluc (Lawrence Berkeley National Laboratory), Jarrod Chapman (Joint Genome Institute), Steven Hofmeyr (Lawrence Berkeley National Laboratory), Chaitanya Aluru (University of California, Berkeley), Rob Egan (Joint Genome Institute), Leonid Olikier (Lawrence Berkeley National Laboratory), Daniel Rokhsar (University of California, Berkeley), Katherine Yelick (Lawrence Berkeley National Laboratory)

De novo whole genome assembly reconstructs genomic sequence from short, overlapping, and potentially erroneous DNA segments and is one of the most important computations

in modern genomics. This work presents the first high-quality end-to-end de novo assembler designed for extreme scale analysis, via efficient parallelization of the Meraculous code. First, we significantly improve scalability of parallel k-mer analysis for complex repetitive genomes that exhibit skewed frequency distributions. Next, we optimize the traversal of the de Bruijn graph of k-mers by employing a novel communication-avoiding parallel algorithm in a variety of use-case scenarios. Finally, we parallelize the Meraculous scaffolding modules by leveraging the one-sided communication capabilities of the Unified Parallel C while effectively mitigating load imbalance. Large-scale results on a Cray XC30 using grand-challenge genomes, demonstrate efficient performance and scalability on thousands of cores. Overall, our pipeline accelerates Meraculous performance by orders of magnitude, creating unprecedented capability for extreme-scale genomic analysis.

A Parallel Connectivity Algorithm for de Bruijn Graphs in Metagenomic Applications

Authors: Patrick Flick, Chirag Jain, Tony Pan, Srinivas Aluru (Georgia Institute of Technology)

Dramatic advances in DNA sequencing technology have made it possible to study microbial environments by direct sequencing of environmental DNA samples. Yet, due to huge volume and high data complexity, current de novo assemblers cannot handle large metagenomic datasets or will fail to perform assembly with acceptable quality. This paper presents the first parallel solution for decomposing the metagenomic assembly problem without compromising post-assembly quality. We transform this problem into that of finding weakly connected components in the de Bruijn graph. We propose a novel distributed memory algorithm to identify the connected subgraphs, and present strategies to minimize the communication volume. We demonstrate scalability of our algorithm on soil metagenome dataset with 1.8 billion reads. Our approach achieves a runtime of 22 minutes using 1280 Intel Xeon cores for 421 GB uncompressed FASTQ dataset. Moreover, our solution is generalizable to finding connected components in arbitrary undirected graphs.

Parallel Distributed Memory Construction of Suffix and Longest Common Prefix Arrays

Authors: Patrick Flick, Srinivas Aluru (Georgia Institute of Technology)

Suffix arrays and trees are fundamental string data structures of importance to many applications in computational biology. Consequently, their parallel construction is an actively studied problem. To date, algorithms with best practical performance lack efficient worst-case run-time guarantees, and vice versa. In addition, much of the recent work targeted low core count, shared memory parallelization. In this paper, we present paral-

lel algorithms for distributed memory construction of suffix arrays and longest common prefix (LCP) arrays that simultaneously achieve good worst-case run-time bounds and superior practical performance. Our algorithms run in $O(T_{\text{sort}}(n,p) \log(n))$ worst-case time where $T_{\text{sort}}(n,p)$ is the run-time of parallel sorting. We present several algorithm engineering techniques that improve performance in practice. We demonstrate the constructions of suffix and LCP arrays of the human genome in less than 8 seconds on 1,024 Intel Xeon cores, reaching speedups of over 110x compared to the best sequential suffix array construction implementation `divsufsort`.

Award: Best Student Paper Finalist

GPU Memory Management

1:30pm-3pm

Room: 19AB

Adaptive and Transparent Cache Bypassing for GPUs

Authors: Ang Li, Gert-Jan van den Braak (Eindhoven University of Technology); Akash Kumar (Technische Universität Dresden), Henk Corporaal (Eindhoven University of Technology)

GPUs have recently emerged to be widely adopted for general-purpose applications. To capture on-chip locality for these applications, modern GPUs have integrated multi-level cache hierarchy in an attempt to reduce the amount and latency of the massive and sometimes irregular memory accesses. However, inferior performance is frequently attained due to the serious congestion in the caches resulting from the huge volume of concurrent threads. In this paper, we propose a novel compiler and runtime framework for adaptive and transparent cache bypassing on GPUs. It uses a simple yet effective approach to control the bypass degree so as to match the size of applications' runtime footprints. We validate the design on several GPUs from different generations using 16 cache sensitive applications. Experiment measurements show that our design can significantly improve overall performance (up to 2.16x on average). And, we summarize some optimization guidelines regarding GPU caches based on the experiment figures.

Award: Best Student Paper Finalist, Best Paper Finalist

ELF: Maximizing Memory-Level Parallelism for GPUs with Coordinated Warp and Fetch Scheduling

Authors: Jason Jong Kyu Park (University of Michigan), Yongjun Park (Hongik University), Scott Mahlke (University of Michigan)

Graphics processing units (GPUs) are increasingly utilized as throughput engines in the modern computer systems. GPUs rely on fast context switching between thousands of threads to hide long latency operations; however, they still stall due to the memory operations. To minimize the stalls, memory operations should be overlapped with other operations as much as possible to maximize memory-level parallelism (MLP). In this paper, we propose Earliest Load First (ELF) warp schedul-

ing, which maximizes the MLP by giving higher priority to the warps that have the fewest instructions to the next memory load. ELF utilizes the same warp priority for the fetch scheduling so that both are coordinated. We also show that ELF reveals its full benefits when there are fewer memory conflicts and fetch stalls. Evaluations show that ELF can improve the performance by 4.1% and achieve total improvement of 11.9% when used with other techniques over commonly-used greedy-then-oldest scheduling.

Memory Access Patterns: the Missing Piece of the Multi-GPU Puzzle

Authors: Tal Ben-Nun, Ely Levy, Amnon Barak, Eri Rubin (Hebrew University of Jerusalem)

With the increased popularity of multi-GPU nodes in modern HPC clusters, it is imperative to develop matching programming paradigms for their efficient utilization. In order to take advantage of the local GPUs and the low-latency high-throughput interconnects that link them, programmers need to meticulously adapt parallel applications with respect to load balancing, boundary conditions and device synchronization. This paper presents MAPS-Multi, an automatic multi-GPU partitioning framework that distributes the workload based on the underlying memory access patterns. The framework consists of host- and device-level APIs that allow programs to efficiently run on a variety of GPU and multi-GPU architectures. The framework implements several layers of code optimization, device abstraction, and automatic inference of inter-GPU memory exchanges. The paper demonstrates that the performance of MAPS-Multi achieves near-linear scaling on fundamental computational operations, as well as real-world applications in deep learning and multivariate analysis.

Scalable Storage Systems

1:30pm-3pm

Room: 18CD

AnalyzeThis: An Analysis Workflow-Aware Storage System

Authors: Hyogi Sim (Virginia Polytechnic Institute and State University), Youngjae Kim, Sudharshan S. Vazhkudai, Devesh Tiwari (Oak Ridge National Laboratory), Ali Anwar, Ali R. Butt (Virginia Polytechnic Institute and State University), Lavanya Ramakrishnan (Lawrence Berkeley National Laboratory)

The need for novel data analysis is urgent in the face of a data deluge from modern applications. Traditional approaches to data analysis incur significant data movement costs, moving data back and forth between the storage system and the processor. Emerging Active Flash devices enable processing on the flash, where the data already resides. An array of such Active Flash devices allows us to revisit how analysis workflows

interact with storage systems. By seamlessly blending together the flash storage and data analysis, we create an analysis workflow-aware storage system, AnalyzeThis. Our guiding principle is that analysis-awareness be deeply ingrained in each and every layer of the storage, elevating data analyses as first-class citizens, and transforming AnalyzeThis into a potent analytics-aware appliance. We implement the AnalyzeThis storage system atop an emulation platform of the Active Flash array. Our results indicate that AnalyzeThis is viable, expediting workflow execution, and minimizing data movement.

Mantle: A Programmable Metadata Load Balancer for the Ceph File System

Authors: Michael A. Sevilla, Noah Watkins, Carlos Maltzahn, Ike Nassi, Scott A. Brandt (University of California, Santa Cruz), Sage A. Weil, Greg Farnum (Red Hat) Sam Fineberg (Hewlett-Packard Development Company, L.P.)

Migrating resources is a useful tool for balancing load in a distributed system, but it is difficult to determine when to move resources, where to move resources, and how much of them to move. We look at resource migration for file system metadata and show how CephFS's dynamic subtree partitioning approach can exploit varying degrees of locality and balance because it can partition the namespace into variable sized units. Unfortunately, the current metadata balancer is complicated and difficult to control because it struggles to address many of the general resource migration challenges inherent to the metadata management problem. To help decouple policy from mechanism, we introduce a programmable storage system that lets the designer inject custom balancing logic. We show the flexibility and transparency of this approach by replicating the strategy of a state-of-the-art metadata balancer and conclude by comparing this strategy to other custom balancers on the same system.

HydraDB: A Resilient RDMA-Driven Key-Value Middleware for In-Memory Cluster Computing

Authors: Yandong Wang, Li Zhang, Jian Tan, Min Li (IBM Corporation), Yuqing Gao (Microsoft Corporation), Xavier Guerin (Tower Research Capital LLC), Xiaoqiao Meng (Pinterest, Inc.), Shicong Meng (Facebook)

In this paper, we describe our experiences and lessons learned from building an in-memory key-value middleware, called HydraDB. HydraDB synthesizes a collection of state-of-the-art techniques, including high-availability, RDMA, as well as multicore-awareness, etc, to deliver a high-throughput, low-latency access service in a reliable manner for cluster computing applications.

The uniqueness of HydraDB lies in its design commitment to exploit RDMA to comprehensively optimize various aspects of

a general-purpose key-value store, including latency-critical operations, read enhancement, and replications for high-availability service, etc. Meanwhile, HydraDB strives to efficiently utilize multicore systems to prevent data manipulation from curbing the performance of RDMA.

Many teams in our organization have adopted HydraDB to improve the execution of their cluster computing frameworks, including MapReduce, Sensemaking analytics and Call Record Processing. In addition, performance evaluation with a variety of YCSB workloads also shows that HydraDB substantially outperforms several existing in-memory key-value stores by an order of magnitude.

Applications - Folding, Imaging, and Proteins

3:30pm-5pm

Room: 18AB

Full Correlation Matrix Analysis of fMRI Data on Intel Xeon Phi Coprocessors

Authors: Yida Wang (Princeton University), Michael J. Anderson (Intel Corporation), Jonathan D. Cohen (Princeton University), Alexander Heinecke (Intel Corporation), Kai Li (Princeton University), Nadathur Satish (Intel Corporation), Narayanan Sundaram (Intel Corporation), Nicholas B. Turk-Browne (Princeton University), Theodore L. Willke (Intel Corporation)

Full correlation matrix analysis (FCMA) is an unbiased approach for exhaustively studying interactions among brain regions in functional magnetic resonance imaging (fMRI) data from human participants. In order to answer neuroscientific questions efficiently, we are developing a closed-loop analysis system with FCMA on a cluster of nodes with Intel Xeon Phi coprocessors. We have proposed several ideas which involve data driven algorithmic modification to improve the performance on the coprocessor. Our experiments with real datasets show that the optimized single-node code runs 5x-16x faster than the baseline implementation using the well-known Intel MKL and LibSVM libraries, and that the cluster implementation achieves near linear speedup on 5760 cores.

A Kernel-Independent FMM in General Dimensions

Authors: William B. March, Bo Xiao, Sameer Tharakan, Chenhan D. Yu, George Biros (The University of Texas at Austin)

We introduce a general-dimensional, kernel-independent, algebraic fast multipole method and apply it to kernel regression. The motivation for this work is the approximation of kernel matrices, which appear in mathematical physics, approximation theory, non-parametric statistics, and machine learning. Existing fast multipole methods are asymptotically optimal, but the underlying constants scale quite badly with

the ambient space dimension. We introduce a method that mitigates this shortcoming; it only requires kernel evaluations and scales well with the problem size, the number of processors, and the ambient dimension---as long as the intrinsic dimension of the dataset is small. We test the performance of our method on several synthetic datasets. As a highlight, our largest run was on an image dataset with 10 million points in 246 dimensions.

Engineering Inhibitory Proteins with InSiPS: The In-Silico Protein Synthesizer

Authors: Andrew Schoenrock, Daniel Burnside, Houman Mote-sharie, Alex Wong, Ashkan Golshani, Frank Dehne, James R. Green (Carleton University)

Engineered proteins are synthetic novel proteins (not found in nature) that are designed to fulfill a predetermined biological function. Such proteins can be used as molecular markers, inhibitory agents, or drugs. For example, a synthetic protein could bind to a critical protein of a pathogen, thereby inhibiting the function of the target protein and potentially reducing the impact of the pathogen. In this paper we present the In-Silico Protein Synthesizer (InSiPS), a massively parallel computational tool for the IBM Blue Gene/Q that is aimed at designing inhibitory proteins. More precisely, InSiPS designs proteins that are predicted to interact with a given target protein (and may inhibit the target's cellular functions) while leaving non-target proteins unaffected (to minimize side-effects). As proof-of-concepts, two InSiPS designed proteins have been synthesized in the lab and their inhibitory properties have been experimentally verified through wet-lab experimentation.

Graph Analytics on HPC Systems

3:30pm-5pm

Room: 19AB

Exploring Network Optimizations for Large-Scale Graph Analytics

Authors: Xinyu Que, Fabio Checconi, Fabrizio Petrini, Xing Liu, Daniele Buono (IBM Corporation)

Graph analytics are arguably one of the most demanding workloads for high-performance interconnection networks. The fine-grained, high-rate communication patterns expose the limits of the protocol stacks. The load and communication imbalance generate hard-to-predict hot-spots and require computational steering due to unpredictable data distributions. In this paper we present a lightweight communication library, implemented on the metal of BlueGene/Q and Power775/PERCS that we have used to support large-scale graph algorithms up to 96K processing nodes and 6 million threads. With this library we have explored several techniques, including computation in the network for special collective

communication patterns such as parallel prefix and wavefront algorithms. The experimental results show significant performance improvements when compared to equally optimized MPI implementations.

GossipMap: A Distributed Community Detection Algorithm for Billion-Edge Directed Graphs

Authors: Seung-Hee Bae, Bill Howe (University of Washington)

In this paper, we describe a new distributed community detection algorithm for billion-edge directed graphs that, unlike modularity-based methods, achieves cluster quality on par with the best-known algorithms in the literature. We show that a simple approximation to the best-known serial algorithm dramatically reduces computation and enables distributed evaluation yet incurs only a very small impact on cluster quality.

We present three main results. First, we show that the clustering produced by our scalable approximate algorithm compares favorably with prior results on small synthetic benchmarks and small real-world datasets (70 million edges). Second, we evaluate our algorithm on billion-edge directed graphs (a 1.5B edge social network graph, and a 3.7B edge web crawl) and show that the results exhibit the structural properties predicted by analysis of much smaller graphs from similar sources. Third, we show that our algorithm exhibits over 90% parallel efficiency on massive graphs in weak scaling experiments.

GraphReduce: Processing Large-Scale Graphs on Accelerator-Based Systems

Authors: Dipanjan Sengupta (Georgia Institute of Technology), Shuaiwen Leon Song (Pacific Northwest National Laboratory), Kapil Agarwal, Karsten Schwan (Georgia Institute of Technology)

Recent work on real-world graph analytics has sought to leverage the massive amount of parallelism offered by GPU devices, but challenges remain due to the inherent irregularity of graph algorithms and limitations in GPU-resident memory for storing large graphs. We present GraphReduce, a highly efficient and scalable GPU-based framework that operates on graphs that exceed the device's internal memory capacity. GraphReduce adopts a combination of edge- and vertex-centric implementations of the Gather-Apply-Scatter programming model and operates on multiple asynchronous GPU streams to fully exploit the high degrees of parallelism in GPUs with efficient graph data movement between the host and device. GraphReduce-based programming is performed via device functions that include gatherMap, gatherReduce, apply, and scatter, implemented by programmers for the graph algorithms they wish to realize. Extensive experimental evaluations for a wide variety of graph inputs and algorithms demonstrate

that GraphReduce significantly outperforms other competing out-of-memory approaches.

Award: Best Student Paper Finalist

MPI/Communication

3:30pm-5pm

Room: 18CD

A Case for Application-Oblivious Energy-Efficient MPI Runtime

Authors: Akshay Venkatesh (Ohio State University), Abhinav Vishnu (Pacific Northwest National Laboratory), Khaled Hamidouche (Ohio State University), Nathan Tallent (Pacific Northwest National Laboratory), Dhabaleswar Panda (Ohio State University), Darren Kerbyson (Pacific Northwest National Laboratory), Adolfo Hoisie (Pacific Northwest National Laboratory)

Power has become the major impediment in designing large scale high-end systems and runtime systems like Message Passing Interface (MPI) that serve as the communication backend for designing applications, and programming models must be made power cognizant. Slack within an MPI call provides a potential for energy and power savings, if an appropriate power reduction technique such as core-idling/DVFS can be applied without perturbing the application's execution time. This paper proposes and implements Energy Aware MPI (EAM) runtime that proves itself energy-efficient in an application-oblivious manner. EAM uses a combination of communication models of common MPI primitives and an online observation of slack for maximizing energy efficiency. We implement EAM using MVAPICH2 and evaluate it on ten applications using up to 4096 processes. Our performance evaluation on an InfiniBand cluster indicates that EAM can reduce energy consumption by 5-41% in comparison to the default approach, with negligible (< 4%) performance loss.

Award: Best Student Paper Finalist

Improving Concurrency and Asynchrony in Multithreaded MPI Applications using Software Offloading

Authors: Karthikeyan Vaidyanathan, Dhiraj D. Kalamkar, Kiran Pamnany, Jeff R. Hammond, Pavan Balaji (Argonne National Laboratory), Dipankar Das, Jongsoo Park (Intel Corporation), Balint Joo (Thomas Jefferson National Accelerator Facility)

We present a new approach for multithreaded communication and asynchronous progress in MPI applications, wherein we offload communication processing to a dedicated thread. The central premise is that given the rapidly increasing core counts on modern systems, the improvements in MPI performance arising from dedicating a thread to drive communication outweigh the small loss of resources for application computation, particularly when overlap of communication and computation can be exploited. Our approach allows application threads to

make MPI calls concurrently, enqueueing these as communication tasks to be processed by a dedicated communication thread. This not only guarantees progress for such communication operations, but also reduces load imbalance. Our implementation additionally significantly reduces the overhead of mutual exclusion seen in existing implementations for applications using MPI THREAD MULTIPLE. Our technique requires no modification to the application, and we demonstrate significant performance improvement (up to 2X) for QCD, FFT, and deep learning CNN.

Practical Scalable Consensus for Pseudo-Synchronous Distributed Systems

Authors: Thomas Herault, Aurelien Bouteiller, George Bosilca (University of Tennessee, Knoxville), Marc Gamell (Rutgers University), Keita Teranishi (Sandia National Laboratories), Manish Parashar (Rutgers University), Jack Dongarra (University of Tennessee, Knoxville)

The ability to consistently handle faults in a distributed environment requires, among a small set of basic routines, an agreement algorithm allowing surviving entities to reach a consensual decision between a bounded set of volatile resources. This paper presents an algorithm that implements an Early Returning Agreement (ERA) in pseudo-synchronous systems, which optimistically allows a process to resume its activity while guaranteeing strong progress. We prove the correctness of our ERA algorithm and expose its logarithmic behavior, which is an extremely desirable property for any algorithm that targets future exascale platforms. We detail a practical implementation of this consensus algorithm in the context of an MPI library, and evaluate both its efficiency and scalability through a set of benchmarks and two fault tolerant scientific applications.

Wednesday, November 18

Cloud Resource Management

10:30am-12pm

Room: 19AB

Monetary Cost Optimizations for MPI-Based HPC Applications on Amazon Clouds: Checkpoints and Replicated Execution

Authors: Yifan Gong, Bingsheng He, Amelie Chi Zhou (Nanyang Technological University)

In this paper, we propose monetary cost optimizations for MPI-based applications with deadline constraints on Amazon EC2 clouds. Particularly, we intend to utilize two kinds of Amazon EC2 instances (on-demand and spot instances). As a spot instance can fail at any time due to out-of-bid events, fault-tolerant executions are necessary. Through detailed studies,

we have found that two common fault tolerant mechanisms, i.e., checkpoint and replicated executions, are complementary with each other for cost-effective MPI executions on spot instances. Therefore, we formulate the optimization problem and propose a novel cost model to minimize the expected monetary cost. The experimental results with NPB benchmarks on Amazon EC2 demonstrate that: (1) it is feasible to run MPI applications with performance constraints on spot instances, and (2) our proposal achieves the significant monetary cost reduction and the necessity of adaptively choosing checkpoint and replication techniques for cost-effective and reliable MPI executions on Amazon EC2.

Elastic Job Bundling: An Adaptive Resource Request Strategy for Large-Scale Parallel Applications

Authors: Feng Liu, Jon B. Weissman (University of Minnesota Twin Cities)

In today's batch queue HPC cluster systems, the user submits a job requesting a fixed number of processors. The system will not start the job until all of the requested resources become available simultaneously. When cluster workload is high, large sized jobs will experience long waiting time due to this policy. In this paper, we propose a new approach that dynamically decomposes a large job into smaller ones to reduce waiting time, and lets the application expand across multiple subjobs while continuously achieving progress. This approach has three benefits: (i) application turnaround time is reduced, (ii) system fragmentation is diminished, and (iii) fairness is promoted. Our approach does not depend on job queue time prediction but exploits available backfill opportunities. Simulation results have shown that our approach can reduce application mean turnaround time by up to 48%.

Fault Tolerant MapReduce-MPI for HPC Clusters

Authors: Yanfei Guo (University of Colorado-Colorado Springs), Wesley Bland, Pavan Balaji (Argonne National Laboratory), Xiaobo Zhou (University of Colorado-Colorado Springs)

Building MapReduce applications using the Message-Passing Interface (MPI) enables us to exploit the performance of large HPC clusters for big data analytics. However, due to the lack of native fault tolerance support in MPI and the incompatibility between the MapReduce fault tolerance model and HPC schedulers, it is very hard to provide a fault tolerant MapReduce runtime for HPC clusters. We propose and develop FT-MRMPI, the first fault tolerant MapReduce framework on MPI for HPC clusters. We discover a unique way to perform failure detection and recovery by exploiting the current MPI semantics and the new proposal of user-level failure mitigation. We design and develop the checkpoint/restart model for fault tolerant MapReduce in MPI. We further tailor the detect/resume model to conserve work for more efficient fault tolerance. The

experimental results on a 256-node HPC cluster show that FT-MRMPI effectively masks failures and reduces the job completion time by 39%.

Interconnection Networks

10:30am-12pm

Room: 18CD

Network Endpoint Congestion Control for Fine-Grained Communication

Authors: Nan Jiang, Larry Dennison, William J. Dally (NVIDIA Corporation)

Endpoint congestion in HPC networks creates tree saturation that is detrimental to performance. Endpoint congestion can be alleviated by reducing the injection rate of traffic sources, but requires fast reaction time to avoid congestion buildup. Congestion control becomes more challenging as application communication shift from traditional two-sided model to potentially fine-grained, one-sided communication embodied by various global address space programming models. Existing hardware solutions, such as Explicit Congestion Notification (ECN) and Speculative Reservation Protocol (SRP), either react too slowly or incur too much overhead for small messages.

In this study we present two new endpoint congestion-control protocols, Small-Message SRP (SMSRP) and Last-Hop Reservation Protocol (LHRP), both targeted specifically for small messages. Experiments show they can quickly respond to endpoint congestion and prevent tree saturation in the network. Under congestion-free traffic conditions, the new protocols generate minimal overhead with performance comparable to networks with no endpoint congestion control.

Cost-Effective Diameter-Two Topologies: Analysis and Evaluation

Authors: Georgios Kathareios, Cyriel Minkenberg, Bogdan Prisacari, German Rodriguez (IBM Corporation), Torsten Hoefler (ETH Zurich)

HPC network topology design is currently shifting from high-performance, higher-cost Fat-Trees to more cost-effective architectures. Three diameter-two designs, the Slim Fly, Multi-Layer Full-Mesh, and Two-Level Orthogonal Fat-Tree excel in this, exhibiting a cost per endpoint of only 2 links and 3 router ports with lower end-to-end latency and higher scalability than traditional networks of the same total cost. However, other than for the Slim Fly, there is currently no clear understanding of the performance and routing of these emerging topologies. For each network, we discuss minimal, indirect random, and adaptive routing algorithms along with deadlock-avoidance mechanisms. Using these, we evaluate

the performance of a series of representative workloads, from global uniform and worst-case traffic to the all-to-all and near-neighbor exchange patterns prevalent in HPC applications. We show that while all three topologies have similar performance, OFTs scale to twice as many endpoints at the same cost as the others.

Profile-Based Power Shifting in Interconnection Networks with On/Off Links

Authors: Shinobu Miwa (University of Electro-Communications), Hiroshi Nakamura (University of Tokyo)

Overprovisioning hardware devices and coordinating power budgets of the devices are proposed to improve performance in the future power-constrained HPC systems. This coordination process is called power shifting. On the other hand, the latest studies uncover that on/off links have the significant capability to save network power in HPC systems, so the future HPC systems will adopt on/off links in addition to overprovisioning. This paper explores power shifting in interconnection networks with on/off links. Since on/off links keep network power low at application runtime, we can transfer the significant amount of power budgets on networks to the other devices before the application runs. Based on this observation, we propose a profile-based power shifting technique which enables HPC users to transfer the power budget remaining on networks to the other devices at the job dispatch. The experimental results show that the proposed technique significantly improves application performance under various power constraints.

State of the Practice: Infrastructure Management

10:30am-12pm

Room: 18AB

Reliability Lessons Learned From GPU Experience with the Titan Supercomputer at Oak Ridge Leadership Computing Facility

Authors: Devesh Tiwari, Saurabh Gupta (Oak Ridge National Laboratory), George Gallano (Christian Brothers University), Jim Rogers, Don Maxwell (Oak Ridge National Laboratory)

The high computational capability of graphics processing units (GPUs) is enabling and driving the scientific discovery process at large-scale. The world's second fastest supercomputer for open science, Titan, has more than 18,000 GPUs that computational scientists use to perform scientific simulations and data analysis. Understanding of GPU reliability characteristics, however, is still in its nascent stage since GPUs have only recently been deployed at large scale. This paper presents a detailed study of GPU errors and their impact on system operations and applications, describing experiences with the 18,688 GPUs

on the Titan supercomputer as well as lessons learned in the process of efficient operation of GPUs at scale. These experiences are helpful to HPC sites which already have large-scale GPU clusters or plan to deploy GPUs in the future.

Big Omics Data Experience

Authors: Patricia Kovatch, Anthony Costa, Zachary Giles, Eugene Fluder, Hyung Min Cho, Svetlana Mazurkova (Icahn School of Medicine at Mount Sinai)

As personalized medicine becomes more integrated into healthcare, the rate at which humans are being sequenced is rising quickly with a concomitant acceleration in compute and data requirements. To achieve the most effective solution for genomic workloads without re-architecting the industry-standard software, we performed a rigorous analysis of usage statistics, benchmarks and available technologies to design a system for maximum throughput. We share our experiences designing a system optimized for Genome Analysis ToolKit pipelines, based on an evaluation of compute, workload and I/O characteristics. The characteristics of genomic-based workloads are vastly different than traditional HPC workloads requiring radically different configurations of the scheduler and I/O to achieve scalability. By understanding how our researchers and clinicians work, we were able to employ new techniques to not only speed their workflow yielding improved and repeatable performance, but we were able to make efficient use of storage and compute nodes.

The Spack Package Manager: Bringing Order to HPC Software Chaos

Authors: Todd Gamblin, Matthew LeGendre, Michael R. Collette, Gregory L. Lee, Adam Moody, Bronis R. de Supinski, Scott Futral (Lawrence Livermore National Laboratory)

Large HPC centers spend considerable time supporting software for thousands of users, but the complexity of HPC software is quickly outpacing the capabilities of existing software management tools. Scientific applications require specific versions of compilers, MPI, and other dependency libraries, so using a single, standard software stack is infeasible. However, managing many configurations is difficult because the configuration space is combinatorial in size.

We introduce Spack, a tool used at Lawrence Livermore National Laboratory to manage this complexity. Spack provides a novel, recursive specification syntax to invoke parametric builds of packages and dependencies. It allows any number of builds to coexist on the same system, and it ensures that installed packages can find their dependencies, regardless of the environment. We show through real-world use cases that Spack supports diverse and demanding applications, bringing order to HPC software chaos.

Applications - Climate and Weather

1:30pm-3pm

Room: 18CD

STELLA: A Domain-Specific Tool for Structured Grid Methods in Weather and Climate Models

Authors: Tobias Gysi, Carlos Osuna (ETH Zurich), Oliver Fuhrer (Federal Office of Meteorology and Climatology MeteoSwiss), Mauro Bianco, Thomas C. Schulthess (Swiss National Supercomputing Center)

Many high-performance computing applications solving partial differential equations (PDEs) can be attributed to the class of kernels using stencils on structured grids. Due to the disparity between floating point operation throughput and main memory bandwidth these codes typically achieve only a low fraction of peak performance. Unfortunately, stencil computation optimization techniques are often hardware dependent and lead to a significant increase in code complexity. We present a domain-specific tool, STELLA, which eases the burden of the application developer by separating the architecture dependent implementation strategy from the user-code and is targeted at multi- and manycore processors. On the example of a numerical weather prediction and regional climate model (COSMO) we demonstrate the usefulness of STELLA for a real-world production code. The dynamical core based on STELLA achieves a speedup factor of 1.8x (CPU) and 5.8x (GPU) with respect to the legacy code while reducing the complexity of the user code.

Improving the Scalability of the Ocean Barotropic Solver in the Community Earth System Model

Authors: Yong Hu, Xiaomeng Huang (Tsinghua University), Alison H. Baker, Yu-heng Tseng, Frank O. Bryan, John M. Dennis (National Center for Atmospheric Research), Guangwen Yang (Tsinghua University)

High-resolution climate simulations require tremendous computing resources. In the Community Earth System Model (CESM), the ocean model is computationally expensive for high-resolution grids and is the least scalable component for most production simulations. In particular, the modified preconditioned Conjugate Gradient, used to solve the elliptic system of equations in the barotropic mode, scales poorly at the high core counts. In this work, we demonstrate that the communication costs in the barotropic solver occupy an increasing portion of the total execution time as core counts are increased. To mitigate this problem, we implement a Chebyshev-type iterative method (CSI) in the ocean model, which requires fewer global reductions, and develop an effective block preconditioner based on the Error Vector Propagation (EVP) method. We demonstrate that CSI with EVP preconditioning improves the scalability of the ocean component and

produces an ocean climate statistically consistent with the original one.

Particle Tracking in Open Simulation Laboratories

Authors: Kalin Kanov, Randal Burns (Johns Hopkins University)

Particle tracking along streamlines and pathlines is a common scientific analysis technique, which has demanding data, computation and communication requirements. It has been studied in the context of high-performance computing due to the difficulty in its efficient parallelization and its high demands on communication and computational load. In this paper, we study efficient evaluation methods for particle tracking in open simulation laboratories. Simulation laboratories have a fundamentally different architecture from today's supercomputers and provide publicly-available analysis functionality. We focus on the I/O demands of particle tracking for numerical simulation datasets 100s of TBs in size. We compare data-parallel and task-parallel approaches for the advection of particles and show scalability results on data-intensive workloads from a live production environment. We have developed particle tracking capabilities for the Johns Hopkins Turbulence Databases, which store computational fluid dynamics simulation data, including forced isotropic turbulence, magnetohydrodynamics, channel flow turbulence and homogenous buoyancy-driven turbulence.

Data Transfers and Data-Intensive Applications

1:30pm-3pm

Room: 19AB

Energy-Aware Data Transfer Algorithms

Authors: Ismail Alan, Engin Arslan, Tefvik Kosar (University at Buffalo)

The amount of data moved over the Internet per year has already exceeded the exabyte scale and soon will hit the zettabyte range. To support this massive amount of data movement across the globe, the networking infrastructure as well as the source and destination nodes consume immense amount of electric power, with an estimated cost measured in billions of dollars. Although considerable amount of research has been done on power management techniques for the networking infrastructure, there has not been much prior work focusing on energy-aware data transfer algorithms for minimizing the power consumed at the end-systems. We introduce novel data transfer algorithms which aim to achieve high data transfer throughput while keeping the energy consumption during the transfers at the minimal levels. Our experimental results show that our energy-aware data transfer algorithms can achieve up to 50% energy savings with the same or higher level of data transfer throughput.

IOrchestra: Supporting High-Performance Data-Intensive Applications in the Cloud via Collaborative Virtualization

Authors: Ron C. Chiang (University of St. Thomas), H. Howie Huang, Timothy Wood (George Washington University), Changbin Liu, Oliver Spatscheck (AT&T)

Multi-tier data-intensive applications are widely deployed in virtualized data centers for high scalability and reliability. As the response time is vital for user satisfaction, this requires achieving good performance at each tier of the applications in order to minimize the overall latency. However, in such virtualized environments, each tier (e.g., application, database, web) is likely to be hosted by different virtual machines (VMs) on multiple physical servers, where a guest VM is unaware of changes outside its domain, and the hypervisor also does not know the configuration and runtime status of a guest VM. As a result, isolated virtualization domains lend themselves to performance unpredictability and variance. In this paper, we propose IOrchestra, a holistic collaborative virtualization framework, which bridges the semantic gaps of I/O stacks and system information across multiple VMs, improves virtual I/O performance through collaboration from guest domains, and increases resource utilization in data centers.

An Elegant Sufficiency: Load-Aware Differentiated Scheduling of Data Transfers

Authors: Rajkumar Kettimuthu (Argonne National Laboratory), Gayane Vardoyan (University of Massachusetts), Gagan Agrawal, P. Sadayappan (Ohio State University), Ian Foster (Argonne National Laboratory)

We investigate the file transfer scheduling problem, where transfers among different endpoints must be scheduled to maximize pertinent metrics. We propose two new algorithms. The first, SEAL, uses runtime information and data-driven models to adapt transfer schedules and concurrency to maximize performance. We implement this algorithm using GridFTP as transfer protocol and evaluate it using real logs in a production WAN environment. Results show that SEAL can improve average slowdowns and turnaround times by up to 25% and worst-case slowdown and turnaround times by up to 50%, compared with the best-performing baseline scheme. Our second algorithm, STEAL, further leverages user-supplied categorization of transfers as either “interactive” (requiring immediate processing) or “batch” (less time-critical). Results show that STEAL reduces the average slowdown of interactive transfers by 63% compared to the best-performing baseline and by 21% compared to SEAL, while allowing batch tasks to use a large portion of the excess bandwidth.

Performance Tools and Models

1:30pm-3pm

Room: 18AB

ScaAnalyzer: A Tool to Identify Memory Scalability Bottlenecks in Parallel Programs

Authors: Xu Liu (College of William & Mary), Bo Wu (Colorado School of Mines)

It is difficult to scale parallel programs in a system that employs a large number of cores. To identify scalability bottlenecks, existing tools principally pinpoint poor thread synchronization strategies and unnecessary data communication. Memory subsystem is one of the key contributors to poor parallel scaling in multicore machines. State-of-the-art tools, however, either lack sophisticated capabilities or are completely ignorant in pinpointing scalability bottlenecks arising from the memory subsystem. To address this issue, we developed a tool—ScaAnalyzer to pinpoint scaling losses due to poor memory access behaviors of parallel programs. ScaAnalyzer collects, attributes, and analyzes memory-related metrics during program execution while incurring very low overhead. ScaAnalyzer provides high-level, detailed guidance to programmers for scalability optimization. We demonstrate the utility of ScaAnalyzer with case studies of three parallel benchmarks. For each benchmark, ScaAnalyzer identifies scalability bottlenecks caused by poor memory access behaviors and provides optimization guidance that yields significant improvement in scalability.

Award: Best Paper Finalist

C²-Bound: A Capacity and Concurrency Driven Analytical Model for Manycore Design

Authors: Yu-Hang Liu, Xian-He Sun (Illinois Institute of Technology)

We propose C²-bound, a data-driven analytical model, that incorporates both memory capacity and data access concurrency factors to optimize manycore design. C²-bound is characterized by combining the newly proposed latency model, concurrent average memory access time (C-AMAT), with the well-known memory-bounded speedup model (Sun-Ni's law) to facilitate computing tasks. Compared to traditional chip designs that lack the notion of memory concurrency and memory capacity, C²-bound model finds memory bound factors significantly impact the optimal number of cores as well as their optimal silicon area allocations. Our model is valuable to the design of new generation manycore architectures that target big data processing, where working sets are usually larger than the conventional scientific computing. With C²-bound, the design space can be narrowed down significantly up to four orders of magnitude. C²-bound analytic results can be either used by reconfigurable hardware or, applied to scheduling and partitioning resources among diverse applications.

Recovering Logical Structure from Charm++ Event Traces

Authors: Katherine E. Isaacs (University of California, Davis), Abhinav Bhatele (Lawrence Livermore National Laboratory), Jonathan Lifflander (University of Illinois at Urbana-Champaign), David Boehme, Todd Gamblin, Martin Schulz (Lawrence Livermore National Laboratory), Bernd Hamann (University of California, Davis), Peer-Timo Bremer (Lawrence Livermore National Laboratory)

Asynchrony and non-determinism in Charm++ programs present a significant challenge in analyzing their event traces. We present a new framework to organize event traces of parallel programs written in Charm++. Our reorganization allows one to more easily explore and analyze such traces by providing context through logical structure. We describe several heuristics to compensate for missing dependencies between events that currently cannot be easily recorded. We introduce a new task reordering that recovers logical structure from the non-deterministic execution order. Using the logical structure, we define several metrics to help guide developers to performance problems. We demonstrate our approach through two proxy applications written in Charm++. Finally, we discuss the applicability of this framework to other task-based runtimes and provide guidelines for tracing to support this form of analysis.

In-Situ (Simulation Time) Analysis

3:30pm-5pm

Room: 18CD

Large-Scale Compute-Intensive Analysis via a Combined In-Situ and Co-Scheduling Workflow Approach

Authors: Christopher Sewell (Los Alamos National Laboratory), Katrin Heitmann, Hal Finkel (Argonne National Laboratory), George Zagaris (Lawrence Livermore National Laboratory), Suzanne T. Parete-Koon (Oak Ridge National Laboratory), Patricia K. Fasel (Los Alamos National Laboratory), Adrian Pope (Argonne National Laboratory), Nicholas Frontiere (University of Chicago), Li-ta Lo (Los Alamos National Laboratory), Bronson Messer (Oak Ridge National Laboratory), Salman Habib (Argonne National Laboratory), James Ahrens (Los Alamos National Laboratory)

Large-scale simulations can produce tens of terabytes of data per analysis cycle, complicating and limiting the efficiency of workflows. Traditionally, outputs are stored on the file system and analyzed in post-processing. With the rapidly increasing size and complexity of simulations, this approach faces an uncertain future. Trending techniques consist of performing the analysis in situ, utilizing the same resources as the simulation, and/or off-loading subsets of the data to a compute-intensive analysis system. We introduce an analysis framework developed for HACC, a cosmological N-body code, that uses

both in situ and co-scheduling approaches for handling petabyte-size outputs. We compare different analysis set-ups ranging from purely off-line, to purely in situ to in situ/co-scheduling. The analysis routines are implemented using the PISTON/VTK-m framework, allowing a single implementation of an algorithm that simultaneously targets a variety of GPU, multi-core, and many-core architectures.

Smart: A MapReduce-Like Framework for In-Situ Scientific Analytics

Authors: Yi Wang, Gagan Agrawal (Ohio State University), Tekin Bicer (Argonne National Laboratory), Wei Jiang (Quantcast)

In-situ analytics has lately been shown to be an effective approach to reduce both I/O and storage costs. Developing an efficient in-situ implementation involves many challenges, including parallelization, data movement or sharing, and resource allocation. Although MapReduce has been widely adopted for parallelizing data analytics, there are several obstacles to applying it to in-situ scientific analytics.

In this paper, we present a novel MapReduce-like framework which supports efficient in-situ scientific analytics. It can load simulated data directly from distributed memory. It leverages a MapReduce-like API for parallelization, while meeting the strict memory constraints of in-situ analytics. It can be launched in the parallel code region of simulation program. We have developed both space sharing and time sharing modes for maximizing the performance in different scenarios. We demonstrate both high efficiency and scalability of our system, by using different simulation and analytics programs on both multi-core and many-core clusters.

Optimal Scheduling of In-Situ Analysis for Large-Scale Scientific Simulations

Authors: Preeti Malakar, Venkatram Vishwanath, Todd Munson, Christopher Knight, Mark Hereld, Sven Leyffer, Michael Papka (Argonne National Laboratory)

Today's leadership computing facilities have enabled the execution of transformative simulations at unprecedented scales. However, analyzing the huge amount of output from these simulations remains a challenge. Most analyses of this output is performed in post-processing mode at the end of the simulation. The time to read the output for the analysis can be significantly high due to poor I/O bandwidth, which increases the end-to-end simulation-analysis time. Simulation-time analysis can reduce this end-to-end time. In this work, we present the scheduling of in-situ analysis as a numerical optimization problem to maximize the number of online analyses subject to resource constraints such as I/O bandwidth, network bandwidth, rate of computation and available memory. We demonstrate the effectiveness of our approach through two application case studies on the IBM Blue Gene/Q system.

Linear Algebra

3:30pm-5pm

Room: 18AB

Exploiting Asynchrony from Exact Forward Recovery for DUE in Iterative Solvers

Authors: Luc Jaulmes, Marc Casas, Miquel Moretó, Eduard Ayguadé, Jesús Labarta, Mateo Valero (Barcelona Supercomputing Center)

This paper presents a method to protect iterative solvers from Detected and Uncorrected Errors (DUE) relying on error detection techniques already available in commodity hardware. Detection operates at the memory page level, which enables the use of simple algorithmic redundancies to correct errors. Such redundancies would be inapplicable under coarse grain error detection, but they become very powerful when the hardware is able to precisely detect errors.

Relations straightforwardly extracted from the solver allow to recover lost data exactly. This method is free of the overheads of backwards recoveries like checkpointing, and does not compromise mathematical convergence properties of the solver as restarting would do. We apply this recovery to three widely used Krylov subspace methods: CG, GMRES and BiCGStab.

We implement and evaluate our resilience techniques on CG, showing very low overheads compared to state-of-the-art solutions. Overlapping recoveries with normal work of the algorithm decreases overheads further.

Award: Best Paper Finalist

High-Performance Algebraic Multigrid Solver Optimized for Multi-Core Based Distributed Parallel Systems

Authors: Jongsoo Park, Mikhail Smelyanskiy (Intel Corporation), Ulrike Meier Yang (Lawrence Livermore National Laboratory), Dheevatsa Mudigere, Pradeep Dubey (Intel Corporation)

Algebraic multigrid (AMG) is a linear solver, well known for its linear computational complexity and excellent scalability. As a result, AMG is expected to be a solver of choice for emerging extreme scale systems. While node level performance of AMG is generally limited by memory bandwidth, achieving high bandwidth efficiency is challenging due to highly sparse irregular computation, such as triple sparse matrix products, sparse-matrix dense-vector multiplications, independent set coarsening algorithms, and smoothers such as Gauss-Seidel. We develop and analyze a highly optimized AMG implementation, based on the well-known HYPRE library. Compared to the HYPRE baseline implementation, our optimized implementation achieves 2.0x speedup on a recent Intel Haswell Xeon processor. Combined with our other multi-node optimizations, this translates into up even higher speedups when weak-scaled to multiple nodes. In addition, our implementation achieves

1.3x speedup compared to AmgX, NVIDIA's high-performance implementation of AMG, running on K40c.

STS-k: A Multilevel Sparse Triangular Solution Scheme for NUMA Multicores

Authors: Humayun Kabir (Pennsylvania State University), Joshua D. Booth (Sandia National Laboratories), Guillaume Aupy, Anne Benoit, Yves Robert (ENS Lyon), Padma Raghavan (Pennsylvania State University)

We consider techniques to improve the performance of parallel sparse triangular solution on non-uniform memory architecture multicores by extending earlier coloring and level set schemes for single-core multiprocessors. We develop STS-k, where k represents a small number of transformations for latency reduction from increased spatial and temporal locality of data accesses. We propose a graph model of data reuse to inform the development of STS-k and to prove that computing an optimal cost schedule is NP-complete. We observe significant speed-ups with STS-3 on a 32-core Intel Westmere-Ex. Execution times are reduced on average by a factor of 6 (83%) for STS-3 with coloring compared to a reference implementation using level sets. Incremental gains solely from the k level transformations in STS-k correspond to reductions in execution times by factors of 1.4 (28%) and 2 (50%) respectively, relative to reference implementations with level sets and coloring.

Management of Graph Workloads

3:30pm-5pm

Room: 19AB

Data Partitioning Strategies for Graph Workloads on Heterogeneous Clusters

Authors: Michael LeBeane, Shuang Song, Reena Panda, Jee Ho Ryoo, Lizy K. John (The University of Texas at Austin)

Large scale graph analytics are an important class of problem in the modern data center. However, while data centers are trending towards a large number of heterogeneous processing nodes, graph analytics frameworks still operate under the assumption of uniform compute resources. In this paper, we develop heterogeneity-aware data ingress strategies for graph analytics workloads using the popular PowerGraph framework. We illustrate how simple estimates of relative node computational throughput can guide heterogeneity-aware data partitioning algorithms to provide balanced graph cutting decisions. Our work enhances five online data ingress strategies from a variety of sources to optimize application execution for throughput differences in heterogeneous data centers. The proposed partitioning algorithms improve the runtime of several popular machine learning and data mining applications by as much as a 65% and on average by 32% as compared to the default balanced partitioning approaches.

Scaling Iterative Graph Computations with GraphMap

Authors: Kisung Lee (Louisiana State University), Ling Liu, Karsten Schwan, Calton Pu, Qi Zhang, Yang Zhou, Emre Yigitoglu (Georgia Institute of Technology), Pingpeng Yuan (Huazhong University of Science and Technology)

Scaling large-scale graph processing has been a heated research topic in recent years. Existing distributed graph systems are based on a distributed memory architecture. These distributed solutions heavily rely on distributed memory and thus suffer from poor scalability when the compute cluster can no longer hold the graph and all the intermediate results in memory. We present GraphMap, a distributed iterative graph computation framework, which effectively utilizes secondary storage to maximize access locality and speed up distributed iterative graph computations. GraphMap has three salient features: (1) We distinguish those data states that are mutable during iterative computations from those that are read-only in all iterations to maximize sequential accesses and minimize random accesses. (2) We devise a two-level graph-partitioning algorithm to enable balanced workloads and locality-optimized data placement. (3) We propose a suite of locality-based optimizations to maximize computation efficiency.

PGX.D: A Fast Distributed Graph Processing System

Authors: Sungpack Hong, Siegfried Depner, Thomas Manhardt, Jan Van Der Lugt (Google), Merijn Varensteen (University of Amsterdam), Hassan Chafi (Oracle Corporation)

Graph analysis is a powerful method in data analysis. In this paper, we present a fast distributed graph processing system, namely PGX.D. We show that PGX.D outperforms other distributed graph systems like GraphLab significantly (3x – 90x). Furthermore, PGX.D on 4 to 16 machines is also faster than the single-machine execution. Using a fast cooperative context-switching mechanism, we implement PGX.D as a low-overhead, bandwidth-efficient communication framework that supports remote data-pulling patterns. Moreover, PGX.D achieves large traffic reduction and good workload balance by applying selective ghost nodes, edge partitioning, and edge chunking, in transparent manners. Our analysis confirms that each of these features is indeed crucial for overall performance of certain kinds of graph algorithms. Finally, we advocate the use of balanced beefy clusters where the sustained random DRAM-access bandwidth in aggregate is matched with the bandwidth of the underlying interconnection fabric.

Award: Best Paper Finalist

Thursday, November 19

Programming Tools

10:30am-12pm

Room: 18AB

CIVL: The Concurrency Intermediate Verification Language

Authors: Stephen F. Siegel, Manchun Zheng, Ziqing Luo, Timothy K. Zirkel, Andre V. Marianiello, John G. Edenhofner (University of Delaware), Matthew B. Dwyer, Michael S. Rogers (University of Nebraska-Lincoln)

There are numerous ways to express parallel programs: message-passing libraries (MPI) and multithreading and GPU language extensions such as OpenMP, Pthreads, and CUDA, are but a few. This multitude creates a serious challenge for developers of software verification tools: it takes enormous effort to develop such tools, but each development effort typically targets one small part of the concurrency landscape, with little sharing of techniques and code among efforts.

To address this problem, we have developed CIVL: the Concurrency Intermediate Verification Language. CIVL provides a general concurrency model which can represent programs in a variety of concurrency dialects, including those listed above.

The CIVL framework includes a tool, based on model checking and symbolic execution, that can verify the absence of deadlocks, race conditions, assertion violations, illegal pointer dereferences and arithmetic, memory leaks, divisions by zero, and out-of-bound array indexing. It can also check that two programs are functionally equivalent.

Clock Delta Compression for Scalable Order-Replay of Non-Deterministic Parallel Applications

Authors: Kento Sato, Dong H. Ahn, Ignacio Laguna, Gregory L. Lee, Martin Schulz (Lawrence Livermore National Laboratory)

The ability to record and replay program execution helps significantly in debugging non-deterministic parallel applications by reproducing message receive orders. However, the large amount of data that traditional record-and-reply techniques record precludes its practical applicability to massively parallel MPI applications. In this paper, we propose a new compression algorithm, Clock Delta Compression (CDC), for scalable record and replay of non-deterministic MPI applications. CDC defines a reference order of message receives based on a totally ordered relation using Lamport clocks and only records the differences between this reference logical-clock order and an observed order. Our evaluation shows that CDC significantly reduces the record data size. For example, when we apply CDC to a Monte Carlo particle transport benchmark (MCB), which represents non-deterministic communication patterns, CDC

reduces the record size by approximately two orders of magnitude compared to traditional techniques and incurs between 13.1% and 25.5% of runtime overhead.

Relative Debugging for a Highly Parallel Hybrid Computer System

Authors: Luiz DeRose, Andrew Gontarek, Aaron Vose, Robert Moench (Cray Inc.), David Abramson, Minh Dinh, Chao Jin (University of Queensland)

Relative debugging traces software errors by comparing two executions of a program concurrently - one code being a reference version and the other faulty. Relative debugging is particularly effective when code is migrated from one platform to another, and this is of significant interest for hybrid computer architectures containing CPUs accelerators or coprocessors. In this paper we extend relative debugging to support porting stencil computation on a hybrid computer. We describe a generic data model that allows programmers to examine the global state across different types of applications, including MPI/OpenMP, MPI/OpenACC, and UPC programs. We present case studies using a hybrid version of the 'stellarator' particle simulation DELTA5D, on Titan at ORNL, and the UPC version of Shallow Water Equations on Crystal, an internal supercomputer of Cray. These case studies used up to 5,120 GPUs and 32,768 CPU cores to illustrate that the debugger is effective and practical.

Resource Management

10:30am-12pm

Room: 19AB

Improving Backfilling by Using Machine Learning to Predict Running Times

Authors: Eric Gaussier (University Grenoble Alpes), David Glessier (BULL), Valentin Reis, Denis Trystram (University Grenoble Alpes)

The job management system is the HPC middleware responsible for distributing computing power to applications. While such systems generate an ever increasing amount of data, they are characterized by uncertainties on some parameters like the job running times. The question raised in this work is: To what extent is it possible/useful to take into account predictions on the job running times for improving the global scheduling?

We present a comprehensive study for answering this question assuming the popular EASY backfilling policy. More precisely, we rely on some classical methods in machine learning and propose new cost functions well-adapted to the problem. Then, we assess our proposed solutions through intensive simulations using several production logs. Finally, we propose a new scheduling algorithm that outperforms the popular EASY backfilling algorithm by 28% considering the average bounded slowdown objective.

Adaptive Data Placement for Staging-Based Coupled Scientific Workflows

Authors: Qian Sun, Tong Jin, Melissa Romanus, Hoang Bui, Fan Zhang (Rutgers University), Hongfeng Yu (University of Nebraska-Lincoln), Hemanth Kolla (Sandia National Laboratories), Scott Klasky (Oak Ridge National Laboratory), Jacqueline Chen (Sandia National Laboratories), Manish Parashar (Rutgers University)

Data staging and in-situ/in-transit data processing are emerging as attractive approaches for supporting extreme scale scientific workflows. These approaches accelerate the data-to-insight process by enabling runtime data sharing between coupled simulations and data analytics components of the workflow. However, the complex and dynamic data exchange patterns exhibited by the workflows coupled with various data access behaviors make efficient data placement within the staging area challenging. In this paper, we present an adaptive data placement approach to address these challenges. Our approach adapts data placement based on application-specific dynamic data access patterns, and applies access pattern-driven and location-aware mechanisms to reduce data access costs and to support efficient data sharing between multiple workflow components. We experimentally demonstrate the effectiveness of our approach on Titan using combustion-analysis workflow. The evaluation results show that our approach can effectively improve data access performance and the overall efficiency of coupled scientific workflows.

Multi-Objective Job Placement in Clusters

Authors: Sergey Blagodurov (Simon Fraser University and Advanced Micro Devices, Inc.), Alexandra Fedorova (Simon Fraser University and The University of British Columbia), Evgeny Vinnik (Simon Fraser University), Tyler Dwyer (Simon Fraser University and The University of British Columbia), Fabien Hermonier (INRIA and the University of Nice - Sophia Antipolis)

One of the key decisions made by both MapReduce and HPC cluster management frameworks, is the placement of jobs within a cluster. To make this decision they consider factors like resource constraints within a node or the proximity of data to a process. However, they fail to account for the degree of collocation on the cluster's nodes. A tight process placement can create contention for the intra-node shared resources, such as shared caches, memory, disk or network bandwidth. A loose placement would create less contention, but exacerbate network delays and increase cluster-wide power consumption. Finding the best job placement is challenging, because among many possible placements, we need to find one that gives us an acceptable trade-off between performance and power consumption. We propose to tackle the problem via multi-objective optimization. Our solution is able to balance conflicting objectives specified by the user and efficiently find a suitable job placement.

Sampling in Matrix Computations

10:30am-11:30am

Room: 18CD

Randomized Algorithm to Update Partial Singular Value Decomposition on a Hybrid CPU/GPU Cluster

Authors: Ichitaro Yamazaki, Jakub Kurzak, Piotr Luszczek, Jack Dongarra (University of Tennessee, Knoxville)

For data analysis a partial singular value decomposition (SVD) of the sparse matrix representing the data is a powerful tool. However, computing the SVD of a large dataset can take a significant amount of time even on a large-scale computer. Hence, there is a growing demand for a novel algorithm that can efficiently process the massive data being generated from many modern applications. To address this challenge, in this paper, we study randomized algorithms to update the SVD as changes are made to the data. Our experimental results demonstrate that these randomized algorithms can obtain the desired accuracy of the SVD with a small number of data accesses, and compared to the state-of-the-art updating algorithm, they often require much lower computational and communication costs. Our performance results on a hybrid CPU/GPU cluster show that these randomized algorithms can obtain significant speedups over the state-of-the-art updating algorithm.

Performance of Random Sampling for Computing Low-Rank Approximations of a Dense Matrix on GPUs

Authors: Theo Mary (University of Toulouse), Ichitaro Yamazaki, Jakub Kurzak, Piotr Luszczek, Stanimire Tomov, Jack Dongarra (University of Tennessee, Knoxville)

Low-rank matrix approximations play an important role in a wide range of applications. To compute a low-rank approximation of a dense matrix, a common approach uses the QR factorization with column pivoting (QRCP). While reliable and efficient, this deterministic approach requires costly communication, which is becoming increasingly expensive on modern computers. We use an alternative approach based on random sampling, which requires much less communication than QRCP. In this paper, we compare the performance of the random sampling with that of QRCP on the NVIDIA Kepler GPU. Our performance results demonstrate that the random sampling method can be up to 13 times faster compared to QRCP while computing an approximation of comparable accuracy. We also present the parallel scaling of random sampling over multiple GPUs, showing a speedup of 5.1 over three GPUs. These results demonstrate the potential of the random sampling as an excellent computational tool for many applications.

Graph Algorithms and Benchmarks

1:30pm-3pm

Room: 18CD

A Work-Efficient Algorithm for Parallel Unordered Depth-First Search

Authors: Umut Acar (Carnegie Mellon University), Arthur Chargueraud, Mike Rainey (French Institute for Research in Computer Science and Automation)

With the increasing processing power of multicore computers, parallel graph search using shared memory machines has become increasingly feasible and important. There has been a lot of progress on parallel breadth-first search, but less attention has been given to algorithms for unordered or loosely ordered traversals. We present a parallel algorithm for unordered depth-first-search on graphs. We prove that the algorithm is work-efficient in a realistic, implementation-ready algorithmic model that includes important scheduling costs. This work-efficiency result applies to all graphs, including those with high diameter and high out-degree. The key components of this result include a new data structure and a new amortization technique for controlling excess parallelism. We present an implementation and experiments that show that the algorithm performs well on a range of graphs and can lead to significant improvements over comparable algorithms.

Enterprise: Breadth-First Graph Traversal on GPUs

Authors: Hang Liu, H. Howie Huang (George Washington University)

The Breadth-First Search (BFS) algorithm serves as the foundation for many big data applications and analytics workloads. While Graphics Processing Unit (GPU) offers massive parallelism, achieving high-performance BFS on GPUs entails efficient scheduling of a large number of GPU threads and effective utilization of GPU memory hierarchy. In this paper, we present a new BFS system, Enterprise, which utilizes three novel techniques to eliminate the performance bottlenecks: (1) streamlined GPU threads scheduling; (2) GPU workload balancing; and (3) GPU based BFS direction optimization. Enterprise achieves up to 76 billion traversed edges per second (TEPS) on a single NVIDIA Kepler K40, and up to 122 billion TEPS on two GPUs that ranks No. 45 in the Graph 500 on November 2014. Enterprise is also very energy-efficient as No. 1 in the Green-Graph 500 (small data category), delivering 446 million TEPS per watt.

GraphBIG: Understanding Graph Computing in the Context of Industrial Solutions

Authors: Lifeng Nai (Georgia Institute of Technology), Yinglong Xia, Ilie G. Tanase (IBM Corporation), Hyesoon Kim (Georgia Institute of Technology), Ching-Yung Lin (IBM Corporation)

With the emergence of data science, graph computing is becoming a crucial tool for processing big connected data. Although efficient implementations of specific graph applications exist, the behavior of full-spectrum graph computing remains unknown. To understand graph computing, we must consider multiple graph computation types, graph frameworks, data representations, and various data sources in a holistic way.

In this paper, we present GraphBIG, a benchmark suite inspired by IBM System G project. To cover major graph computation types and data sources, GraphBIG selects representative workloads and data sets from 21 real-world use cases of multiple application domains. Besides, by utilizing System G design, GraphBIG workloads incorporate modern graph frameworks and data representations. We characterized GraphBIG on real machines and observed extremely irregular memory patterns and significant diverse behaviors across different computation. GraphBIG helps users understand the architectural behavior of modern graph computing and enables future graph architecture and system research.

Resilience

1:30pm-3pm

Room: 19AB

Local Recovery and Failure Masking for Stencil-Based Applications at Extreme Scales

Authors: Marc Gamell (Rutgers University), Keita Teranishi, Michael A. Heroux, Jackson Mayo, Hemanth Kolla, Jacqueline Chen (Sandia National Laboratories), Manish Parashar (Rutgers University)

Application resilience is a key challenge that has to be addressed to realize the exascale vision. Online recovery can dramatically reduce the overhead of failures as compared to the more traditional approach where the job is terminated and restarted from the last checkpoint. In this paper we explore how online local recovery can be used for certain application classes to further reduce resilience overheads. Specifically we develop programming support and scalable runtime mechanisms to enable online and transparent local recovery for stencil-based parallel applications. We also show how multiple independent failures can be masked to effectively reduce the impact on the total time to solution. We integrate these mechanisms with the S3D combustion simulation and experimentally demonstrate (using the Titan Cray-XK7 system at ORNL) the ability to tolerate high failure rates (i.e., node

failures every 5 seconds) with low overhead while sustaining performance, at scales up to 262,144 cores.

VOCL-FT: Introducing Techniques for Efficient Soft Error Coprocessor Recovery

Authors: Antonio J. Peña, Wesley Bland, Pavan Balaji (Argonne National Laboratory)

Popular accelerator programming models rely on offloading computation operations and their corresponding data transfers to the coprocessors, leveraging synchronization points where needed. In this paper we identify and explore how such a programming model enables optimization opportunities not utilized in traditional checkpoint/restart systems, and we analyze them as the building blocks for an efficient fault-tolerant system for accelerators. Although we leverage our techniques to protect from detected but uncorrected ECC errors in the device memory in OpenCL-accelerated applications, coprocessor reliability solutions based on different error detectors and similar API semantics can directly adopt the techniques we propose. Adding error detection and protection involves a tradeoff between runtime overhead and recovery time. Although optimal configurations depend on the particular application, the length of the run, the error rate, and the temporary storage speed, our test cases reveal a good balance with significantly reduced runtime overheads.

Understanding the Propagation of Transient Errors in HPC Applications

Authors: Rizwan Ashraf (University of Central Florida), Roberto Gioiosa, Gokcen Kestor (Pacific Northwest National Laboratory), Ronald DeMara (University of Central Florida), Chen-Yong Cher, Pradip Bose (IBM Corporation)

Resiliency of exascale systems has quickly become an important concern for the scientific community. Despite its importance, much remains to be determined regarding how faults disseminate or at what rate do they impact HPC applications. The understanding of where and how fast faults propagate could lead to more efficient implementation of application-driven error detection and recovery.

In this work, we propose a fault propagation framework to analyze how faults propagate in MPI applications and to understand their vulnerability to faults. We employ a combination of compiler-level code transformation and instrumentation, along with a runtime checker. Using the information provided by our framework, we employ machine learning technique to derive application fault propagation models that can be used to estimate the number of corrupted memory locations at runtime.

State of the Practice: Measuring Systems

1:30pm-3pm

Room: 18AB

Scientific Benchmarking of Parallel Computing Systems

Authors: Torsten Hoefler, Roberto Belli (ETH Zurich)

Measuring and reporting performance of parallel computers constitutes the basis for scientific advancement of high-performance computing (HPC). Most scientific reports show performance improvements of new techniques and are thus obliged to ensure reproducibility or at least interpretability. Our investigation of a stratified sample of 120 papers across three top conferences in the field shows that the state of the practice is not sufficient. For example, it is often unclear if reported improvements are in the noise or observed by chance. In addition to distilling best practices from existing work, we propose statistically sound analysis and reporting techniques and simple guidelines for experimental design in parallel computing. We aim to improve the standards of reporting research results and initiate a discussion in the HPC field. A wide adoption of this minimal set of rules will lead to better reproducibility and interpretability of performance results and improve the scientific culture around HPC.

Node Variability in Large-Scale Power Measurement: Perspectives from the Green500, Top500 and EEHPCWG

Authors: Thomas R. W. Scogland (The Green500 and Lawrence Livermore National Laboratory), Jonathan Azose (University of Washington), David Rohr (Goethe University Frankfurt), Suzanne Rivoire (Sonoma State University), Natalie Bates (Energy Efficient HPC Working Group), Daniel Hackenberg (Technical University of Dresden), Torsten Wilde (Leibniz Supercomputing Center), James H. Rogers (Oak Ridge National Laboratory)

The last decade has seen power consumption move from an afterthought to the foremost design constraint of new supercomputers. Measuring the power of a supercomputer can be a daunting proposition and, as a result, many published measurements are extrapolated. This paper explores the validity of these extrapolations in the context of inter-node power variability and power variations over time within a run. We characterize power variability across nodes in systems at eight supercomputer centers across the globe. Based on these, we find that the current requirement for measurements submitted to the Green500 and others is insufficient, allowing variations of as much as 20% from time of measurement and a further 10-15% from insufficient sample sizes. This paper proposes new measurement requirements to ensure consistent accuracy for power and energy measurements of supercomputers, some of which have been accepted for use by the Green500 and Top500.

A Practical Approach to Reconciling Availability, Performance, and Capacity in Provisioning Extreme-Scale Storage Systems

Authors: Lipeng Wan (University of Tennessee, Knoxville), Feiyi Wang, Sarp Oral, Devesh Tiwari, Sudharshan S. Vazhkudai (Oak Ridge National Laboratory), Qing Cao (University of Tennessee, Knoxville)

The increasing data demands from high-performance computing applications significantly accelerate the capacity, capability, and reliability requirements of storage systems. As systems scale, component failures and repair times increase, with significant impact to data availability. A wide array of decision points must be balanced in designing such systems.

We propose a systematic approach that balances and optimizes both initial and continuous spare provisioning based on a detailed investigation of the anatomy and field failure data analysis of extreme-scale storage systems. We consider both the component failure characteristics, the cost and the impact at the system level simultaneously. We build a provisioning tool to evaluate different provisioning schemes, and the results demonstrate that our optimized provisioning can reduce the unavailable duration by as much as 52% under a fixed budget. We also observe that non-disk components have much higher failure rates than disks, therefore warrant careful considerations in the overall provisioning process.

Power Constrained Computing

3:30pm-5pm

Room: 19AB

Analyzing and Mitigating the Impact of Manufacturing Variability in Power-Constrained Supercomputing

Authors: Yuichi Inadomi (Kyushu University), Tapasya Patki (University of Arizona), Koji Inoue, Mutsumi Aoyagi (Kyushu University), Barry Rountree, Martin Schulz (Lawrence Livermore National Laboratory), David Lowenthal (University of Arizona), Yasutaka Wada (Meisei University), Keiichiro Fukazawa (Kyoto University), Masatsugu Ueda (Kyushu University), Masaaki Kondo (University of Tokyo), Ikuo Miyoshi (Fujitsu)

A key challenge in next-generation supercomputing is to effectively schedule limited power resources. Modern processors suffer from increasingly large power variations owing to the chip manufacturing process. These variations lead to power inhomogeneity in current systems and manifest into performance inhomogeneity in power constrained environments, drastically limiting supercomputing performance. We present a first-of-its-kind manufacturing variability study on four production HPC systems spanning four microarchitectures, analyze its impact on HPC applications, and propose a novel variation-aware power budgeting scheme to maximize effective application performance. Our low-cost and scalable

budgeting algorithm strives to achieve performance homogeneity under a power constraint by deriving application-specific, module-level power allocations. Experimental results using a 1920 socket system show up to 5.4x speedup, with an average speedup of 1.8x across all benchmarks, compared to a variation-unaware power allocation scheme.

Finding the Limits of Power-Constrained Application Performance

Authors: Peter E. Bailey, Aniruddha Marathe, David K. Lowenthal (University of Arizona), Barry Rountree, Martin Schulz (Lawrence Livermore National Laboratory)

As we approach exascale systems, power is turning from an optimization goal to a critical constraint. With power bounds imposed by both stakeholders and the limitations of existing infrastructure, we need new techniques that extract maximum performance from limited power. In this paper, we find the theoretical upper bound of computational performance on a per-application basis in hybrid MPI + OpenMP applications. We use a linear programming formulation to optimize application schedules under various power constraints, where a schedule consists of a DVFS state and number of OpenMP threads. We also provide a mixed integer-linear formulation and show that the resulting schedules closely match schedules from the LP formulation. Across applications, we use our LP-derived upper bounds to show that current approaches trail optimal, power-constrained performance by up to 41.1%. Our LP formulation provides future optimization approaches with a quantitative optimization target.

Dynamic Power Sharing for Higher Job Throughput

Authors: Daniel A. Ellsworth, Allen D. Malony (University of Oregon), Barry Rountree, Martin Schulz (Lawrence Livermore National Laboratory)

Current trends for high-performance systems are leading us toward hardware over-provisioning where it is no longer possible to run all components at peak power without exceeding a system or facility wide power bound. Static power scheduling is likely to lead to inefficiencies with over and under provisioning of power to components at runtime. In this paper we investigate the performance and scalability of POWsched, an application agnostic runtime power scheduler capable of enforcing a system-wide power limit. Our experimental results show POWsched can improve overall runtime by as much as 14%. We also contribute a model and simulation, POWsim, for investigating dynamic power scheduling and enforcement at scale.

Programming Systems

3:30pm-5pm

Room: 18AB

Regent: A High-Productivity Programming Language for HPC with Logical Regions

Authors: Elliott Slaughter, Wonchan Lee, Sean Treichler (Stanford University), Michael Bauer (NVIDIA Corporation), Alex Aiken (Stanford University)

We present Regent, a high-productivity programming language for high performance computing with logical regions. Regent users compose programs with tasks (functions eligible for parallel execution) and logical regions (hierarchical collections of structured objects). Regent programs appear to execute sequentially, require no explicit synchronization, and are trivially deadlock-free. Regent's type system catches many common classes of mistakes and guarantees that a program with correct serial execution produces identical results on parallel and distributed machines.

We present an optimizing compiler for Regent that translates Regent programs into efficient implementations for Legion, an asynchronous task-based model. Regent employs several novel compiler optimizations to minimize the dynamic overhead of the runtime system and enable efficient operation. We evaluate Regent on three benchmark applications and demonstrate that Regent achieves performance comparable to hand-tuned Legion.

Bridging OpenCL and CUDA: A Comparative Analysis and Translation

Authors: Junghyun Kim, Thanh Tuan Dao, Jaehoon Jung, Jinyoung Joo, Jaejin Lee (Seoul National University)

Heterogeneous systems are widening their user-base, and heterogeneous computing is becoming popular in supercomputing. Among others, OpenCL and CUDA are the most popular programming models for heterogeneous systems. Although OpenCL inherited many features from CUDA and they have almost the same platform model, they are not compatible with each other. In this paper, we present similarities and differences between them and propose an automatic translation framework for both OpenCL to CUDA and CUDA to OpenCL. We describe features that make it difficult to translate from one to the other and provide our solution. We show that our translator achieves comparable performance between the original and target applications in both directions. Since each programming model separately has a wide user-base and large code-base, our translation framework is useful to extend the code-base for each programming model and unifies the efforts to develop applications for heterogeneous systems.

CilkSpec: Optimistic Concurrency for Cilk

Authors: Shaizeen Aga (University of Michigan, Ann Arbor), Sriram Krishnamoorthy (Pacific Northwest National Laboratory), Satish Narayanasamy (University of Michigan, Ann Arbor)

Recursive parallel programming models such as Cilk strive to simplify the task of parallel programming by enabling a simple divide-and-conquer model of programming. This model is effective in recursively partitioning work into smaller parts and combining their results. However, recursive work partitioning can impose additional constraints on concurrency than is implied by the true dependencies in a program. In this paper, we present a speculation-based approach to alleviate the concurrency constraints imposed by such recursive parallel programs. We design a runtime infrastructure that supports speculative execution and a predictor to accurately learn and identify opportunities to relax extraneous concurrency constraints. Experimental evaluation demonstrates that speculative relaxation of concurrency constraints can deliver gains of up to 1.6 times on 30 cores over baseline Cilk.

Tensor Computation

3:30pm-4:30pm

Room: 18CD

An Input-Adaptive and In-Place Approach to Dense Tensor-Times-Matrix Multiply

Authors: Jiajia Li, Casey Battaglini, Ioakeim Perros, Jimeng Sun, Richard Vuduc (Georgia Institute of Technology)

This paper describes a novel framework, called InTensLi (“intensely”), for producing fast single-node implementations of dense tensor-times-matrix multiply (Ttm) of arbitrary dimension. Whereas conventional implementations of Ttm rely on explicitly converting the input tensor operand into a matrix—in order to be able to use any available and fast general matrix-matrix multiply (Gemm) implementation—our framework’s strategy is to carry out the Ttm in-place, avoiding this copy. As the resulting implementations expose tuning parameters, this paper also describes a heuristic empirical model for selecting an optimal configuration based on the Ttm’s inputs. When compared to widely used Ttm implementations that are available in the Tensor Tool-box and Cyclops Tensor Framework (Ctf), InTensLi’s in-place and input-adaptive Ttm implementations achieve 4x and 13x speedups, showing Gemm-like performance on a variety of input sizes.

Scalable Sparse Tensor Decompositions in Distributed Memory Systems

Authors: Oguz Kaya, Bora Ucar (ENS Lyon)

We investigate efficient parallelization of the most common iterative sparse tensor decomposition algorithms on distributed memory systems. A key operation in each iteration of these algorithms is the matricized tensor times Khatri-Rao product (MTTKRP), which amounts to element-wise vector multiplication and reduction depending on the sparsity of the tensor. We investigate fine and coarse-grain task definitions for this operation, and propose hypergraph partitioning-based methods for these to achieve load balance as well as reduce the communication requirements. We also design distributed memory sparse tensor library, HyperTensor, which implements a well-known algorithm for the CANDECOMP/PARAFAC(CP) decomposition utilizing these task definitions and partitions. We use this library to test the scalability of the proposed implementation of MTTKRP in CP decomposition context up to 1024 MPI ranks. We observed up to 194 fold speedups using 512 MPI processes on a real world data, and significantly better scalability than a state of the art implementation.



Research Posters/ ACM Student Research Competition Posters

**Tuesday, November 17 –
Thursday November 19**

Research Posters Exhibit ACM Student Research Competition Posters Exhibit

Chairs: Michela Becchi (University of Missouri), Dorian C. Arnold (University of New Mexico), Manish Parashar (Rutgers University)

8:30am-5pm

Room: Level 4 – Concourse

Reception & Exhibit

Tuesday, November 17

5:15pm-7pm

Room: Level 4 – Concourse

Research Posters

Resource Usage Characterization for Social Network Analytics on Spark

Authors: Irene Manotas (University of Delaware); Rui Zhang, Min Li, Renu Tewari, Dean Hildebrand (IBM Corporation)

Platforms for Big Data Analytics such as Hadoop, Spark, and Storm have gained large attention given their easy-to-use programming model, scalability, and performance characteristics. Along with the wide adoption of these big data platforms, Online Social Networks (OSN) have evolved as one of the major sources of information given the large amount of data being generated in a daily basis from different online communities such as Twitter, Facebook, etc. However, current benchmarks neither consider the evaluation of big data platforms with workloads targeted for OSN analytics nor the usage of OSN

data as input. Hence, there are no studies that characterize the resource utilization of algorithms used for OSN analytics on big data platforms. This poster presents the resource characterization of two workloads for OSN data. Our results show the data patterns and major resource demands that could appear when analyzing OSN data.

A Coding Based Optimization for Hadoop

Authors: Zakia Asad (University of Toronto), Mohammad Asad Rehman Chaudhry (soptimizer), David Malone (The Hamilton Institute)

The rise of cloud and distributed data-intensive (“Big Data”) applications puts pressure on data center networks due to the movement of massive volumes of data. Reducing volume of communication is pivotal for embracing greener data exchange by efficiently utilizing network resources. This work proposes the use of coding techniques working in tandem with software-defined network control as a means of dynamically-controlled reduction in volume of communication. We introduce motivating real-world use-cases, and present a novel sparse coding algorithm for the data center networks. Moreover, we bridge the gap between theory and practice by performing a proof-of-concept implementation of the proposed system in a real world data center. We use Hadoop as our target framework. The experimental results show advantage of proposed system compared to “vanilla Hadoop implementation”, “in-network combiner”, and “Combine-N-Code” in terms of volume of communication, goodput, and number of bits that can be transmitted per Joule of energy.

BLAST Motivated Small Dense Linear Algebra Library Comparison

Authors: Pate Motter (University of Colorado Boulder); Ian Karlin, Christopher Earl (Lawrence Livermore National Laboratory)

Future computing architectures will be more memory bandwidth bound than current machines. Higher-order algorithms are more compute intense and can take advantage of the extra compute capabilities of future machines to produce higher

quality answers. In this poster we focus on BLAST, an arbitrary order arbitrary Lagrangian-Eulerian (ALE) finite element code. Typical of finite element codes, BLAST requires both global sparse and local dense matrix solves. The dense matrix solves and sparse matrix assembly perform numerous small, dense matrix calculations and consume most of its runtime. Many libraries focus on optimizing small linear algebra operations. We created a benchmark suite that mimics BLAST's most computationally intensive portions, currently implemented using MFEM. We use this suite to explore the performance of these libraries. For our benchmarks Armadillo, Blaze, and a template-based version of MFEM all produce promising results. Eigen's feature set made it promising; however, its performance was not competitive.

Numerical Tools for Multiscale Wave Propagation in Geophysics

Authors: Jose Camata (Federal University of Rio de Janeiro); Lucio de Abreu Correa, Luciano de Carvalho Paludo, Regis Cottereau (French National Center for Scientific Research); Alvaro Coutinho (Federal University of Rio de Janeiro)

Current methods of elastic tomography in geophysical media are mainly based on the arrival times of the first waves, which correspond to a particular homogenization regime. The late coda, that corresponds to a diffusion-like regime is mostly disregarded. We want to complement the classical tomography methods with statistical information gathered from the coda. Such an objective requires the construction of numerical tools that can efficiently adapt to the required scale of study, in particular (i) a scalable mesh device that automatically considers topographical details up to a parameterizable level (ii) a spectral element solver for elastic wave propagation, and (iii) a scalable random field generator for elastic parameters. We address item (i) with an octree-based meshing algorithm accounting for the depth-dependent velocity structure of the Earth. Item (iii) is dealt with by a superposition algorithm adapted when the propagation length is large compared to the correlation length.

Scalable Mesh Generation for HPC Applications

Authors: Rajeev Jain, Navamita Ray, Iulian Grindeanu, Danqing Wu, Vijay Mahadevan (Argonne National Laboratory)

Computational solvers simulating physical phenomena on complex domain geometries need well resolved, high-quality meshes to tackle the discrete problem accurately and efficiently. High fidelity real-world HPC solvers have intricate mesh generation needs, which is accomplished through access to geometry data and scalable algorithms to produce optimal quality unstructured meshes. We provide an overview of the open-source computational workflow enabled through SIGMA toolkit and focus on scalability of parallel mesh generation

algorithms implemented in MeshKit library. The algorithmic efficiency in MeshKit is evaluated and component-wise profiling is performed to investigate bottlenecks in the underlying parallel mesh infrastructure (MOAB).

Neuroscience Gateway - Enabling HPC for Computational Neuroscience

Authors: Subhashini Sivagnanam, Amitava Majumdar (San Diego Supercomputer Center); Pramod Kumbhar (Swiss Federal Institute of Technology in Lausanne), Michael Hines (Yale University), Kenneth Yoshimoto (San Diego Supercomputer Center), Ted Carnevale (Yale University)

In this poster, we describe the Neuroscience Gateway that enables HPC access for the computational neuroscience community since 2013. A central challenge in neuroscience is to understand how brain function emerges from interactions of a large number of biological processes at multiple physical and temporal scales. Computational modeling is an essential tool for developing this understanding. Driven by a rapidly expanding body of empirical observations, models and simulation protocols are becoming increasingly complex. This has stimulated development of powerful, open source computational neuroscience simulators, which run efficiently on HPC systems. This has also resulted in critical need for access to HPC by computational neuroscientists. The Neuroscience gateway hides the complexities of using HPC systems directly and allows researchers to seamlessly access computational neuroscience codes that are made available on various HPC resources. We also report on performance of the NEURON application on HPC machines.

Scalable Spatial Data Synthesis in Chameleon: Towards Performance Models for Elastic Cloud Appliances

Authors: Luis Pineda-Morales (Microsoft Corporation and French Institute for Research in Computer Science and Automation); Balaji Subramaniam, Kate Keahey (Argonne National Laboratory); Gabriel Antoniu, Alexandru Costan (French Institute for Research in Computer Science and Automation); Shaowen Wang, Anand Padmanabhan, Aiman Soliman (University of Illinois at Urbana-Champaign)

Several scientific domains rely on the ability to synthesize spatial data, embedded with geographic references. Economics and sociology, for instance, use spatial data to analyze and describe population dynamics. As the sources of spatial data, such as sensors and social media, have become more accurate and numerous, the generated data has considerably grown in size and complexity over the past years. As a consequence, larger computing capabilities are required for storing, processing and visualizing the data. In the recent years, cloud computing has emerged as a convenient infrastructure for supporting current spatial data synthesis needs, since they offer

dynamically provisioned and fairly inexpensive resources. In this poster, we aim to leverage cloud computing resources for enabling spatial data synthesis.

Utility-Based Data Transfers Scheduling Between Distributed Computing Facilities

Authors: Xin Wang (Illinois Institute of Technology); Wei Tang, Rajkumar Kettimuttu (Argonne National Laboratory); Zhiling Lan (Illinois Institute of Technology)

Today's scientific applications increasingly involve large amounts of input/output data that must be moved among multiple computing facilities via wide-area networks (WANs). The bandwidth of WANs, however, is growing at a much smaller rate and thus becoming a bottleneck. Moreover, the network bandwidth has not been viewed as a limited resource, and thus coordinated allocation is lacking. Uncoordinated scheduling of competing data transfers over shared network links results in suboptimal system performance and poor user experiences. To address these problems, we propose a data transfer scheduler to coordinate data transfers between distributed computing facilities over WANs. The scheduler allocates resources to data transfer requests based on user-centric utility functions in order to achieve maximum overall user satisfaction. We conducted trace-based simulation and demonstrated that our data transfer scheduling algorithms can considerably improve data transfer performance as well as quantified user satisfaction compared with traditional first-come, first-serve or short-job-first approaches.

FPGA Based OpenCL Acceleration of Genome Sequencing Software

Authors: Ashish Sirasao, Elliott Delaye, Ravi Sunkavalli, Stephen Neuendorffer (Xilinx Inc.)

The Smith-Waterman algorithm produces the optimal pairwise alignment between two sequences of proteins or nucleotides and is frequently used as a key component of alignment and variation detection tools for next-generation sequencing data. In this paper an efficient and scalable implementation of the Smith-Waterman algorithm is written in OpenCL and implemented on a Xilinx Virtex-7 FPGA which shows >2x compute performance and 18x-20x performance per watt advantage compared with 12 core and 60 core CPUs. These results were achieved using an off the shelf PCIe accelerator card by optimizing kernel throughput of a systolic array architecture with compiler pipelining of the OpenCL kernels and adjusting the number of systolic nodes per compute unit to fully utilize the FPGA resources.

Benchmarking High Performance Graph Analysis Systems with Graph Mining and Pattern Matching Workloads

Authors: Seokyeong Hong (North Carolina State University); Seung-Hwan Lim, Sangkeun Lee, Sreenivas R. Sukumar (Oak Ridge National Laboratory); Ranga R. Vatsavai (North Carolina State University)

The increases in volume and inter-connectivity of graph data result in the emergence of recent scalable and high performance graph analysis systems. Those systems provide different graph representation models, a variety of querying interfaces and libraries, and several underlying computation models. As a consequence, such diversities complicate in-situ choices of optimal platforms for data scientists given graph analysis workloads. In this poster, we compare recent high performance and scalable graph analysis systems in distributed and supercomputer-based processing environments with two important graph analysis workloads: graph mining and graph pattern matching. We also compare those systems in terms of expressiveness and suitability of their querying interfaces for the two distinct workloads.

Large-Scale MO Calculation with a GPU-Accelerated FMO Program

Authors: Hiroaki Umeda (University of Tsukuba), Toshihiro Hanawa (University of Tokyo); Mitsuo Shoji, Taisuke Boku, Yasuteru Shigeta (University of Tsukuba)

A GPU-enabled Fragment Molecular Orbital (FMO) program has been developed with CUDA and has executed performance benchmarks, including the first large-scale GPU-accelerated FMO calculation.

FMO is one of ab initio molecular orbital methods for large molecules and can be executed on a modern HPC computer system, such as a GPU cluster. There are two hotspots in the FMO calculation: (a) Fock matrix preparation and (b) electrostatic potential (ESP) calculation. GPU-enable Fock matrix preparation is implemented with the full advantage of a two-electron integral symmetric property without costly exclusive accumulation to a shared matrix. For ESP calculation, four-center inter-fragment Coulomb interaction is implemented for the GPGPU with a same strategy as Fock matrix preparation.

Performance Benchmark shows 3.8x speedup from CPU for on-the-fly calculation. With larger benchmarks, FMO calculation of 23,460 and 46,641 atomic molecular systems were performed with 256 and 512 GPU systems, respectively, and these large-scale GPU-accelerated FMO calculations successfully executed in two hours.

Fast Classification of MPI Applications Using Lamport's Logical Clocks

Authors: Zhou Tong (Florida State University)

We propose a tool for fast classification and identification of performance limiting factors in MPI applications. Our tool is based on Lamport's logical clock and takes DUMPI traces as input. By simulating the DUMPI traces for an application in one pass, the tool is able to determine whether the application is computation-bound, communication-bound, or has synchronization issues in different network configurations. The key idea is to use Lamport's logical clock and maintain many logical clock counters that keep track of application time in different configurations. This allows one-pass simulation to produce performance estimations for many different network configurations. By observing the application performance for different network configurations, the performance characteristics of the application as well as the performance limiting factors of the application are revealed. The techniques used in the tool are detailed, and our experiments with the classification of 9 DOE miniapps are described.

HPX Applications and Performance Adaptation

Authors: Alice Koniges (Lawrence Berkeley National Laboratory), Jayashree Ajay Candadai (Indiana University), Hartmut Kaiser (Louisiana State University), Kevin Huck (University of Oregon), Jeremy Kemp (University of Houston), Thomas Heller (Friedrich-Alexander University Erlangen-Nürnberg), Matthew Anderson (Indiana University), Andrew Lumsdaine (Indiana University), Adrian Serio (Louisiana State University), Ron Brightwell (Sandia National Laboratories), Thomas Sterling (Indiana University)

This poster focuses on application performance under HPX. Developed world-wide, HPX is emerging as a critical new programming model combined with a runtime system that uses an asynchronous style to escape the traditional static communicating sequential processes execution model, namely MPI, with a fully dynamic and adaptive model exploiting the capabilities of future generation performance-oriented runtime systems and hardware architecture enhancements. We focus on comparing application performance on standard supercomputers such as a Cray XC30 to implementations on next generation testbeds. We also describe performance adaptation techniques and legacy application migration schemes. We discuss which applications benefit substantially from the asynchronous formulations, and which, because of communication patterns and inherent synchronization points, and other details will require rewriting for newer architectures. Most of our applications show improvement in their new implementations, and improved performance on the next generation hardware is even more pronounced.

Virtualizing File Transfer Agents for Increased Throughput on a Single Host

Authors: Thomas Stitt (Pennsylvania State University), Amanda Bonnie (Los Alamos National Laboratory), Zach Fuerst (Dakota State University)

Single Lustre File Transfer Agent (FTA) performance is known to under utilize the bandwidth of Infiniband (IB) cards. The utility and viability of multiple Lustre FTA Virtual Machines (VMs) for improved network throughput on a single node was investigated. It is proposed that having multiple VMs on a single node will help achieve better usage of the IB card. Kernel-based Virtual Machines (KVM) was chosen as the hypervisor because of its ease of use, popularity, and compatibility with the drivers needed for the Mellanox ConnectX-3 IB card. Single Root - I/O Virtualization (SR-IOV) was configured so that the IB card could be divided up among the VMs. SR-IOV allows direct access to PCIe hardware, bypassing the hypervisor, leading to reduced latency and improved performance. Our results lead us to conclude that this method of provisioning FTAs should be further explored for HPC production use.

Mitos: A Simple Interface for Complex Hardware Sampling and Attribution

Authors: Alfredo Gimenez (University of California, Davis); Benafsh Husain, David Boehme, Todd Gamblin, Martin Schulz (Lawrence Livermore National Laboratory)

As high-performance architectures become more intricate and complex, so do the capabilities to monitor their performance. In particular, hardware sampling mechanisms have extended beyond the traditional scope of code profiling to include performance data relevant to the data address space, hardware topology, and load latency. However, these new mechanisms do not adhere to a common specification and as such, require architecture-specific knowledge and setup to properly function. In addition, the incoming data is often low-level and difficult to attribute to any useful context without a high level of expertise. This has resulted in a destructively steep barrier to entry, both in data acquisition and analysis.

Mitos provides a simple interface for modern hardware sampling mechanisms and the ability to define attributions of low-level samples to intuitive contexts. We present the Mitos API and demonstrate how to use it to create detailed yet understandable visualizations of performance data for analysis.

HPC Enabled Real-Time Remote Processing of Laparoscopic Surgery

Authors: *Karan Sapra, Zahra Ronaghi, Ryan Izard, David M. Kwartowitz, Melissa C. Smith (Clemson University)*

Laparoscopic surgery is a minimally invasive surgical technique where surgeons insert a small video camera into the patient's body to visualize internal organs and use small tools to perform surgical procedures. However, the benefit of small incisions has a drawback of limited subsurface tissue visualization. Image-guided surgery (IGS) uses images to map subsurface structures and can reduce the limitations of laparoscopic surgery. One particular laparoscopic camera system is the vision system of the daVinci robotic surgical system. The video streams generate approximately 360 MB of data per second, demonstrating a trend towards increased data sizes in medicine. Processing this huge stream of data on a single or dual node setup is a challenging task, thus we propose High Performance Computing (HPC) enabled framework for laparoscopic surgery. We utilize high-speed networks to access computing clusters to perform various operations on pre- and intra-operative images in a secure, reliable and scalable manner.

Consistent Hashing Distance Metrics for Large-Scale Object Storage

Authors: *Philip Carns, Kevin Harms, John Jenkins, Misbah Mubarak, Robert B. Ross (Argonne National Laboratory); Christopher Carothers (Rensselaer Polytechnic Institute)*

Large-scale storage systems often use object-level replication to protect against server failures. This calls for the development of efficient algorithms to map replicas to servers at scale. We introduce a modular, open-source consistent hashing library, called "libch-placement," and use it to investigate design trade-offs in CPU efficiency, replica declustering, and object sharding for replicated object placement. We identify a multiring hashing configuration that strikes a balance between all three properties. We also demonstrate that large-scale object placement can be calculated at an aggregate rate of tens of billions of objects per second with 8,192 servers. These technologies can be leveraged by large-scale storage systems to improve performance, improve reliability, and manage data placement with per-object granularity.

Dynamic Adaptively Refined Mesh Simulations on 1M+ Cores

Authors: *Brian T. N. Gunney (Lawrence Livermore National Laboratory)*

Patch-based structured adaptive mesh refinement (SAMR) is widely used for high-resolution simulations. Combined with modern supercomputers, it can provide simulations of unprecedented size and resolution, but managing dynamic SAMR meshes at large scales is challenging. Distributed mesh

management is a scalable approach, but early distributed algorithms still had trouble scaling past 10^5 MPI tasks. This poster presents two critical regridding algorithms, integrated into that approach to ensure efficiency of the whole. The clustering algorithm is an extension of the tile-clustering approach, making it more flexible and efficient in both clustering and parallelism. The partitioner is a new algorithm designed to mitigate the network congestion experienced by its predecessor. We evaluated performance using weak- and strong-scaling benchmarks designed to be difficult for dynamic adaptivity. Results show good scaling on up to 1.5M cores and 2M MPI tasks. Detailed timing diagnostics suggest scaling would continue well past that.

Scalable and Highly SIMD-Vectorized Molecular Dynamics Simulation Involving Multiple Bubble Nuclei

Authors: *Hiroshi Watanabe, Satoshi Morita (The University of Tokyo); Hajime Inaoka (RIKEN Advanced Institute for Computational Science), Haruhiko Matsuo (Research Organization for Information Science and Technology); Synge Todo, Nobuyasu Ito (The University of Tokyo)*

While understanding the behaviors of gas-liquid multiphase systems is important for the field of engineering, the simulation of systems involving multiple bubble nuclei is challenging since (1) the system involves travels, creations and annihilations of phase boundaries, and (2) a huge system is required in order to investigate the interactions between bubbles. We overcame the above difficulties by implementing a scalable molecular dynamics code adopting hybrid parallelism which allows us to perform simulations involving 13.7 billion Lennard-Jones atoms on the full nodes of the K computer. The simulation code is highly tuned to SPARC architecture and 98.8 of arithmetic operations are SIMDized. A performance of 2.44 PFlops representing 23.0 of peak is achieved. We present an unprecedented simulation of cavitation involving multiple bubble nuclei. The detailed behavior of bubbles in the early stage of cavitation is captured for the first time.

Transition to Trinity

Authors: *Kathryn S. Protin, Susan K. Coulter, Jesse E. Martinez, Alex F. Montañó, Charles D. Wilder (Los Alamos National Laboratory)*

This poster will discuss the work done by the network team in the High Performance Computing Systems (HPC-3) group at Los Alamos National Laboratory (LANL) to prepare a network infrastructure suitable for the Trinity supercomputer. The team transitioned from our previous backbone, Parallel Scalable Backbone (PaScaLBB), to the Next Generation Backbone (NGBB) upon Trinity's arrival in June 2015, with its unprecedented performance of over 40 petaflops and an attached 80-petabyte parallel file system. This poster outlines past supercomputing work done at LANL, the need for an improved

network infrastructure that will last until approximately 2025, the past and future network backbones, and the planning and execution of the transition to the new backbone. In conclusion, we review the challenges faced and lessons learned as we migrated production supercomputers and file systems to a completely new network backbone with minimal disruption to LANL's HPC environment and its users.

Comparison of Machine-Learning Techniques for Handling Multicollinearity in Big Data Analytics and High-Performance Data Mining

Authors: Gerard Dumancas (Oklahoma Baptist University), Ghalib Bello (Virginia Commonwealth University)

Big data analytics and high-performance data mining have become increasingly popular in various fields. They focus on the automated analysis of large-scale data, a process ideally involving minimal human input. A typical big data analytic scenario involves the use of thousands of variables, many of which will be highly correlated. Using mortality and moderately correlated lipid profile data from the NHANES database, we compared the predictive capabilities of individual parametric and nonparametric machine-learning techniques, as well as 'stacking', an ensemble learning technique. Our results indicate that orthogonal-partial least squares-discriminant analysis offers the best performance in the presence of multicollinearity, and that the use of stacking does not significantly improve predictive performance. The insights gained from this study could be useful in selecting machine-learning methods for automated pre-processing of thousands of correlated variables in high-performance data mining.

Beating cuBLAS: Automatically Generating Bespoke Matrix Multiplication Kernels Using GiMMiK

Authors: Freddie D. Witherden, Bartosz D. Wozniak, Francis P. Russell, Peter E. Vincent, Paul H. J. Kelly (Imperial College London)

Matrix multiplication is a fundamental performance primitive ubiquitous in all areas of science and engineering. In this work we present GiMMiK: a generator of bespoke matrix multiplication kernels for block by panel type multiplications where the block matrix is constant. GiMMiK exploits a priori knowledge of this matrix to generate highly performant CUDA code for NVIDIA GPUs. The performance of GiMMiK kernels is particularly apparent when the matrix has some degree of sparsity. GiMMiK embeds matrix entries directly in the code and eliminates multiplies by zeros. Together with the ability of GiMMiK kernels to avoid poorly optimised cleanup code, GiMMiK is able to outperform cuBLAS on a variety of real-world problems. Speedups of 10 times are found on a K40c for a 294×1029 matrix with 99% sparsity. It is open source and released under a three clause BSD license.

GPU-Accelerated VLSI Routing Using Group Steiner Trees

Authors: Basileal Imana, Venkata Suhas Maringanti, Peter Yoon (Trinity College)

The problem of interconnecting nets with multi-port terminals in VLSI circuits is a direct generalization of the Group Steiner Problem (GSP). The GSP is a combinatorial optimization problem which arises in the routing phase of VLSI circuit design. This problem has been intractable, making it impractical to be used in real-world VLSI applications. This poster presents our initial work on designing and implementing a parallel approximation algorithm for the GSP based on an existing heuristic on a distributed architecture. Our implementation uses a CUDA-aware MPI-based approach to compute the approximate minimum-cost Group Steiner tree for several industry-standard VLSI graphs. Our implementation achieves up to 302x speedup compared to the best known serial work for the same graph. We present the speedup results for graphs up to 3K vertices. We also investigate some performance bottleneck issues by analyzing and interpreting the program performance data.

Improving Throughput by Dynamically Adapting Concurrence of Data Transfer

Authors: Prasanna Balaprakash, Vitali Morozov, Rajkumar Kettimuthu (Argonne National Laboratory)

Improving the throughput of data transfer over high-speed long-distance networks has become increasingly difficult and complex. Numerous factors such as varying congestion scenarios, external factors which are hard to characterize analytically, and dynamics of the underlying transfer protocol, contribute to this difficulty. In this study, we consider optimizing memory-to-memory transfer via TCP, where the data is transferred from a source memory to a destination memory using TCP. Inspired by the simplicity and the effectiveness of additive increase and multiplicative decrease scheme of TCP variants, we propose a tuning algorithm that can dynamically adapt the number of parallel TCP streams to improve the aggregate throughput of data transfers. We illustrate the effectiveness of the proposed algorithm on a controlled environment. Preliminary results show significant throughput improvement under congestion.

PERMON Toolbox Combining Discretization, Domain Decomposition, and Quadratic Programming

Authors: Vaclav Hapla, David Horak, Lukas Pospisil (IT4Innovations National Supercomputing Center)

Quadratic programming (QP) problems result from certain methods in image processing or particle dynamics, or finite element discretization of contact problems of mechanics. Domain decomposition methods solve an original large problem by splitting it into smaller subdomain problems that are almost

independent, allowing naturally massively parallel computations on supercomputers. We are interested in combining non-overlapping DDMs, namely FETI (Finite Element Tearing and Interconnecting), with optimal QP algorithms. FETI combines both iterative and direct solvers and allows highly accurate computations scaling up to tens of thousands of processors. Due to limitations of commercial packages, problems often have to be adapted to be solvable and it takes a long time before recent numerical methods needed for HPC are implemented into such packages. These issues lead us to establish the PERMON (Parallel, Efficient, Robust, Modular, Object-oriented, Numerical) toolbox.

Scaling Uncertainty Quantification Studies to Millions of Jobs

Authors: *Tamara L. Dahlgren, David Domyancic, Scott Brandon, Todd Gamblin, John Gyllenhaal, Rao Nimmakayala, Richard Klein (Lawrence Livermore National Laboratory)*

Computer simulations for assessing the likelihood of certain outcomes often involve hundreds to millions of executions and can take months to complete using current resources. Reducing the turn-around time by scaling concurrent simulation ensemble runs to millions of processors is hindered by resource constraints. Our poster describes a preliminary investigation of mitigating the impacts of these limitations in the LLNL Uncertainty Quantification Pipeline (UQP) using CRAM. CRAM virtualizes MPI by splitting an MPI program into separate groups, each with its own sub-communicator for running a simulation. We launched a single process problem on all 1.6 million Sequoia cores in under 40 minutes versus 4.5 days. The small problem size resulted in 400,000 concurrent ensemble simulations. Preparations are underway to demonstrate using CRAM to maximize the number of simulations we can run within our allocated partition for an ensemble using a multi-physics package.

Overcoming Distributed Debugging Challenges in the MPI+OpenMP Programming Model

Authors: *Lai Wei (Rice University); Ignacio LagunaDong H. Ahn, Matthew P. LeGendre, Gregory L. Lee (Lawrence Livermore National Laboratory)*

There is a general consensus that exascale computing will embrace a wider range of programming models to harness the many levels of architectural parallelism. To aid programmers in managing the complexities arising from multiple programming models, debugging tools must allow programmers to identify errors at the level of the programming model where the root cause of a bug was introduced. However, the question of what are the effective levels for debugging in hybrid distributed models remains unanswered. In this work, we share our lessons learned from incorporating OpenMP awareness into a highly-scalable, lightweight debugging tool

for MPI applications: the Stack Trace Analysis Tool (STAT). Our framework leverages OMPD, an emerging debugging interface for OpenMP, to provide easy-to-understand stack trace views for MPI+OpenMP programs. Our tool helps users debug their programs at the user code level by mapping the stack traces to the high-level abstractions provided by programming models.

Emulating In-Memory Data Rearrangement for HPC Applications

Authors: *Christopher W. Hajas (University of Florida); G. Scott Lloyd, Maya B. Gokhale (Lawrence Livermore National Laboratory)*

As bandwidth requirements for scientific applications continue to increase, new and novel memory architectures to support these applications are required. The Hybrid Memory Cube is a high-bandwidth memory architecture containing a logic layer with stacked DRAM. The logic layer aids the memory transactions; however, additional custom logic functions to perform near-memory computation are the subject of various research endeavors.

We propose a Data Rearrangement Engine in the logic layer to accelerate data-intensive, cache unfriendly applications containing irregular memory accesses by minimizing DRAM latency through the coalescing of disjoint memory accesses. Using a custom FPGA emulation framework, we found 1.4x speedup on a Sparse-Matrix, Dense-Vector benchmark (SpMV). We investigated the multi-dimensional parameter space to achieve maximum speedup and determine optimal cache invalidation and memory access coalescing schemes on various sizes and densities of matrices.

C++ Abstraction Layers – Performance, Portability and Productivity

Authors: *Dennis C. Dingo, Simon D. Hammond, Christian R. Trott, Harold C. Edwards (Sandia National Laboratories)*

Programming applications for next-generation supercomputing architectures is a critical and challenging goal for HPC developers. Developers may elect to use hardware-specific programming models to obtain the highest levels of performance possible while sacrificing portability, alternatively, they may choose more portable solutions such as directive-based models to gain portability often at the cost of cross-platform performance.

In this poster we present quantitative runtime performance and qualitative discussion of source code modifications for ports of the LULESH benchmark using the Kokkos C++ programming model and library. We show, when compared with the original OpenMP benchmark, we are able to demonstrate at least equivalent, if not greater, performance than the original code with approximately similar levels of code modification.

Crucially however, applications written in Kokkos are hardware portable and we are able to use a single application source to provide efficient execution on GPUs, many-core and multi-core devices.

STATuner: Efficient Tuning of CUDA Kernels Parameters

Authors: Ravi Gupta (Purdue University); Ignacio Laguna, Dong H. Ahn, Todd Gamblin (Lawrence Livermore National Laboratory); Saurabh Bagchi, Felix Xiaozhu Lin (Purdue University)

CUDA programmers need to decide the block size to use for a kernel launch that yields the lowest execution time. However, existing models to predict the best block size are not always accurate and involve a lot of manual effort from programmers. We identify a list of static metrics that can be used to characterize a kernel and build a Machine Learning model to predict block size that can be used in a kernel launch to minimize execution time. We use a set of kernels to train our model based on these identified static metrics and compare its predictions with the well-known NVIDIA tool called Occupancy Calculator on test kernels. Our model is able to predict block size that gives average error of 4.4% in comparison to Occupancy Calculator that gives error of 6.6%. Our model requires no trial runs of the kernel and less effort compared to Occupancy Calculator.

An Approach to the Highest Efficiency of the HPCG Benchmark on the SX-ACE Supercomputer

Authors: Kazuhiko Komatsu, Ryusuke Egawa (Tohoku University); Yoko Isobe, Ryusei Ogata (NEC Corporation); Hiroyuki Takizawa, Hiroaki Kobayashi (Tohoku University)

This paper reports how we have achieved the world's highest efficiency on the HPCG benchmark by fully exploiting the high potential of the SX-ACE supercomputer. To achieve efficient vector calculations and memory accesses, several optimization techniques such as various sparse matrix data packing, the multicolor ordering method and the hyperplane method for eliminating data dependencies, selective caching for the on-chip memory, and problem size tuning have been discussed and evaluated. Evaluation results clarify that these optimization techniques are beneficial to achieve high sustained performance of HPCG on SX-ACE since the memory access performance has a strong impact on the overall performance of HPCG. Furthermore, by employing the most effective combination of the optimization techniques, SX-ACE achieves an 11.4% efficiency in the case of 512 nodes, which is the highest efficiency among all of the supercomputers in the latest HPCG ranking.

Performance, Power, and Energy of In-Situ and Post-Processing Visualization: A Case Study in Climate Simulation

Authors: Vignesh Adhinarayanan (Virginia Polytechnic Institute and State University); Scott Pakin, David Rogers (Los Alamos National Laboratory); Wu-chun Feng (Virginia Polytechnic Institute and State University), James Ahrens (Los Alamos National Laboratory)

The energy consumption of off-chip data movement for an exascale system is projected to be two orders of magnitude higher than on-chip data movement. However, in an attempt to increase the fidelity of simulation, scientists increasingly run high-resolution simulations, producing large amounts of data for visualization and analysis. Consequently, the I/O subsystem is expected to consume a significant chunk of the power budget available for a supercomputer. To combat this problem, researchers have proposed in-situ techniques where the visualization is performed alongside the simulation using data residing in memory. This study quantifies the savings in performance, power, and energy from adopting in-situ visualization using a climate simulation application on a server-grade node. We offer deeper insights from subsystem-level power measurements and make projections for a larger cluster. We plan to complement these projections with actual measurements on a supercomputer in the final version of the poster.

Large Scale Artificial Neural Network Training Using Multi-GPUs

Authors: Linnan Wang (Georgia Institute of Technology), Wei Wu (University of Tennessee, Knoxville), Alex Zhang (Rutgers University New Brunswick), Jianxiong Xiao (Princeton University)

This paper describes a method for accelerating large scale Artificial Neural Networks (ANN) training using multi-GPUs by reducing the forward and backward passes to matrix multiplication. We propose an out-of-core multi-GPU matrix multiplication and integrate the algorithm with ANN. The experiments demonstrate that the proposed algorithm can achieve linear speedup with multiple inhomogeneous GPUs.

Simulating and Visualizing Traffic on the Dragonfly Network

Authors: Abhinav Bhatele (Lawrence Livermore National Laboratory), Nikhil Jain (University of Illinois at Urbana-Champaign), Yarden Livnat (University of Utah), Valerio Pascucci (University of Utah), Peer-Timo Bremer (Lawrence Livermore National Laboratory)

The dragonfly topology is becoming a popular choice for building high-radix, low-diameter networks with high-bandwidth links. Even with a powerful network, preliminary experiments on Edison at NERSC have shown that for communication heavy

applications, job interference and thus presumably job placement remains an important factor. In this poster, we explore the effects of job placement, job sizes, parallel workloads and network configurations on network throughput to better understand inter-job interference. We use a simulation tool called Damsely to model the network behavior of Edison and study the impact of various system parameters on network throughput. Parallel workloads based on five representative communication patterns are used and the simulation studies on up to 131,072 cores are aided by a new visualization of the dragonfly network.

Development of Explicit Moving Particle Simulation Framework and Zoom-Up Tsunami Analysis System

Authors: Kohei Murotani, Seiichi Koshizuka (University of Tokyo); Masao Ogino (Nagoya University), Ryuji Shioya (Tokyo University)

In this research, we are developing LexADV_EMPS which is the Explicit Moving Particle Simulation (MPS) framework. The LexADV_EMPS supports “Domain decomposition”, “Halo exchange” and “Dynamic load balance” as parallel computing functions for particle methods of the continuum mechanics. We have been able to solve large scale realistic tsunami analyses using distributed memory parallel Explicit MPS method implemented by the LexADV_EMPS. Today, Ishinomaki city, Kesennuma city and Fukushima Daiichi Nuclear Power Station have been successfully solved by our system using the K computer of RIKEN, FX10 of the University of Tokyo and CX400 of Nagoya University.

Reliable Performance Auto-Tuning in Presence of DVFS

Authors: Md Rakib Hasan (Louisiana State University); Eric Van Hensbergen, Wade Walker (ARM Research Labs, LLC)

In an era where exascale systems are imminent, maintaining a power budget for such systems is one of the most critical problems to overcome. Along with much research on balancing performance and power, Dynamic Voltage and Frequency Scaling (DVFS) is being used extensively to save idle-time CPU power consumption. The drawback is that the inherent random behavior of DVFS makes walltime unreliable to be used as a performance metric which causes random performance from libraries (e.g. ATLAS) that rely on machine-specific auto-tuning of several characteristics for the best performance. In this poster: 1) We show that a suboptimal selection (not the worst case) of kernel and block size during auto-tuning can cause ATLAS to lose 40% of DGEMM performance and 2) We present a more reliable performance metric in presence of DVFS that can estimate the same performance as no-DVFS yielding proper autotuning.

LIBXSMM: A High Performance Library for Small Matrix Multiplications

Authors: Alexander Heinecke, Hans Pabst, Greg Henry (Intel Corporation)

In this work we present a library, LIBXSMM, that provides a high-performance implementation of small sparse and dense matrix multiplications on latest Intel architectures. Such operations are important building blocks in modern scientific applications and general math libraries are normally tuned for all dimensions being large. LIBXSMM follows a matrix multiplication code generation approach specifically matching the applications’ needs. By providing several interfaces, the replacement of BLAS calls is simple and straightforward. We show that depending on the application’s characteristics, LIBXSMM can either leverage the entire DRAM bandwidth or reaches close to the processor’s computational peak performance. Our performance results of CP2K and SeisSol therefore demonstrate that using LIBXSMM as a highly-efficient computational backend, leads to speed-ups of greater than two compared to compiler generated inlined code or calling highly-optimized vendor math libraries.

A High-Performance Approach for Solution Space Traversal in Combinatorial Optimization

Authors: Wendy K. Tam Cho, Yan Y. Liu (University of Illinois at Urbana-Champaign)

Many interesting problems, including redistricting and causal inference models, have heroically large solutions spaces and can be framed within a large-scale optimization framework. Their solution spaces are vast, characterized by a series of expansive plateaus rather than a rapid succession of precipices. The expansive plateaus coupled with the sheer size of the solution space prove challenging for standard optimization methodologies. Standard techniques do not scale to problems where the number of possible solutions far exceeds a googol. We demonstrate how a high performance computing environment allows us to extract insights into these important problems. We utilize multiple processors to collaboratively hill climb, broadcast messages to one another about the landscape characteristics, and request aid in climbing particularly difficult peaks. Massively parallel architectures allow us to increase the efficiency of landscape traversal, allowing us to synthesize, characterize, and extract information from these massive landscapes.

A Performance Evaluation of Kokkos and RAJA using the TeaLeaf Mini-App

Authors: Matthew Martineau, Simon McIntosh-Smith (University of Bristol); Wayne Gaudin (Atomic Weapons Establishment), Mike Boulton (University of Bristol), David Beckingsale (Lawrence Livermore National Laboratory)

In this research project we have taken the TeaLeaf 'mini app' and developed several ports using a mixture of new and mature parallel programming models. Performance data collected on modern HPC devices demonstrates the capacity for each to achieve portable performance.

We have discovered that RAJA is a promising model with an intuitive development approach, that exhibits good performance on central processing units (CPUs), but does not currently support other devices. Kokkos is significantly more challenging to develop, but presents a highly competitive option for performance portability on CPUs and on NVIDIA general-purpose graphics processing units (GPUs). The results show that Kokkos can exhibit portable performance to within 5% of OpenMP and hand-optimised CUDA for some solvers, but currently performs poorly on the Intel Xeon Phi.

Our poster presents highlights of the results collected during this research to enable open discussion about the benefits of each model for application developers.

Efficient Large-Scale Sparse Eigenvalue Computations on Heterogeneous Hardware

Authors: Moritz Kreutzer (Friedrich-Alexander University Erlangen-Nürnberg); Andreas Pieper, Andreas Alvermann, Holger Fehske (Ernst Moritz Arndt University of Greifswald); Georg Hager, Gerhard Wellein (Friedrich-Alexander University Erlangen-Nürnberg)

In quantum physics it is often required to determine spectral properties of large, sparse matrices. For instance, an approximation to the full spectrum or a number of inner eigenvalues can be computed with algorithms based on the evaluation of Chebyshev polynomials. We identify relevant bottlenecks of this class of algorithms and develop a reformulated version to increase the computational intensity and obtain a potentially higher efficiency, basically by employing kernel fusion and vector blocking. The optimized algorithm requires a manual implementation of compute kernels. Guided by a performance model, we show the capabilities of our fully heterogeneous implementation on a petascale system. Based on MPI+OpenMP/CUDA, our approach utilizes all parts of a heterogeneous CPU+GPU system with high efficiency. Finally, our scaling study on up to 4096 heterogeneous nodes reveals a performance of half a petaflop/s, which corresponds to 11% of LINPACK performance for an originally bandwidth-bound sparse linear algebra problem.

Integrated Co-Design of Future Exascale Software

Authors: Bjoern Gmeiner (University Erlangen-Nuremberg); Markus Huber, Lorenz John (Technical University of Munich); Ulrich Ruede (University Erlangen-Nuremberg); Christian Waluga, Barbara Wohlmuth (Technical University of Munich); Holger Stengel (University Erlangen-Nuremberg)

The co-design of algorithms for the numerical approximation of partial differential equations is essential to exploit future exascale systems. Here, we focus on key attributes such as node performance, ultra scalable multigrid methods, scheduling techniques for uncertain data, and fault tolerant iterative solvers. In the case of a hard fault, we combine domain partitioning with highly scalable geometric multigrid schemes to obtain fast fault-robust solvers. The recovery strategy is based on a hierarchical hybrid concept where the values on lower dimensional primitives such as faces are stored redundantly and thus can be recovered easily. The lost volume unknowns are re-computed approximately by solving a local Dirichlet problem on the faulty subdomain. Different strategies are compared and evaluated with respect to performance, computational cost, and speed up. Locally accelerated strategies resulting in asynchronous multigrid iterations can fully compensate faults.

mrCUDA: Low-Overhead Middleware for Transparently Migrating CUDA Execution from Remote to Local GPUs

Authors: Pak Markthub, Akihiro Nomura, Satoshi Matsuoka (Tokyo Institute of Technology)

mrCUDA is a state-of-the-art remote CUDA execution middleware that enables CUDA applications running on one node to transparently use GPUs on other nodes. With this capability, applications can use nodes that do not have enough unoccupied GPUs by using mrCUDA to borrow idle GPUs from some other nodes. However, those applications may suffer from mrCUDA's overhead; especially for applications that frequently call CUDA kernels or have to transfer a lot of data, mrCUDA's overhead can be detrimentally large. We propose mrCUDA, a middleware for transparently live-migrating CUDA execution from remote to local GPUs, and show that mrCUDA's overhead is negligibly small compared with rCUDA's overhead. Hence, mrCUDA enables applications to run on nodes that do not have enough unoccupied GPUs (by using rCUDA) and later migrate the work to local GPUs (thus, get rid of rCUDA's overhead) when available.

Adapting Genome-Wide Association Workflows for HPC Processing at Pawsey

Authors: Jindan (Charlene) Yang, Christopher Harris (Pawsey Supercomputing Center); Sylvia Young, Grant Morahan (The University of Western Australia)

As part of Bioinformatics Petascale Pioneers Program at Pawsey Supercomputing Centre, this project adapts a number of genome-wide association workflows to a unique Cray

supercomputing environment, transforms the performance of these workflows and significantly boosts the research lifecycle. Pairwise and third order gene-gene interaction studies are sped up by hundreds of times through massive parallelization and GPU computing, laying the ground for larger datasets and higher order analyses in the future.

Out-of-Core Sorting Acceleration using GPU and Flash NVM

Authors: Hitoshi Sato, Ryo Mizote, Satoshi Matsuoka (Tokyo Institute of Technology)

We propose a sample-sort-based out-of-core sorting acceleration technique, called xtr2sort, that deals with multi-level memory hierarchy of GPU, CPU and Flash NVM, as an instance for future computing systems with deep memory hierarchy. Our approach splits the input records into several chunks to fit on GPU and overlaps I/O operations between Flash NVM and CPU, data transfers between CPU and GPU, and sorting on GPU in an asynchronous manner. Experimental results show that xtr2sort can sort up to 64 times larger record size than in-core GPU sorting and 4 times larger record size than in-core CPU sorting, and achieve 2.16 times faster than out-of-core CPU sorting using 72 threads, even the input records cannot fit on CPU and GPU. The results indicate that I/O chunking and latency hiding approach work really well for GPU and Flash NVM, and is a possible approach for future big data processing with extreme computing techniques.

Performance Comparison of the Multi-Zone Scalar Pentadiagonal (SP-MZ) NAS Parallel Benchmark on Many-Core Parallel Platforms

Authors: Christopher P. Stone (Computational Science and Engineering, LLC), Bracy Elton (Engility Corporation)

The NAS multi-zone scalar-pentadiagonal (SP-MZ) benchmark is representative of many CFD applications. Offloading this class of algorithm to many-core accelerator devices should boost application throughput and reduce time-to-solution. OpenACC and OpenMP compiler directives provide platform portability, hierarchical thread and vector parallelism, and simplified development for legacy applications. We examine the performance of the SP-MZ benchmark on clusters comprised of NVIDIA GPU and Intel Xeon Phi accelerators. We found that offloading the SP-MZ application to the accelerators was straightforward using the compiler directives. However, significant code restructuring was required to attain acceptable performance on the many-core accelerator devices. We implemented similar optimizations for the Intel Xeon Phi, via OpenMP, and the NVIDIA Kepler GPU, with OpenACC, in order to increase both thread and vector parallelism. We observed comparable performance between the two many-core accelerator devices and to HPC-grade multi-core host CPUs.

OPESCI: Open Performance portable Seismic Imaging

Authors: Marcos de Aguiar (SENAI CIMATEC), Gerard Gorman (Imperial College London), Renato Miceli (SENAI CIMATEC); Christian Jacobs, Michael Lange, Tianjiao Sun (Imperial College London); Felipe Zacarias (SENAI CIMATEC)

Entering the exascale era, HPC architectures are rapidly changing and diversifying, to continue delivering performance increases. These changes offer opportunities while also demanding disruptive software changes to harness the full hardware potential. The question is how to achieve an acceptable degree of performance portability across diverse, rapidly evolving architectures, in spite of the sharp trade-off between easy-to-maintain, portable software written in high-level languages, and highly optimized, parallel codes for target architectures. The solution proposed is to leverage domain-specific languages (DSL), and code generators to introduce multiple layers of abstraction. At the highest level, application developers write algorithms in clear, concise manner, while at the lower levels source-to-source compilers transforms DSL codes into highly optimized native codes that can be compiled for target platforms for near-to-peak performance. The framework provides layers, decoupling domain experts from code tuning specialists, where different optimized code generator back ends can be replaced.

Parallel Execution of Workflows Driven by a Distributed Database Management System

Authors: Renan Souza, Vitor Silva (Federal University of Rio de Janeiro), Daniel de Oliveira (Fluminense Federal University), Patrick Valduriez (French Institute for Research in Computer Science and Automation), Alexandre A. B. Lima, Marta Mattoso (Federal University of Rio de Janeiro)

Scientific Workflow Management Systems (SWfMS) that execute large-scale simulations need to manage many task computing in high performance environments. With the scale of tasks and processing cores to be managed, SWfMS require efficient distributed data structures to manage data related to scheduling, data movement and provenance data gathering. Although related systems store these data in multiple log files, some existing approaches store them using a Database Management System (DBMS) at runtime, which provides powerful analytical capabilities, such as execution monitoring, anticipated result analyses, and user steering. Despite these advantages, approaches relying on a centralized DBMS introduce a point of contention, jeopardizing performance in large-scale executions. In this paper, we propose an architecture relying on a distributed DBMS to both manage the parallel execution of tasks and store those data at runtime. Our experiments show an efficiency of over 80% on 1,000 cores without abdicating the analytical capabilities at runtime.

Using MuMMI to Model and Optimize Energy and Performance of HPC Applications on Power-Aware Supercomputers

Authors: Xingfu Wu (Texas A&M University), Valerie Taylor (Texas A&M University)

Hardware performance counters have been used as effective proxies to estimate power consumption and runtime. In this poster, we present revisions to the MuMMI tool for analyzing performance and power tradeoffs. The revisions entail adding techniques that allow the user to easily identify the high priority counters for application refinements. The performance models focus on four metrics: runtime, system power, CPU power and memory power. We rank the counters from these models to identify the most important counters for application optimization focus, then explore the counter-guided optimizations with an aerospace application PMLB executed on two power-aware supercomputers, Mira (at Argonne National Laboratory), and SystemG (at Virginia Tech). The counter-guided optimizations result in a reduction in energy by average 18.28% on up to 32,768 cores on Mira and average 11.28% on up to 128 cores on SystemG.

High Level Synthesis of SPARQL Queries

Authors: Marco Minutoli, Vito Giovanni Castellana, Antonino Tumeo (Pacific Northwest National Laboratory)

RDF databases naturally map to labeled, directed graphs. SPARQL is a query language for RDF databases that expresses queries as graph pattern matching operations. GEMS is a RDF database that, differently from other solutions, employs graph methods at all levels of its stack. Graph methods are inherently task parallel, but they exhibit an irregular behavior. In this poster we discuss an approach to accelerate GEMS with reconfigurable devices. The proposed approach automatically generates parallel hardware implementations of SPARQL queries using a customized High Level Synthesis (HLS) flow. The flow has been enhanced with solutions to address limitations of graph methods with conventional HLS methods, enhancing TLP extraction and management of concurrent memory operations. We have validated our approach by synthesizing and simulating seven queries from LUBM. We show that the proposed approach provides an average speed up of 2.1 with respect to the serial version of the hardware accelerators.

User Environment Tracking and Problem Detection with XALT

Authors: Kapil Agrawal, Gregory Peterson (University of Tennessee, Knoxville); Mark Fahey (Argonne National Laboratory), Robert McLay (The University of Texas at Austin)

This work enhances our understanding of individual users' software needs, then leverages that understanding to help

stakeholders conduct business in a more efficient, effective and systematic manner. XALT is designed to track linkage and execution information for applications that are compiled and executed on any Linux cluster, workstation, or high-end supercomputers. XALT allows administrators and other support staff to consider demand when prioritizing what to install, support and maintain. Datasets, dashboards, and historical reports generated by XALT and the systems with which it interoperates will preserve institutional knowledge and lessons learned so that users, developers, and support staff need not reinvent the wheel when issues that have already been encountered arise.

Task-Based Parallel Computation of the Density Matrix in Quantum-Based Molecular Dynamics Using Graph Partitioning

Authors: Sergio Pino, Matthew Kroonblawd, Purnima Ghale, Georg Hahn, Vivek Sardeshmukh, Guangjie Shi, Hristo Djidjev, Christian Negre, Robert Pavel, Benjamin Bergen, Susan Mniszewski, Christoph Junghans (Los Alamos National Laboratory)

Quantum molecular dynamics (QMD) simulations are highly accurate, but are computationally expensive due to the calculation of the ground-state electronic density matrix P via a $O(N^3)$ diagonalization. Second-order spectral projection (SP2) is an efficient $O(N)$ alternative to obtain P . We parallelize SP2 with a data parallel approach based on an undirected graph representation of P in which graph-partitioning techniques are used to divide the computation into smaller independent partitions. Undesirable load imbalances arise in standard MPI/OpenMP-based implementations as the partitions are generally not of equal size. Expressing the algorithm using task-based programming models allows us to approach the load-balancing problems by scheduling parallel computations at runtime. We develop CnC and Charm++ implementations that can be integrated into existing QMD codes. Our approach is applied to QMD simulations of representative biological protein systems with more than 10,000 atoms, exceeding size limitations of diagonalization by more than an order of magnitude.

Investigating Prefetch Potential on the Xeon Phi with Autotuning

Authors: Saami Rahman, Ziliang Zong, Apan Qasem (Texas State University)

Prefetching is a well-known technique that is used to hide memory latency. Modern compilers analyze the program and insert prefetch instructions in the compiled binary. The Intel C Compiler (ICC) allows the programmer to specify two parameters that can help the compiler insert more accurate and timely prefetch instructions. The two parameters are `-opt-prefetch` and `-opt-prefetch-distance`. When unspecified, ICC uses default heuristics. In this work, we present the results of autotuning the two mentioned parameters and its effect

on performance and energy. Choosing these parameters using analysis by hand can be challenging and time consuming as it requires knowledge of memory access patterns as well as significant time investment. We have developed a simple autotuning framework for the Xeon Phi architecture that automatically tunes these two parameters for any given program. We used the framework on 4 memory intensive programs and gained up to 1.47 speedup and 1.39 greenup.

Automating Sparse Linear Solver Selection with Lighthouse

Authors: Kanika Sood (University of Oregon), Pate Motter (University of Colorado Boulder), Elizabeth Jessup (University of Colorado Boulder), Boyana Norris (University of Oregon)

Solving large, sparse linear systems efficiently is a challenging problem in scientific computing. Taxonomies and high-performance numerical linear algebra solutions help to translate algorithms to software. However, accessible, comprehensive, and usable tools for high quality code production are not available. To address this challenge, we present an extensible methodology for classifying iterative algorithms for solving sparse linear systems. Lighthouse is the first framework that offers an organized taxonomy of software components for linear algebra that enables functionality and performance-based search and generates code templates and optimized low-level kernels. It enables the selection of a solution method that is likely to converge and perform well. We describe the integration of PETSc and Trilinos iterative solvers into Lighthouse. We present a comparative analysis of solver classification results for a varied set of input problems and machine learning methods achieving up to 93% accuracy in identifying the best-performing linear solution methods.

Geometric-Aware Partitioning on Large-Scale Data for Parallel Quad Meshing

Authors: Xin Li (Louisiana State University)

We present a partitioning algorithm to decompose complex 2D data into small simple subregions for effective parallel quad meshing. We formulate the partitioning problem for effective parallel quad meshing, which leads to an expensive quadratic integer optimization problem with linear constraints. Directly solving this problem is prohibitive for large-scale data partitioning. Hence, we suggest a more efficient two-step algorithm to obtain an approximate solution. First, we partition the region into a set of square-like cells using L_infinity Centroidal Voronoi Tessellation (CVT), then we solve a graph partitioning on the dual graph of this CVT to minimize the total boundary length of the partitioning, while enforcing the load balancing and each subregion's connectivity. With this geometry-aware decomposition, subregions are distributed to multiple processors for parallel quadrilateral mesh generation.

Analysis of Node Failures in High Performance Computers Based on System Logs

Authors: Siavash Ghiasvand (Technical University of Dresden), Florina M. Ciorba (University of Basel); Ronny Tschueter, Wolfgang E. Nagel (Technical University of Dresden)

The growth in size and complexity of HPC systems leads to a rapid increase of their failure rates. In the near future, it is expected that the mean time between failures of HPC systems becomes too short and that current failure recovery mechanisms will no longer be able to recover the systems from failures. Early failure detection is, thus, essential to prevent their destructive effects. Based on measurements of a production system at TU Dresden over an 8-month time period, we study the correlation of node failures in time and space. We infer possible types of correlations and show that in many cases the observed node failures are directly correlated. The significance of such a study is achieving a clearer understanding of correlations between observed node failures and enabling failure detection as early as possible. The results aimed to help system administrators minimize (or prevent) the destructive effects of failures.

Directive-Based Pipelining Extension for OpenMP

Authors: Xuewen Cui (Virginia Polytechnic Institute and State University); Thomas R. W. Scogland, Bronis R. de Supinski (Lawrence Livermore National Laboratory); Wu-chun Feng (Virginia Polytechnic Institute and State University)

Heterogeneity continues to increase in computing applications, with the rise of accelerators such as GPUs, FPGAs, APUs, and other co-processors. They have also become common in state-of-the-art supercomputers on the TOP500 list. Programming models, such as CUDA, OpenMP, OpenACC and OpenCL are designed to offload compute intensive workloads to co-processors efficiently. However, the naive offload model, synchronously copying and executing, in sequence is inefficient. However, pipelining these activities reduces programmability.

We propose an easy-to-use directive-based pipelining extension for OpenMP. Our extension offers a simple interface to overlap data transfer and kernel computation with an autotuning scheduler. We achieve performance improvements between 40% and 60% for a Lattice QCD application.

Matrices Over Runtime Systems @ Exascale

Authors: Emmanuel Agullo, Olivier Aumage (French Institute for Research in Computer Science and Automation); George Bosilca (University of Tennessee, Knoxville), B  renger Bramas (French Institute for Research in Computer Science and Automation), Alfredo Buttari (French National Center for Scientific Research and Toulouse Institute of Computer Science Research), Olivier Coulaud (French Institute for Research in Computer Science and Automation), Eric Darve (Stanford University), Jack Dongarra (University of Tennessee, Knoxville), Mathieu Faverge (French Institute for Research in Computer Science and Automation), Nathalie Furmento (French National Center for Scientific Research and Toulouse Institute of Computer Science Research), Luc Giraud (French Institute for Research in Computer Science and Automation), Abdou Guermouche (University of Bordeaux and French Institute for Research in Computer Science and Automation), Julien Langou (University of Colorado Denver), Florent Lopez (French National Center for Scientific Research and Toulouse Institute of Computer Science Research), Hatem Ltaief (King Abdullah University of Science & Technology); Samuel Pitoiset, Florent Pruvost, Marc Sergent (French Institute for Research in Computer Science and Automation); Samuel Thibault (University of Bordeaux and French Institute for Research in Computer Science and Automation), Stanimire Tomov (University of Tennessee, Knoxville)

The goal of the Matrices Over Runtime Systems @ Exascale (MORSE) project is to design dense and sparse linear algebra methods that achieve the fastest possible time to an accurate solution on large-scale multicore systems with GPU accelerators, using all the processing power that future high end systems can make available. We propose a framework for describing matrix algorithms with a sequential expression at a high level of abstraction and delegating the actual execution to a runtime system. In this poster we show that this model allows for (1) achieving an excellent scalability on heterogeneous clusters, (2) designing advanced numerical algorithms and (3) being compliant with standards such as OpenMP 4.0 or possible extensions of this standard. We illustrate our methodology on three classes of problems: dense linear algebra, sparse direct methods and fast multipole methods. The resulting codes have been incorporated into the Chameleon, qr_mumps and ScalFMM solvers, respectively.

Parallelization, Acceleration, and Advancement of Dissipative Particle Dynamics (DPD) Methods

Authors: Timothy I. Mattox, James P. Larentzos (Engility Corporation); Christopher P. Stone (Computational Science and Engineering, LLC), Sean Ziegeler (Engility Corporation), John K. Brennan (U.S. Army Research Laboratory), Martin L  sal (Institute of Chemical Process Fundamentals of the Academy of Sciences of the Czech Republic and J. E. Purkyne University)

Our group has developed micro- and mesoscale modeling capabilities necessary to represent salient physical and chemical features of material microstructure. We built novel simulation tools that incorporate coarse-grain (CG) models upscaled from quantum-based models, which are then coupled with continuum level approaches. We further developed a suite of discrete-particle modeling tools, based upon the Dissipative Particle Dynamics (DPD) method, for simulation of materials at isothermal, isobaric, isoenergetic, and isoenthalpic conditions. In addition to the ability to simulate at those various conditions, our method has a particularly unique capability to model chemical reactivity within the CG particles of the DPD simulation. We recently integrated these methods into LAMMPS, enabling the utilization of HPC resources to model many DoD-relevant materials at previously impractical time and spatial scales for unprecedented simulations of phenomena not previously possible. Our poster presents the parallelization, acceleration, and advancement of these DPD methods within the LAMMPS code.

Reduced-Precision Floating-Point Analysis

Authors: Michael O. Lam (James Madison University), Jeffrey K. Hollingsworth (University of Maryland)

Floating-point computation is ubiquitous in scientific computing, but rounding error can compromise the results of extended calculations. In previous work, we presented techniques for automated mixed-precision analysis and configuration. We now present new techniques that use binary instrumentation and modification to do fine-grained floating-point precision analysis, simulating any level of precision less than or equal to the precision of the original program. These techniques have lower overhead and provide more general insights than previous mixed-precision analyses. We also present a novel histogram-based visualization of a program's floating-point precision sensitivity, and an incremental search technique that gives the user more control over the precision analysis process.

GPU Acceleration of a Non-Hydrostatic Ocean Model Using a Mixed Precision Multigrid Preconditioned Conjugate Gradient Method

Authors: Takateru Yamagishi (Research Organization for Information Science and Technology), Yoshimasa Matsumura (Hokkaido University)

To meet the demand for fast and detailed calculations in numerical ocean simulation, we implemented a non-hydrostatic ocean model for the GPU, using basic and essential methods such as exploitation of large-scale parallelism in threads and instructions, coalesced access to the global memory, and minimization of memory transfer between the CPU host and GPU device. We also studied and evaluated the application of mixed precision calculation to the multigrid preconditioning of a Poisson/Helmholtz solver, which would deteriorate in performance due to the small number of threads in the multigrid method. The GPU-implemented model ran 4.3 times faster on NVIDIA K20C than on Fujitsu SPARC64 VIIIfx. The application of mixed precision achieved a 16% acceleration of the Poisson/Helmholtz solver, which is consistent with the decrease of global memory transfer due to the switch to mixed precision. No specific errors were found in the outputs.

Parallelization of Tsunami Simulation on CPU, GPU and FPGAs

Authors: Fumiya Kono, Naohito Nakasato, Kensaku Hayashi, Alexander Vazhenin, Stanislav Sedukhin (University of Aizu); Kohei Nagasu, Kentaro Sano (Tohoku University); Vasily Titov (National Oceanic and Atmospheric Administration)

Tsunamis are known to be among the most serious disasters. MOST (Method of Splitting Tsunami) is a numerical solver for modeling tsunami waves. Prediction of the arrival time of a tsunami is critical to evacuate people from coastal area. Therefore, fast computation of MOST enabled by parallel processing is important. We have developed a tsunami propagation code based on MOST and implemented different algorithms for parallelization by OpenMP and OpenACC. We have benchmarked these parallelized codes on various architectures such as multi-core CPU systems, Many Integrated Core architecture, and GPU. In this poster, we compare the performance of various parallel implementations. We found that a parallelized code that applied spatial blocking on OpenACC demonstrated the best performance at present. Concurrently, we are developing an accelerator of the MOST computation with Field-Programmable Gate Arrays (FPGAs) for high-performance and power-efficient simulation. We also present preliminary evaluation of its prototype implementation with a 28nm FPGA.

Large-Scale Ultrasound Simulations with Local Fourier Basis Decomposition

Authors: Jiri Jaros, Matej Dohnal (Brno University of Technology); Bradley E. Treeby (University College London)

The simulation of ultrasound wave propagation through biological tissue has a wide range of practical applications including planning therapeutic ultrasound treatments. The major challenge is to ensure the ultrasound focus is accurately placed at the desired target since the surrounding tissue can significantly distort it. Performing accurate ultrasound simulations, however, requires the simulation code to be able to exploit thousands of processor cores to deliver treatment plans within 24 hours.

This poster presents a novel domain decomposition based on the Local Fourier basis that reduces the communication overhead by replacing global all-to-all communications with the local nearest-neighbour communication patterns yet exploiting the benefits of spectral methods. The method was implemented using both pure-MPI and hybrid OpenMP/MPI approaches and the performance investigated on the SuperMUC cluster. We managed to increase the number of usable computer cores from 512 up to 8192 while reducing the simulation time by a factor of 8.55.

libSkylark: A Framework for High-Performance Matrix Sketching for Statistical Computing

Authors: Georgios Kollias, Yves Ineichen, Haim Avron (IBM Corporation); Vikas Sindhwani (Google); Ken Clarkson, Costas Bekas, Alessandro Curioni (IBM Corporation)

Matrix-based operations lie at the heart of many tasks in machine learning and statistics. Sketching the corresponding matrices is a way to compress them while preserving their key properties. This translates to dramatic reductions in execution time when the tasks are performed over the sketched matrices, while at the same time retaining provable bounds within practical approximation brackets. libSkylark is a high-performance framework enabling the sketching of potentially huge, distributed matrices and then applying the machinery of associated statistical computing flows. Sketching typically involves projections on randomized directions computed in parallel. libSkylark integrates state-of-the-art parallel pseudorandom number generators and their lazily computed streams with communication-minimization techniques for applying them on distributed matrix objects and then chaining the output into distributed numerical linear algebra and machine learning kernels. Scalability results for the sketching layer and example applications of our framework in natural language processing and speech recognition are presented.

Heuristic Dynamic Load Balancing Algorithm Applied to the Fragment Molecular Orbital Method

Authors: Yuri Alexeev, Prasanna Balaprakash (Argonne National Laboratory)

Load balancing for large-scale systems is an important NP-hard difficulty problem. We propose a heuristic dynamic load-balancing algorithm (HDLB), employing covariance matrix adaptation evolution strategy (CMA-ES), a state-of-the-art heuristic algorithm, as an alternative to the default dynamic load balancing (DLB) and previously developed heuristic static load balancing algorithms (HSLB). The problem of allocating CPU cores to tasks is formulated as an integer nonlinear optimization problem, which is solved by using an optimization solver. On 16,384 cores of Blue Gene/Q, we achieved an excellent performance of HDLB compared to the default load balancer for an execution of the fragment molecular orbital method applied to model protein system quantum-mechanically. HDLB is shown to outperform default load balancing by at least a factor of 2, thus motivating the use of this approach on other coarse-grained applications.

SLAP: Making a Case for the Low-Powered Cluster by Leveraging Mobile Processors

Authors: Dukyun Nam, Jik-Soo Kim, Hoon Ryu, Gibeom Gu, Chan Yeol Park (Korea Institute of Science and Technology Information)

In this poster, we present our empirical study of building a low-powered cluster called SLAP by leveraging mobile processors. To investigate the reliability and usability of our system, we have conducted various performance analyses based on HPL benchmark, a real semiconductor engineering application and the application of distributed file system. Our experience shows potential benefits, possibilities and limits of applying energy efficient alternative solutions to supercomputing which can result in many interesting research issues.

Characterizing Memory Throttling Using the Roofline Model

Authors: Bo Li (Virginia Polytechnic Institute and State University), Edgar A. Leon (Lawrence Livermore National Laboratory), Kirk W. Cameron (Virginia Polytechnic Institute and State University)

Memory bandwidth is one of the most important factors contributing to the performance of many HPC applications. Characterizing their sensitivity to this resource may help application and system developers understand the performance tradeoffs when running on multiple systems with different memory characteristics. In this work, we use memory throttling as a means to analyze the impact of the memory system on the performance of applications. We make two contributions. First, we characterize memory throttling in terms of the

roofline model of performance. This shows that with throttling we can emulate different memory systems. Second, by identifying the pattern between memory bandwidth utilization of an application and the maximum amount of throttling without affection performance (optimal), we propose an accurate model for predicting the optimal throttling for a given code region. This model can be employed to balance power between components on a power-limited system.

Argo: An Exascale Operating System and Runtime

Authors: Swann Perarnau, Rinku Gupta, Pete Beckman (Argonne National Laboratory)

New computing technologies are expected to change the high-performance computing landscape dramatically. Future exascale systems will comprise hundreds of thousands of compute nodes linked by complex networks. Compute nodes are expected to host both general-purpose and special-purpose processors or accelerators, with more complex memory hierarchies. At these scales the HPC community expects that we will also require new programming models, to take advantage of both intra-node and inter-node parallelism.

In this context, the Argo Project is developing a new operating system and runtime for exascale machines. It is designed from the ground up to run future HPC application at extreme scales. At the heart of the project are four key innovations: dynamic reconfiguring of node resources in response to workload changes, allowance for massive concurrency, a hierarchical framework for management of nodes, and a cross-layer communication infrastructure that allows resource managers and optimizers to communicate efficiently across the platform.

Cost Effective Programmable H/W Based Data Plane Acceleration: Linking PCI-Express Commodity I/O H/W with FPGAs

Authors: Woong Shin, Heon Y. Yeom (Seoul National University)

In this poster, we present a cost effective FPGA-based data plane design which reuses existing I/O devices such as 10GbE NICs or NVM-e SSDs as data ingress and egress ports. We achieved this by building a FPGA based device driver logic which is capable of exploiting PCI-e point to point communication. FPGA H/W design support such as C based High Level Synthesis tools enabled us to implement complex device drivers within FPGAs. Our design avoids re-implementing the performance and stability of existing ASIC based commodity I/O devices, already installed in our systems, thus reducing data plane implementation costs.

Verification of Resilient Communication Models for the Simulation of a Highly Adaptive Energy-Efficient Computer

Authors: Mario Bielert, Kim Feldhoff (Dresden University of Technology); Florina M. Ciorba (University of Basel); Stefan Pfennig, Elke Franz, Thomas Ilsche, Wolfgang E. Nagel (Dresden University of Technology)

Delivering high performance in an energy-efficient manner is of great importance in conducting research in computational sciences and in daily use of technology. From a computing perspective, a novel concept (the HAEC Box) has been proposed that utilizes innovative ideas of optical and wireless chip-to-chip communication to allow a new level of runtime adaptivity for future computers, which is required to achieve high performance and energy efficiency. HAEC-SIM is an integrated simulation environment designed for the study of the performance and energy costs of the HAEC Box running communication-intensive applications. In this work, we conduct a verification of the implementation of three resilient communication models in HAEC-SIM. The verification involves two NAS Parallel Benchmarks and their simulated execution on a 3D torus system with 16x16x16 nodes with Infiniband links. The simulation results are consistent with those of an independent implementation. Thus, the HAEC-SIM based simulations are accurate in this regard.

Network-Attached Accelerators: Host-Independent Accelerators for Future HPC Systems

Authors: Sarah Marie Neuwirth, Dirk Frey, Ulrich Bruening (University of Heidelberg)

The emergence of accelerator technology in current supercomputing systems is changing the landscape of supercomputing architectures. Accelerators like GPGPUs and coprocessors are optimized for parallel computation while being more energy efficient. Today's accelerators come with some limitations. They require a local host CPU to configure and operate them. This limits the number of accelerators per host. Another problem is the unbalanced communication between distributed accelerators. Network-attached accelerators are an architectural approach for scaling the number of accelerators and host CPUs independently. The design enables remote initialization, control of the accelerator devices, and host-independent accelerator-to-accelerator direct communication. Workloads can be dynamically assigned to CPUs and accelerators at run-time in an N to M ratio. An operative prototype implementation, based on the Intel Xeon Phi coprocessor and the EXTOLL NIC, is used to evaluate the latency, bandwidth, performance of the MPI communication, and communication time of the LAMMPS molecular dynamics simulator.

Exploring Asynchronous Many-Task Runtime Systems Toward Extreme Scales

Authors: Samuel Knight (Florida Institute of Technology), Marc Gamell (Rutgers University), Gavin Baker (California Polytechnic State University); David Hollman, Gregory Sjaadema, Keita Teranishi, Hemanth Kolla, Jeremiah Wilke, Janine C. Bennett (Sandia National Laboratories)

Major exascale computing reports indicate a number of software challenges to meet the drastic change of system architectures in the near future. While several-orders-of-magnitude increase in parallelism is the most commonly cited of those, hurdles also include performance heterogeneity of compute nodes across the system, increased imbalance between computational capacity and I/O capabilities, frequent system interrupts, and complex hardware architectures. Asynchronous, task-parallel programming models show great promise in addressing these issues, but are not yet fully understood nor developed sufficiently for computational science and engineering application codes. We address these knowledge gaps through quantitative and qualitative exploration of leading candidate solutions in the context of engineering applications at Sandia. In this poster, we evaluate MiniAero code ported to three leading candidate programming models (Charm++, UIN-TAH, and Legion) to examine both the feasibility and tradeoffs of a technical solution that permits insertion of new programming model elements into an existing code base.

Memory Hotplug for Energy Savings of HPC systems

Authors: Shinobu Miwa, Hiroki Honda (University of Electro-Communications)

Many supercomputer centers need energy savings by their systems. For increased energy savings, reducing the standby power of hardware devices is indispensable because many devices in an HPC system are unused for long periods. Although many techniques like per-core power gating have succeeded in reduction of CPU standby power, power savings of memories is still modest regardless of its considerable amount of standby power. To bridge this gap, we propose exploiting memory hotplug for power savings of memories within HPC systems. Memory hotplug, which is supported by Linux kernel 3.9 or later, is to plug and pull DIMMs in a running computer system. With this state-of-the-art technology, this poster proposes a technique to control memory size, which is referred to on-demand memory hot-add. Our experiment with HPCC benchmarks demonstrates that the proposed technique allows a server node to turn off many slices of memories.

Inverse Modeling Nanostructures from X-Ray Scattering Data through Massive Parallelism

Authors: Abhinav Sarje, Dinesh Kumar, Singanallur Venkatakrishnan, Slim Chourou, Xiaoye S. Li, Alexander Hexemer (Lawrence Berkeley National Laboratory)

We consider the problem of reconstructing material nanostructures from grazing-incidence small-angle X-ray scattering (GISAXS) data obtained through experiments at synchrotron light-sources. This is an important tool for characterization of macromolecules and nano-particle systems applicable to applications such as design of energy-relevant nano-devices. Computational analysis of experimentally collected scattering data has been the primary bottleneck in this process. We exploit the availability of massive parallelism in leadership-class supercomputers with multi-core and graphics processors to realize the compute-intensive reconstruction process. To develop a solution, we employ various optimization algorithms including gradient-based LMVM, derivative-free trust region-based POUNDERs, and particle swarm optimization, and apply these in a massively parallel fashion. We compare their performance in terms of both quality of solution and computational speed. We demonstrate the effective utilization of up to 8,000 GPU nodes of the Titan supercomputer for inverse modeling of organic-photovoltaics (OPVs) in less than 15 minutes.

Accelerating Tridiagonal Matrix Inversion on the GPU

Authors: Bemnet Demere, Peter Yoon, Ebenezer Hormenou (Trinity College)

Inverting a matrix is a more computationally challenging process than solving a linear system. However, in fields such as structural engineering, dynamic systems, and cryptography, computing the inverse of a matrix is inevitable. In this poster, we present an accelerated procedure for computing the inverse of diagonally dominant tridiagonal matrices on the GPU. The algorithm is based on the recursive application of the Sherman-Morrison formula for tridiagonal matrices. The preliminary experimental results on NVIDIA Tesla K20c GPUs show that our GPU implementation of the inversion procedure outperforms the conventional CPU-based implementations with a speedup of up to 24x.

Quantifying Productivity---Toward Development Effort Estimation in HPC

Authors: Sandra Wienke, Tim Cramer, Matthias S. Müller (RWTH Aachen University), Martin Schulz (Lawrence Livermore National Laboratory)

With increasing expenses for future HPC centers, the need to look at their productivity, defined as amount of science per total cost of ownership (TCO), grows. This includes development costs which arise from development effort spent to parallelize, tune or port an application to a certain architecture. Develop-

ment effort estimation is popular in software engineering, but cannot be applied directly to (non-commercial) HPC setups due to their particular target of performance. In our work-in-progress, we illustrate a methodology to qualify and quantify development effort parameters and hence how to estimate development effort and productivity. Here, the main challenge is to account for the numerous impact factors on development effort. We show preliminary results for two case studies: Questionnaires reveal development effort parameters with high impact and statistical tests help us to derive further details (here: comparing programming models). Additionally, we provide an online survey to engage the HPC community.

Optimizing CUDA Shared Memory Usage

Authors: Shuang Gao, Gregory D. Peterson (University of Tennessee, Knoxville)

CUDA shared memory is fast on-chip storage. However, the bank conflict issue could cause a performance bottleneck. Current NVIDIA Tesla GPUs support memory bank accesses with configurable bit-widths. While this feature provides an efficient bank mapping scheme for 32-bit and 64-bit data types, it becomes trickier to solve the bank conflict problem through manual code tuning. This paper presents a framework for automatic bank conflict analysis and optimization. Given static array access information, we calculate the conflict degree, and then provide optimized data access patterns. Basically, by searching among different combinations of inter- and intra-array padding, along with bank access bit-width configurations, we can efficiently reduce or eliminate bank conflicts. From RODINIA and the CUDA SDK we selected 13 kernels with bottlenecks due to shared memory bank conflicts. After using our approach, these benchmarks achieve 5%-35% improvement in runtime.

Optimization of Stencil-Based Fusion Kernels on Tera-Flops Many-Core Architectures

Authors: Yuuichi Asahi (Japan Atomic Energy Agency), Guillaume Latu (French Alternative Energies and Atomic Energy Commission); Takuya Ina, Yasuhiro Idomura (Japan Atomic Energy Agency); Virginie Grandgirard, Xavier Garbet (French Alternative Energies and Atomic Energy Commission)

Plasma turbulence is of great importance in fusion science. However, turbulence simulations are costly so that more computing power is needed for simulating the coming fusion device ITER. This motivates us to develop advanced algorithms and optimizations of fusion codes on tera-flops many-core architectures like Xeon Phi, NVIDIA Tesla and Fujitsu FX100. We evaluate the kernel performance extracted from the hot spots of fusion codes, GYSELA and GT5D, with different numerical algorithms.

For the former kernel, which is based on a semi-Lagrangian scheme with high arithmetic intensity, high-performance is obtained on accelerators (Xeon Phi and Tesla) by applying SIMD optimization. On the other hand, in the latter kernel, which is based on a finite difference scheme, a large shared cache plays a critical role in improving arithmetic intensity, and thus, multi-core CPUs (FX100) with a large shared cache give higher performance.

GPU-STREAM: Benchmarking the Achievable Memory Bandwidth of Graphics Processing Units

Authors: Tom Deakin, Simon McIntosh-Smith (University of Bristol)

Many scientific codes consist of memory bandwidth bound kernels--- the dominating factor of the runtime is the speed at which data can be loaded from memory into the Arithmetic Logic Units. Generally Programmable Graphics Processing Units (GPUs) and other accelerator devices such as the Intel Xeon Phi offer an increased memory bandwidth over CPU architectures. However, as with CPUs, the peak memory bandwidth is often unachievable in practice and so benchmarks are required to measure a practical upper bound on expected performance.

We present GPU-STREAM as an auxiliary tool to the standard STREAM benchmark to provide cross-platform comparable results of achievable memory bandwidth between multi- and many-core devices.

Our poster will present the cross-platform validity of these claims, and also a short quantification on the effect of ECC memory on memory bandwidth.

Optimization of an Ocean Model Using Performance Tools

Authors: Oriol Tintó Prims, Miguel Castrillo, Harald Servat, German Llorca, Kim Serradell (Barcelona Supercomputing Center); Oriol Mula-Valls (Catalan Institute of Climate Sciences); Francisco Doblas-Reyes (Barcelona Supercomputing Center)

Do the model developers know about the computational performance of the models they are developing? The landscape of the computational science has a lot of examples of scientists using huge amounts of computational resources to run models without worrying too much about computational performance. The proper use of these resources has too often been the last of the modelers' concerns, and therefore it is not strange if someone that finally decides to study the computational performance of his model finds out that it is far from being good. This is not at all surprising when one realizes that scientific applications are usually old and usually based on inherited older code developed by scientists without a proper background in computer science. In this work, we would like to

show that a performance analysis of scientific models is useful and, without any doubt, is worth the effort.

Parallel Cardiac Electrophysiology Modeling Framework

Authors: Jacob Pollack, Xiaopeng Zhao, Kwai Wong (University of Tennessee, Knoxville)

Cardiovascular diseases are the leading cause of death worldwide. Using computer simulations to accurately model the dynamics of the heart provides a platform for developing methods to better detect and treat heart disease. We present a heart electrophysiology modeling framework designed to run on HPC platforms. The model is capable of simulating the generation and propagation of electrical signals throughout the heart under a variety of circumstances.

The action potential of each simulated cell can be described using a variety of different models. However, prior implementations of the cardiac modeling framework only supported the use of the Beeler-Reuter model. Our work refactoring the cardiac model allows the model to be extended to a multitude of electrical models, including the O'Hara-Rudy model. This expansion of functionality dramatically increases the simulation's usefulness, as many applications require the use of novel or complex electrical models.

Performing Large Scale Density Functional Theory Calculations using a Subspace Iterative Eigensolver

Authors: Sachin Nanavati (RWTH Aachen University), Mario Berljafa (The University of Manchester); Daniel Wortmann, Edoardo Di Napoli (Juelich Research Center)

One of the most accurate technique within the framework of Density Functional Theory is the Linearized Augmented Plane Wave (LAPW) method. Typically, the computationally intensive part of LAPW-based simulations, involve the solution of sequences of dense generalized eigenvalue problems. Traditionally these problems are solved in parallel using direct eigensolvers from standard libraries like ScaLAPACK. Here we introduce an eigensolver, based on subspace iteration method accelerated with Chebyshev polynomials of optimal degree (ChASE). ChASE exploits the correlation between successive eigenvectors of the self-consistent field sequence. Capitalizing on the frequent use of BLAS3 subroutines, the results obtained by ChASE are consistently competitive with direct eigensolvers. In the present work, we illustrate numerical results from the integration of distributed and shared memory parallelization of ChASE within FLEUR, a LAPW based software. Our results show that ChASE improves FLEUR's scaling behavior and performance for calculations of large physical system on modern supercomputers.

A Standard Debug Interface for OpenMP Target Regions

Authors: Andreas Erik Hindborg (Technical University of Denmark), Ignacio Laguna (Lawrence Livermore National Laboratory), Sven Karlsson (Technical University of Denmark), Dong H. Ahn (Lawrence Livermore National Laboratory)

As OpenMP covers more on-node parallelism, debugging OpenMP programs is becoming challenging. Debuggers often debug OpenMP at a system-thread level, leaving programmers with limited state information on OpenMP threads. The OpenMP debugging interface, OMPD, has been proposed as a standard interface to overcome these limitations. However, the current specification works only for OpenMP 3.1 -- it cannot handle OpenMP 4.0 programs because it does not consider OpenMP target constructs. In this work, we study how to extend OMPD to OpenMP 4.0 and propose a set of OMPD extensions. We identify crucial scalability issues with the way OMPD currently represents threads and propose a new representation, Construction of Online Lightweight Thread handle, COLT, as a scalable solution. Our evaluation shows that our scheme is feasible as far as basic OMPD support is provided for OpenMP 4 runtime component and can also significantly reduce the memory and performance overheads of the debuggers.

Advanced Tiling Techniques for Memory-Starved Streaming Numerical Kernels

Authors: Tareq Malas (King Abdullah University of Science & Technology), Georg Hager (Erlangen Regional Computing Center); Hatem Ltaief, David Keyes (King Abdullah University of Science & Technology)

Many temporal blocking techniques for stencil algorithms have been suggested for speeding up memory-bound code via improved temporal locality. Most of the established work concentrates on updating separate cache blocks per thread, which works on all types of shared memory systems, regardless of whether there is a shared cache among the cores or not. The downside of this approach is that the cache space for each thread can become too small for accommodating a sufficient number of updates and eventually decouple from memory bandwidth. We introduce a generalized multi-dimensional intra-tile parallelism scheme for shared-cache multicore processors that results in a significant reduction of cache size requirements. It ensures data access patterns that allow optimal hardware prefetching and TLB utilization. In this poster we describe the approach and some implementation details. We also conduct a performance comparison with the state-of-the-art stencil frameworks PLUTO and Pochoir

MINIO: An I/O Benchmark for Investigating High Level Parallel Libraries

Authors: James Dickson (University of Warwick), Satheesh Maheswaran (Atomic Weapons Establishment), Steven Wright (University of Warwick), Andy Herdman (Atomic Weapons Establishment), Stephen Jarvis (University of Warwick)

Input/output (I/O) operations are among the biggest challenges facing scientific computing as it transitions to exascale. The traditional software stack---comprising parallel file systems, middleware, and high level libraries -- has evolved to enable applications to better cope with the demands of enormous datasets. This software stack makes high performance parallel I/O easily accessible to application engineers. However, it is important to ensure best performance is not compromised through attempts to enrich these libraries.

We present MINIO, a benchmark for the investigation of I/O behaviour focusing on understanding overheads and inefficiencies in high level library usage. MINIO uses HDF5 and TyphonIO to explore I/O at scale using different application behavioural patterns.

A case study is performed using MINIO to identify performance limiting characteristics present in the TyphonIO library as an example of performance discrepancies in the I/O stack.

Energy-Efficient Graph Traversal on Integrated CPU-GPU Architecture

Authors: Heng Lin, Jidong Zhai, Wenguang Chen (Tsinghua University)

Recently, architecture designers tend to integrate CPUs and GPUs on the same chip to produce energy-efficient designs. On the other hand, graph applications are becoming increasingly important for big data analysis. Among the graph analysis algorithms, BFS is the most representative one and also an important building block for other algorithms. Despite previous efforts, it remains an important problem to get optimal performance for BFS on integrated architectures.

In this poster, we propose an adaptive algorithm to atomically find the optimal algorithm on the suitable devices. Our algorithm can get 1.6x speedup compared with the state-of-the-art algorithms in energy consumption.

Comparison of Virtualization and Containerization Techniques for High-Performance Computing

Authors: Yuyu Zhou (University of Pittsburgh); Balaji Subramaniam, Kate Keahey (Argonne National Laboratory); John Lange (University of Pittsburgh)

HPC users have traditionally used dedicated clusters hosted in national laboratories and universities to run their scientific applications. Recently, the use of cloud computing for such scientific applications has become popular, as exemplified by Amazon providing HPC instances. Nonetheless, HPC users have approached cloud computing cautiously due to performance overhead associated with virtualization and interference brought by virtual instances co-location. Recent improvements and developments in virtualization and containerization have alleviated some of these concerns regarding performance. However, the applicability of such technologies to HPC applications has not yet been thoroughly explored. Furthermore, scalability of scientific applications in the context of virtualized or containerized environments is not well studied. In this poster, we seek to understand the applicability of virtualization (exemplified by KVM) and containerization (exemplified by Docker) technologies to HPC applications.

Rendezview: An Interactive Visual Mining Tool for Discerning Flock Relationships in Social Media Data

Authors: Melissa J. Bica (University of Colorado Boulder), Kyoung-Sook Kim (National Institute of Advanced Industrial Science and Technology)

Social media data provides insight into people's opinions, thoughts, and reactions about real-world events. However, this data is often analyzed at a shallow level with simple visual representations, making much of this insight undiscoverable. Our approach to this problem was to create a framework for visual data mining that enables users to find implicit patterns and relationships within their data, focusing particularly on flock phenomena in social media. Rendezview is an interactive visualization framework that consists of three visual components: a spatiotemporal 3D map, a word cloud, and a Sankey diagram. These components provide individual functions for data exploration and interoperate with each other based on user interaction. The current version of Rendezview represents local topics and their co-occurrence relationships from geo-tagged Twitter messages. This work will be presented by focusing on visualizations, information representation, and interactions, and a live demo will be available to show the system in use.

A Deadlock Detection Concept for OpenMP Tasks and Fully Hybrid MPI-OpenMP Applications

Authors: Tobias Hilbrich (Dresden University of Technology), Bronis R. de Supinski (Lawrence Livermore National Laboratory); Andreas Knuepfer, Robert Dietrich (Dresden University of Technology); Christian Terboven, Felix Muenchhalfen (RWTH Aachen University); Wolfgang E. Nagel (Dresden University of Technology)

Current high-performance computing applications often combine the Message Passing Interface (MPI) with threaded parallel programming paradigms such as OpenMP. MPI allows fully hybrid applications in which multiple threads of a process issue MPI operations concurrently. Little study on deadlock conditions for this combined use exists. We propose a wait-for graph approach to understand and detect deadlock for such fully hybrid applications. It specifically considers OpenMP 3.0 tasking support to incorporate OpenMP's task-based execution model. Our model creates dependencies with deadlock criteria that can be visualized to support comprehensive deadlock reports. We use a model checking approach to investigate wide ranges of valid execution states of example programs to verify the soundness of our wait-for graph construction.

Design and Modelling of Cloud-Based Burst Buffers

Authors: Tianqi Xu (Tokyo Institute of Technology), Kento Sato (Lawrence Livermore National Laboratory), Satoshi Matsuoka (Tokyo Institute of Technology)

With the growing of data size, public clouds have been gathering more and more interest due to their capabilities of public big data processing. However, applications running on cloud are suffering from low I/O bandwidth as well as a loose consistency model in shared cloud storages. In previous work, we have proposed cloud-based burst buffers (CloudBB) to accelerate the I/O and strengthen the consistency while using shared cloud storages. In this work, we introduce the performance models to predict the performance and help users to determine the optimal configuration while using our system. We focus on two aspects: execution time and cost. Our model predicts the optimal configuration according to characteristics of applications and execution environment. We validate our model using a real HPC application on a real public cloud, Amazon EC2/S3. The results show that our model can predict the performance and help users to determine the optimal configuration.

Towards Scalable Graph Analytics on Time Dependent Graphs

Authors: Suraj Poudel (University of Alabama at Birmingham); Roger Pearce, Maya Gokhale (Lawrence Livermore National Laboratory)

The objective of this study is to annotate a distributed memory scalefree graph with temporal metadata to model a time dependent graph and to perform time dependent graph analytics. In our approach, each edge can have multiple metadata fields and each metadata field encodes temporal information. We performed experiments with large scale CAIDA datasets (anonymized internet traces) on a Linux HPC cluster (Catalyst of LLNL) by deriving flow between two IP's and annotating the flow information as metadata into the graph topology using an asynchronous visitor model of HavoqGT. Using this dataset, graph representation and HavoqGT framework, we implemented and evaluated time dependent Single Source Shortest Path (TD-SSSP) and Connected Vertices.

Bellerophon: A Computational Workflow Environment for Real-time Analysis, Artifact Management, and Regression Testing of Core-Collapse Supernova Simulations

Authors: Eric Lingerfelt, Bronson Messer (Oak Ridge National Laboratory)

We present an overview of the Bellerophon software system, which has been built to support CHIMERA, a production-level HPC application that simulates the evolution of core-collapse supernovae. Developed over the last 5 years at ORNL, Bellerophon enables CHIMERA's geographically dispersed team of collaborators to perform job monitoring and real-time data analysis from multiple supercomputing resources, including platforms at OLCF, NERSC, and NICS. Its multi-tier architecture provides an encapsulated, end-to-end software solution that enables the CHIMERA team to quickly and easily access highly customizable animated and static views of results from anywhere in the world via a web-deliverable, cross-platform desktop application. Bellerophon has quickly evolved into the CHIMERA team's de facto work environment for analysis, artifact management, regression testing, and other workflow tasks.

Caliper: Composite Performance Data Collection in HPC Codes

Authors: David Boehme, Todd Gamblin, Peer-Timo Bremer, Olga T. Pearce, Martin Schulz (Lawrence Livermore National Laboratory)

Correlating performance metrics with program context information is key to understanding HPC application behavior. Given the composite architecture of modern HPC applications, metrics and context information must be correlated from

independent places across the software stack. Current data-collection approaches either focus on singular performance aspects, limiting the ability to draw correlations, or are not flexible enough to capture custom, application-specific performance factors. With the Caliper framework, we introduce (1) a flexible data model that can efficiently represent arbitrary performance-related data, and (2) a library that transparently combines performance metrics and program context information provided by source-code annotations and automatic measurement modules. Measurement modules and source-code annotations in different program and system components are independent of each other and can be combined in an arbitrary fashion. This composite approach allows us to easily create powerful measurement solutions that facilitate the correlation of performance data across the software stack.

A Real-Time Tsunami Inundation Forecast System for Tsunami Disaster Prevention and Mitigation

Authors: Akihiro Musa, Hiroshi Matsuoka, Osamu Watanabe (NEC Corporation); Yoichi Murashima (Kokusai Kogyo Co., Ltd.); Shunichi Koshimura, Ryota Hino (Tohoku University), Yusaku Ohta, Hiroaki Kobayashi (Tohoku University)

The tsunami disasters of Indonesia, Chile and Japan have occurred in the last decade, and inflicted casualties and damaged social infrastructures. Therefore, tsunami forecasting systems are urgently required worldwide for disaster prevention and mitigation. For this purpose, we have developed a real-time tsunami inundation forecast system that can complete a tsunami inundation simulation at the level of 10-meter grids within 20 minutes. A HPC system is essential to complete such a huge simulation. As the tsunami inundation simulation program is memory-intensive, we incorporate the high memory-bandwidth NEC vector supercomputer SX-ACE into the system. In this poster, we describe an overview of the system and the characteristics of SX-ACE, the performance evaluation of tsunami simulation on SX-ACE, and an emergency job management mechanism for real-time simulation on SX-ACE. The performance evaluation indicates that the performance of SX-ACE with 512 cores is equivalent to that of K computer with 9469 cores.

Multi-GPU Graph Analytics

Authors: Yuechao Pan, Yangzihao Wang, Yuduo Wu, Carl Yang, John D. Owens (University of California, Davis)

We present Gunrock, a multi-GPU graph processing library that enables easy graph algorithm implementation and extension onto multiple GPUs, for scalable performance on large graphs with billions of edges. Our high-level data-centric abstraction focuses on vertex or edge frontier operations. With this abstraction, Gunrock balances between performance and low programming complexity, by coupling high performance GPU computing primitives and optimization strategies. Our

multi-GPU framework only requires programmers to specify a few algorithm-dependent blocks, hiding most multi-GPU related implementation details. The framework effectively overlaps computation and communication, and implements a just-enough memory allocation scheme that allows memory usage to scale with more GPUs. We achieve 22GTEPS peak performance for BFS, which is the best of all single-node GPU graph libraries, and demonstrate a 6x speed-up with 2x total memory consumption on 8 GPUs. We identify synchronization / data communication patterns, graph topologies, and partitioning algorithms as limiting factors to further scalability.

Multi-Level Blocking Optimization for Fast Sparse Matrix Vector Multiplication on GPUs

Authors: Yusuke Nagasaka, Akira Nukada, Satoshi Matsuoka (Tokyo Institute of Technology)

Many scientific and industrial simulations require solving large linear equations, whose bottleneck is sparse matrix vector multiplication (SpMV). Although some previous work has shown improvement of SpMV performance on GPU, the critical bottlenecks such as requirement of high memory bandwidth and low cache hit ratio due to random memory access to input vector still remain. We propose the state of the art sparse matrix format reducing memory access for GPU. Adaptive Multi-level Blocking (AMB) format compresses the column index by using 16-bit integer and several blocking optimizations, and we also devise effective SpMV kernel. We evaluated the performance of our approach for 62 positive definite large size matrices in single precision. AMB format achieves significant speedup of 2.83x on maximum and 1.75x on average compared to cuSparse library and 1.38x on maximum and 1.0x on average compared to yaSpMV, which is a recently proposed fast SpMV library.

GLOVE: An Interactive Visualization Service Framework with Multi-Dimensional Indexing on the GPU

Authors: Jinwoong Kim (Ulsan National Institute of Science and Technology); Sehoon Lee, Joong-Youn Lee (Korea Institute of Science and Technology Information); Beomseok Nam (Ulsan National Institute of Science and Technology), Min Ah Kim (Korea Institute of Science and Technology Information)

In this poster, we present an interactive visualization service framework - GLOVE for massively large scale scientific datasets. GLOVE is designed to support multiple user workloads and to provide interactive user interfaces. GLOVE improves the scalability of the framework and reduces the query latency by managing datasets in large scale distributed memories and by employing multi-dimensional indexing trees that help navigate through massively large datasets. In this poster, we primarily focus on our design and implementation of GPU-based multi-dimensional indexing for GLOVE. Our massively parallel tree traversal algorithm - MPRTS for the GPU avoids irregular

memory access patterns and recursive backtracking and significantly improves the warp efficiency of the GPU and reduces query response time. Our preliminary experiments show GPU-based indexing accelerates the query performance by an order of magnitude.

Molecular Electrostatic Potential Evaluation with the Fragment Molecular Orbital Method

Authors: Yuri Alexeev (Argonne National Laboratory), Dmitri Fedorov (National Institute of Advanced Industrial Science and Technology)

The molecular electrostatic potential (MEP) is a useful tool to analyze intermolecular electrostatic interactions and the properties of the chemical system. The most accurate way to compute MEP is to use quantum mechanics methods, but it is prohibitively computationally expensive for large chemical systems. Presently, the ability to compute MEP accurately for large systems is in high demand because of the recent advances in X-ray, cryo-electron microscopy, NMR, and mass-spectrometry techniques for elucidation of structure and conformation. The solution is to use linearly scaling QM methods, like the fragment molecular orbital (FMO) method. The major problems are accurate computation of MEP, the storage of electron density and electrostatic potential in memory, and scalability of the code. To address these issues, we implemented different MEP algorithms and compared their performance. It was found that the new fragment cube method (FCM) produces accurate MEP at a fraction of the cost.

Benchmarking and Experimental Testbed Studies of AWGR-Based, Multi-Layer Photonic Interconnects for Low-Latency, Energy-Efficient Computing Architectures

Authors: Paolo Grani, Roberto Proietti, Zheng Cao, S. J. Ben Yoo (University of California, Davis)

In this work, we show future research directions about optical computing systems based on Arrayed Waveguide Grating Routers. We analyze the optical integration at different layers of large-scale architectures. At board-level, we aim to demonstrate significant execution time improvements and energy efficiency of optically interconnected systems running real benchmark applications on a cycle-accurate simulator. We simulate tiled optical Multi-Socket Blades (MSBs), and we implement multiple optimization techniques to achieve different tradeoffs between performance and energy consumption and to prove the importance of chip-level optical integration. In the large-scale we show, through a network simulator, significant latency and throughput improvement running uniform distributed traffic for both intra- and inter-cluster optical-based architecture. A software framework, based on trace-driven approach, will be developed for benchmarking large-scale solutions. We show an experimental hardware setup for next future testbeds to demonstrate the efficiency, in terms of

throughput and energy consumption, of the proposed optical architectures.

Large-Scale and Massively Parallel Phase-Field Simulations of Pattern Formations in Ternary Eutectic Alloys

Authors: Johannes Hötzer (Karlsruhe University of Applied Sciences); Martin Bauer (FAU Erlangen Nuremberg); Marcus Jainta, Philipp Steinmetz, Marco Berghoff (Karlsruhe Institute of Technology); Florian Schornbaum, Christian Godenschwager, Harald Köstler, Ulrich Rüde (FAU Erlangen Nuremberg); Britta Nestler (Karlsruhe Institute of Technology)

Various patterns form during directional solidification of ternary eutectic alloys. These macroscopic patterns strongly influence the material properties. To study the influence of the material and process parameters for a wide range of ternary alloys during solidification, simulations are conducted to gain new insights. In this poster, we present the results of massive parallel phase-field simulations on up to 262144 cores, based on the HPC framework waLBerla. Results of optimization techniques are shown, starting from the model up to the code level, including buffering strategies and vectorization. The approach comprises systematic node level performance engineering and gains in a speedup of factor 80 compared to the original code. Excellent weak scaling results on the currently largest German supercomputers, SuperMUC, Hornet and Juqueen, are presented. Novel methods like principle component analysis and graph based approaches are applied to compare the computed microstructural patterns with experimental Al-Ag-Cu micrographs.

Accurate and Efficient QM/MM Molecular Dynamics on 86,016 Cores of SuperMUC Phase 2

Authors: Magnus Schwörer (Ludwig Maximilian University of Munich); Momme Allalen, Ferdinand Jamitzky, Helmut Satzger (Leibniz Supercomputing Center); Gerald Mathias (Ludwig Maximilian University of Munich)

We have recently presented a hybrid approach for molecular dynamics (MD) simulations, in which the atomic forces are calculated quantum-mechanically by grid-based density functional theory (DFT) for a solute molecule, while a polarizable molecular mechanics (PMM) force field is applied to the solvent, thus explicitly modeling electronic polarization effects. In particular, we combine the PMM-MD driver IPHIGENIE and the DFT program CPMD into one MPI/OpenMP-parallel executable. The IPHIGENIE/CPMD program package is now available to the community. In the latest version of the algorithm, the performance of calculating long-range interactions using hierarchically nested fast multipole expansions was enhanced by one order of magnitude. Furthermore, a generalized ensemble method was adopted for DFT/PMM-MD and now enables efficient conformational sampling of biomolecules using state-of-the-art DFT/PMM models. The poster presents the algorithm,

the scaling of IPHIGENIE/CPMD up to 86,016 Intel Haswell cores of SuperMUC Phase 2, and preliminary results of a large-scale biomolecular application.

A Splitting Approach for the Parallel Solution of Large Linear Systems on GPU Cards

Authors: Ang Li, Radu Serban, Dan Negrut (University of Wisconsin-Madison)

We discuss a GPU solver for sparse or dense banded linear systems $Ax=b$, with A possibly nonsymmetric, sparse, and moderately large. The split and parallelize (SaP) approach seeks to partition the matrix A into P diagonal sub-blocks which are independently factored in parallel. The solution may choose to consider or to ignore the off-diagonal coupling blocks. This approach, along with the Krylov iterative methods that it preconditions, are implemented in the SaP::GPU solver, which runs entirely on the GPU except for several stages involved in preliminary row-column permutations.

SaP-GPU compares well in terms of efficiency with three commonly used sparse direct solvers: PARDISO, SuperLU, and MUMPS. Compared to Intel's MKL, SaP-GPU proves to also be performant on dense banded systems that are close to being diagonally dominant.

SaP-GPU is available open source under a permissive BSD3 license.

MLTUNE: A Tool-Chain for Automating the Workflow of Machine-Learning Based Performance Tuning

Authors: Biplab Kumar Saha, Saami Rahman, Apan Qasem (Texas State University)

Recent interests in machine learning-based methods have produced several sophisticated models for performance optimization and workload characterization. Generally, these models are sensitive to architectural parameters and are most effective when trained on the target platform. Training of these models, however, is a fairly involved process. It requires knowledge of statistical techniques that the end-users of such tools may not possess. This poster presents MLTUNE, a tool-chain that automates the workflow for developing machine learning algorithms for performance tuning. Leveraging existing open-source software, the tool-chain provides automated mechanisms for sample generation, dynamic feature extraction, feature selection, data labeling, validation and model selection. MLTUNE can also be used by performance engineers to build their own learning models. The poster highlights the key design features of MLTUNE, which sacrifices some sophistication for generalization and automation. The system's applicability is demonstrated with an auto-generated model for predicting profitable affinity configurations for parallel workloads.

High Performance Data Structures for Multicore Environments

Authors: *Giuliano Laccetti, Marco Lapegna (University of Naples)*

Heap-based priority queues are common dynamical data structures used in several fields, ranging from operating systems to scientific applications. However, the rise of new multicore CPUs introduced new challenges in the process of design of these data structures: in addition to traditional requirements like correctness and progress, the scalability is of paramount importance. It is a common opinion that these two demands are partially in conflict with each other so that it is necessary to relax the requirements about a behavior identical to the corresponding sequential data structures. In this poster we introduce a loosely coordinated approach for the management of heap-based priority queues on multicore CPUs, with the aim to realize a tradeoff between efficiency and sequential correctness. The approach is based on a sharing of information only among a small number of cores, so as to improve performance without completely losing the features of the data structure.

BurstFS: A Distributed Burst Buffer File System for Scientific Applications

Authors: *Teng Wang (Florida State University); Kathryn Mohror, Adam Moody (Lawrence Livermore National Laboratory); Weikuan Yu (Florida State University), Kento Sato (Lawrence Livermore National Laboratory)*

Large-scale scientific applications on leadership computing facilities usually generate colossal amount of scientific datasets. These datasets impose a great challenge to the backend storage systems for timely data service. Recently, the idea of burst buffer was proposed as a promising storage solution to cope with the exploding data pressure. Node-local burst buffers, which locate on the individual compute node, have been already deployed on several HPC systems. Despite their great potential for massive scalability, there is still a lack of software solution to leverage such storage architecture. In this study, we propose BurstFS, a distributed burst buffer file system, to exploit the node-local burst buffers, and provide scientific applications with efficient, scalable I/O service. Our initial evaluation demonstrated that BurstFS increased bandwidths by 1.83x and its performance scaled linearly with the number of involved processes.

Increasing Fabric Utilization with Job-Aware Routing

Authors: *Jens Domke (Dresden University of Technology)*

The InfiniBand (IB) technology became one of the most widely used interconnects for HPC systems in recent years. The achievable communication throughput for parallel applications depends heavily on the available number of links and switches of the fabric. These numbers are derived from the

quality of the used routing algorithm, which usually optimizes the forwarding tables for global path balancing. However, in a multi-user/multi-job HPC environment this results in suboptimal usage of the shared network by individual jobs. We extend an existing routing algorithm to factor in the locality of running parallel applications, and we create an interface between the batch system and the subnet manager of IB to drive necessary re-routing steps for the fabric. As a result, our job-aware routing allows each running parallel application to make better use of the shared IB fabric, and therefore increase the application performance and the overall fabric utilization.

Exploiting Domain Knowledge to Optimize Mesh Partitioning for Multi-Scale Methods

Authors: *Muhammad Hasan Jamal, Milind Kulkarni, Arun Prakash (Purdue University)*

Multi-scale computational methods are widely used for complex scientific computing problems that span multiple spatial and temporal scales. These problem meshes are decomposed into multiple subdomains that are solved independently at different timescales and granularity and are then coupled back to get the desired solution. The computational cost associated with different scales can vary by multiple orders of magnitude. Hence the problem of finding an optimal mesh partitioning, choosing appropriate timescales for the partitions, and determining the number of partitions at each timescale is non-trivial. Existing partitioning tools overlook the constraints posed by multi-scale methods, leading to sub-optimal partitions with a high performance penalty. Our partitioning approach exploits domain knowledge to handle multi-scale problems appropriately and produces optimized mesh partitioning automatically. Our approach produces decompositions that perform as well as, if not better than, decompositions produced by state-of-the-art partitioners, like METIS, and even those that are manually constructed by domain scientists.

Efficient GPU Techniques for Processing Temporally Correlated Satellite Image Data

Authors: *Tahsin A. Reza (University of British Columbia), Dipayan Mukherjee (Indian Institute of Technology Kharagpur), Matei Ripeanu (University of British Columbia)*

Spatio-temporal processing has many usages in different scientific domains, e.g., geostatistical processing, video processing and signal processing. Spatio-temporal processing typically operates on massive volume multi-dimensional data that make cache-efficient processing challenging. In this paper, we present highlights of our ongoing work on efficient parallel processing of spatio-temporal data on massively parallel many-core platforms, GPUs. Our example application solves a unique problem within Interferometric Synthetic Aperture Radar (InSAR) processing pipeline. The goal is selecting objects that appear stable across a set of satellite images taken over

time. We present three GPU approaches that differ in terms of thread mapping, parallel efficiency and memory access patterns. We conduct roofline analysis to understand how the most time consuming GPU kernel can be improved. Through detailed benchmarking using hardware counters, we gain insights into runtime performance of the GPU techniques and discuss their tradeoffs.

PDE Preconditioner Resilient to Soft and Hard Faults

Authors: *Francesco Rizzi, Karla Morris, Kathryn Dahlgren, Khachik Sargsyan (Sandia National Laboratories); Paul Mycek (Duke University), Cosmin Safta (Sandia National Laboratories); Olivier LeMaitre, Omar Knio (Duke University); Bert Debusschere (Sandia National Laboratories)*

We present a resilient domain-decomposition preconditioner for partial differential equations (PDEs). The algorithm reformulates the PDE as a sampling problem, followed by a solution update through data manipulation that is resilient to both soft and hard faults. The algorithm exploits data locality to reduce global communication. We discuss a server-client implementation where all state information is held by the servers, and clients are designed solely as computational units. Focusing on the stages of the algorithm that are most intensive in communication and computation, we explore the scalability of the actual code up to 12k cores, and build an SST/macro skeleton allowing us to extrapolate up to 50k cores. We show the resilience under simulated hard and soft faults for a 2D linear Poisson equation.

Evaluating DVFS and Concurrency Throttling on IBM's Power8 Architecture

Authors: *Wei Wang (University of Delaware), Edgar A. Leon (Lawrence Livermore National Laboratory)*

Two of the world's next-generation of supercomputers at the U.S. national laboratories will be based on IBM's Power architecture. Early insights of how this architecture impacts applications are beneficial to extract the best performance. In this work, we investigate the effects of DVFS and thread concurrency throttling on the performance of applications for an IBM OpenPower, Power8 system. We apply these techniques dynamically on a per region basis to three HPC codes: miniFE, LULESH, and Graph500. Our empirical results offer the following insights. First, concurrency throttling provides significant performance improvements. Second, 4-way simultaneous multi-threading (SMT) performs as well as 8-way SMT. This may result in greater power-efficiency. And, third, applying specific frequency and concurrency throttling combinations on memory-bound regions results in greater performance and energy efficiency.

Analyzing the Performance of a Sparse Matrix Vector Multiply for Extreme Scale Computers

Authors: *Amanda Bienz, Jon Calhoun, Luke Olson (University of Illinois at Urbana-Champaign)*

As HPC systems continue to progress towards extreme scale, scalability of applications becomes critical. Scalability is not purely an algorithmic property but is dependent on machine architecture e.g. network contention. Fundamental to a large class of HPC applications is a sparse-matrix vector multiply (SpMV). We investigate the performance and scalability of SpMV routines in widely used software packages PETSc, Hypre, and Trilinos. By the use of an asynchronous algorithm, we improve both scalability and performance for a SpMV by 9.86x against some packages when there is network contention.

PPP: Parallel Programming with Pictures

Authors: *Annette C. Feng, Wu Feng, Eli Tilevich (Virginia Polytechnic Institute and State University)*

Multicore computers are becoming the norm in our ubiquitous computing systems, and the need for programmers who can write codes for these systems and realize requisite, measurable performance gains continues to rise. However, programmers spend many years learning their craft using sequential languages before ever being introduced to parallel programming. By then, it is difficult for many programmers to "think in parallel." Parallel programming constructs ought to be as fundamental as if-then-else statements and should be taught from the outset rather than being delayed until a senior-year systems course in college. Thus, we introduce explicitly parallel programming constructs to a (sequential) block-based language called Snap!, which was derived from Scratch at MIT, and show that this approach can be a successful way to introduce parallel programming to K-12 students, college students, and even professionals desiring re-training.

Accelerating the B-Spline Evaluation in Quantum Monte Carlo

Authors: *Ye Luo, Anouar Benali, Vitali Morozov (Argonne National Laboratory)*

In Quantum Monte Carlo simulations, the many-body Schrödinger equation can be solved with its wavefunction represented by B-splines which is computationally less intensive $O(N^2)$ than the commonly used planewaves $O(N^3)$. Despite the high efficiency of B-splines, the wavefunction evaluation still takes over 20% of the total application time. We recently improved the algorithm by fully taking advantage of vectorization and optimizing the memory access on BG/Q. We achieved about three-fold speedup in the subroutines calculating multiple B-splines. Threading capability is also added to the new algorithm to maximize the single node performance. According to the specifications of the upcoming HPC systems (long vector

units, more integrated cores, higher memory bandwidth), all the methods used to design the new algorithm make it ready to efficiently exploit the new features of these systems.

Design of a NVRAM Specialized Degree Aware Dynamic Graph Data Structure

Authors: Keita Iwabuchi (Tokyo Institute of Technology); Roger A. Pearce, Brian Van Essen, Maya Gokhale (Lawrence Livermore National Laboratory); Satoshi Matsuoka (Tokyo Institute of Technology)

Large-scale graph processing, where the graph is too large to fit in main-memory, is often required in diverse application fields. Recently, Non-Volatile Random Access Memories (NVRAM) devices, e.g., NAND Flash, have enabled the possibility to extend main-memory capacity without extremely high cost and power consumption. Our previous work demonstrated that NVRAM can be a powerful storage media for graph applications.

In many graph applications, the structure of the graph changes dynamically over time and may require real time analysis, such as genome analysis and network security. However, constructing large graphs is expensive especially for NVRAM devices, due to sparse random memory access patterns.

We describe a NVRAM specialized degree aware dynamic graph data structure using open addressing compact hash tables in order to minimize the number of page misses. We demonstrate that our dynamic graph structure can scale near-linearly on an out-of-core dynamic edge insertion workload.

ACM Student Research Competition Posters

An Analysis of Network Congestion in the Titan Supercomputer's Interconnect

Author: Jonathan Freed (University of South Carolina)

The Titan supercomputer is used by computational scientists to run large-scale simulations. These simulations often run concurrently, thus sharing system resources. A saturated system can result in network congestion—negatively affecting interconnect bandwidth. Our project analyzed data collected by testing the bandwidth between different node pairs. In particular, we searched for correlations when the bandwidth utilization was low. We investigated the direct path between the two test nodes, as well as the neighborhood of nodes that are one connection away from the direct path. For each set of nodes, we analyzed the effects of the number of “busy” nodes (nodes currently allocated to jobs), the jobs that were running, the date and time the test took place, and the distance between the nodes. By understanding job interference, we can develop job-scheduling strategies that lower such interference and lead to more efficient use of Titan's resources and faster computations for researchers.

Lessons from Post-Processing Climate Data on Modern Flash-Based HPC Systems

Author: Adnan Haider (Illinois Institute of Technology)

Post-processing climate data applications are necessary tools for climate scientists. Because these applications are I/O bound, flash devices can provide a significant and needed boost in performance. However, the tradeoffs associated with different flash architectures is unclear, making the choice of an appropriate flash architecture deployment for post-processing software difficult. Thus, we analyzed the performance of a local and pooled flash architecture to quantify the performance tradeoffs. We learned three main concepts. First, we found that an incorrect matching between storage architecture and I/O workload can hide the benefits of flash by increasing runtime by 2x. Second, after tuning Gordon's architecture, we found that a local flash architecture could be a cost-effective alternative to a pooled architecture if scalability and interconnect bottlenecks are alleviated. Third, the benefits of running post-processing applications on the latest data-intensive systems which lack flash devices can provide significant speedups, lessening the need for flash.

UtiliStation: Increasing the Utilization of the International Space Station with Big Data Analytics for Stowage

Author: Ellis Giles (Rice University)

The International Space Station is one of the most expensive and complex objects ever constructed. It benefits humanity with a unique microgravity and extreme research environment.

Sending experiments and materials to the ISS is costly and logistically complex. The high dollar and time cost, along with increasing demand, has led to large amounts of items being stowed onboard the ISS. However, stowage space onboard the ISS is very limited, and, as more items are stowed, the more time it takes to retrieve an item - up to 25% of a crew member's time.

This research involves creating a software system called UtiliStation, which may increase the utilization of the ISS by optimizing onboard stowage of resources over time. It can reduce the complexity of cargo management while reducing the ground and astronaut time by using Map-Reduce functions on inventory and crew procedure data.

Rapid Replication of Multi-Petabyte File Systems

Author: Justin G. Sybrandt (National Energy Research Scientific Computing Center)

As file systems grow larger, tools which were once industry standard become unsustainable at scale. Today, large data sets containing hundreds of millions of files often take longer to traverse than to copy. The time needed to replicate a file

system has grown from hours to weeks, an unrealistic wait for a backup. Distsync is our new utility that can quickly update an out-of-date file system replica. By utilizing General Parallel File System (GPFS) policy scans, distsync finds changed files without navigating between directories. It can then parallelize work across multiple nodes, maximizing the performance of a GPFS. NERSC is currently using distsync to replicate file systems of over 100 million inodes and over 4 petabytes.

Integrating STELLA & MODESTO: Definition and Optimization of Complex Stencil Programs

Author: Tobias Gysi (ETH Zurich)

An efficient implementation of complex stencil programs requires data-locality transformations such as loop tiling and loop fusion. When applying these transformations we face two main challenges: 1) the direct application is costly and typically results in code that is optimized for a specific target architecture and 2) the selection of good transformations is not straightforward. We address these challenges using STELLA a stencil library that abstracts data-locality transformations and MODESTO a model-driven stencil optimization framework. In particular, MODESTO represents different stencil program implementation variants using a stencil algebra and selects good transformations based on a compile-time performance model. We evaluate the effectiveness of the approach using example kernels from the COSMO atmospheric model. Compared to naive and expert-tuned variants we attain a 2.0-3.1x and a 1.0-1.8x speedup respectively.

High Performance Model Based Image Reconstruction

Author: Xiao Wang (Purdue University)

Computed Tomography (CT) is an important technique used in a wide range of applications, ranging from explosive detection, medical imaging, and scientific imaging to non-destructive testing. Among available reconstruction methods, Model Based Iterative Reconstruction (MBIR) produces higher quality images than commonly used Filtered Backprojection (FBP) but at a very high computational cost. We describe a new MBIR implementation, PSV-ICD that significantly reduces the computational costs of MBIR while retaining its benefits. It uses a novel organization of the scanner data into supervoxels (SV) that, combined with a supervoxel buffer (SVB), dramatically increases locality and prefetching. Further performance improvements are obtained by a novel parallelization across SVs. Experimental data is presented showing a speedup of 187x resulting from these techniques from the parallelization on 20 cores.

Forecasting Storms in Parallel File Systems

Author: Ryan McKenna (University of Delaware)

Large-scale scientific applications rely on the parallel file system (PFS) to store checkpoints and outputs. When the PFS is over-utilized, applications can slow down significantly as they compete for scarce bandwidth. To prevent this type of “filesystem storm”, schedulers must avoid running many IO-intensive jobs at the same time. To effectively implement such a strategy, schedulers must predict the IO workload and runtime of future jobs. In this poster, we explore the use of machine learning methods to forecast file system usage and to predict the runtimes of queued jobs using historical data. We show that our runtime predictions achieve over 80% accuracy to within 10 minutes of actual runtime.

PEAK: Parallel EM Algorithm using Kd-tree

Author: Laleh Aghababae Beni (University of California, Irvine)

The data mining community voted the Expectation Maximization (EM) algorithm as one of the top ten algorithms having the most impact on data mining research. EM is a popular iterative algorithm for learning mixture models with applications in various areas from computer vision and astronomy, to signal processing. We introduce a new high-performance parallel algorithm on modern multicore systems that impacts all stages of EM. We use tree data-structures and user-controlled approximations to reduce the asymptotic runtime complexity of EM with significant performance improvements. PEAK utilizes the same tree and algorithmic framework for all the stages of EM.

Experimental results show that our parallel algorithm significantly outperforms the state-of-the-art algorithms and libraries on all dataset configurations (varying number of points, dimensionality of the dataset, and number of mixtures). Looking forward, we identify approaches to extend this idea to a larger scale of similar problems.

A High-Performance Preconditioned SVD Solver for Accurately Computing Large-Scale Singular Value Problems in PRIMME

Author: Lingfei Wu (William & Mary)

The dramatic increase in the demand for solving large scale singular value problems has rekindled interest in iterative methods for the SVD. Unlike the remarkable progress in dense SVD solvers, some promising recent advances in large scale iterative methods are still plagued by slow convergence and accuracy limitations for computing smallest singular triplets. Furthermore, their current implementations in MATLAB cannot address the required large problems. Recently, we presented a preconditioned, two-stage method to effectively and accurately compute a small number of extreme singular

lar triplets. In this research, we present high-performance software, PRIMME_SVDS, that implements our hybrid method based on the state-of-the-art eigensolver package PRIMME for both largest and smallest singular values. PRIMME_SVDS fills a gap in production level software for computing the partial SVD, especially with preconditioning. The numerical experiments demonstrate its superior performance compared to other state-of-the-art methods and its good scalability performance under strong and weak scaling.

Improving Application Concurrency on GPUs by Managing Implicit and Explicit Synchronizations

Author: Michael C. Butler (University of Missouri)

GPUs have progressively become part of shared computing environments, such as HPC servers and clusters. Commonly used GPU software stacks (e.g., CUDA and OpenCL), however, are designed for the dedicated use of GPUs by a single application, possibly leading to resource underutilization. In recent years, several node-level runtime components have been proposed to allow the efficient sharing of GPUs among concurrent applications; however, they are limited by synchronizations embedded in the applications or implicitly introduced by the GPU software stack. In this work, we analyze the effect of explicit and implicit synchronizations on application concurrency and GPU utilization, design runtime mechanisms to bypass these synchronizations, and integrate these mechanisms into a GPU virtualization runtime named Sync-Free GPU (SF-GPU). The resultant runtime removes unnecessary blockages caused by multitenancy, ensuring any two applications running on the same device experience limited to no interference. Finally, we evaluate the impact of our proposed mechanisms.

Performance Analysis and Optimization of the Weather Research and Forecasting Model (WRF) on Intel Multicore and Manycore Architectures

Author: Samuel J. Elliott (University of Colorado Boulder)

The Weather Research and Forecasting (WRF) Model is a mesoscale numerical weather prediction system designed for both atmospheric research and operational forecasting needs. The benchmarks used in this study were run on the Texas Advanced Computing Center (TACC) Stampede cluster, which utilizes Intel Xeon E5-2660 CPU's and Xeon Phi SE10P coprocessors. Many aspects of WRF optimization were analyzed, many of which contributed significantly to optimization on Xeon Phi. These optimizations show that symmetric execution on Xeon and Xeon Phi can be used for highly efficient WRF simulations that significantly speed up execution relative to running on either homogeneous architecture.

Performance Analysis and Optimization of a Hybrid Distributed Reverse Time Migration Application

Author: Sri Raj Paul (Rice University)

Applications to analyze seismic data employ scalable parallel systems to produce timely results. To fully exploit emerging processor architectures, programs will need to employ threaded parallelism within a node and message passing across nodes. Today, MPI+OpenMP is the preferred programming model for this task. However, tuning hybrid programs for clusters is difficult. Performance tools can help users identify bottlenecks and uncover opportunities for improvement. This poster describes our experiences of applying Rice University's HPCToolkit and hardware performance counters to gain insight into an MPI+OpenMP code that performs Reverse Time Migration (RTM) on a cluster of multicore processors. The tools provided us with insights into the effectiveness of the domain decomposition strategy, the use of threaded parallelism, and functional unit utilization in individual cores. By applying the insights obtained from the tools, we were able to improve the performance of the RTM code by roughly 30 percent.

AccFFT: A New Parallel FFT Library for CPU and GPU Architectures

Author: Amir Gholami (The University of Texas at Austin)

We present a new distributed-FFT library. Despite the extensive work on FFTs, we achieve significant speedups. Our library uses novel all-to-all communication algorithms to overcome this barrier. These schemes are modified for GPUs to effectively hide PCI-e overhead. Even though we do not use GPU-Direct technology, the GPU results are either better or almost the same as the CPU times (corresponding to 16 or 20 CPU cores). We present performance results on the Maverick and Stampede platforms at the Texas Advanced Computing Center (TACC) and on the Titan system at the Oak Ridge National Laboratory (ORNL). Comparison with P3DFFT and PFFT libraries show a consistent 2-3x speedup across a range of processor counts and problem sizes. Comparison with FFTE library (GPU only) shows a similar trend with 2x speedup. The library is tested up to 131K cores and 4,096 GPUs of Titan, and up to 16K cores of Stampede.

Power Balancing Complements Load Balancing

Author: Neha Gholkar (North Carolina State University)

Until recently the research community has focused on minimizing energy usage of super computers. Considering the US DOE's mandate of power constraint of 20 MW for the exascale sites, effort needs to be directed toward minimizing the wasteful usage of power while maximizing performance under this constraint.

Most of the workloads on supercomputers are often coupled parallel scientific simulations. Real applications like ParaDis tend to have imbalanced load distribution. We have also observed that processors tend to be inhomogeneous under a power constraint. Conventional load balancing algorithms are oblivious of the power budget and the inhomogeneity of the processors and hence they do not lead to optimal solutions. For the same reasons, naive uniform power capping is not a solution either.

We propose an algorithm that optimizes a job for performance under a power constraint. It has a combination of load balancing and power balancing.

Portable Performance of Large-Scale Physics Applications: Toward Targeting Heterogeneous Exascale Architectures Through Application Fitting

Author: William Killian (University of Delaware)

Physics simulations are one of the driving applications for supercomputing and this trend is expected to continue as we transition to exascale computing. Modern and upcoming hardware design exposes tens to thousands of threads to applications, and achieving peak performance mandates harnessing all available parallelism in a single node. In this work we focus on two physics micro-benchmarks representative of kernels found in multi-physics codes. We map these onto three target architectures: Intel CPUs, IBM BlueGene/Q, and NVIDIA GPUs. Speedups on CPUs were up to 12x over our baseline while speedups on BlueGene/Q and GPUs peaked at 40x and 18x, respectively. We achieved 54% of peak performance on a single core. Using compiler directives with additional architecture-aware source code utilities allowed for code portability. Based on our experience, we list a set of guidelines for programmers and scientists to follow towards attaining a single, performance portable implementation.

IGen: The Illinois Genomics Execution Environment

Author: Subho Sankar Banerjee (University of Illinois at Urbana-Champaign)

There has been a great optimism for the usage of DNA sequence data in clinical practice, notably for diagnostics and developing personalized treatments tailored to an individual's genome. This poster, presents a study of software tools used in identifying and characterizing mutations in a genome. We present IGen, a runtime framework which executes this workflow as a data-flow graph over a partitioned global address space. Preliminary results on the Blue Waters supercomputer show that IGen is able to accelerate single node performance (alignment - 1.2x, variant calling - 9x), as well as distribute the computation across several machines with near-linear scaling. Theoretical models for performance of the entire workflow suggest that IGen will have a ~3x improvement in runtime on a single node with near linear scaling across multiple nodes.

Hilbert Curve Based Flexible Dynamic Partitioning Scheme for Adaptive Scientific Computations

Author: Milinda Fernando (University of Utah)

Space Filling Curves (SFC) are commonly used by the HPC community for partitioning data and for resource allocations. Among the various SFCs, Hilbert curves have been found to be in general superior to the more ubiquitous Morton or Z-Curve. However, their adoption in large-scale HPC codes, especially for partitioning applications, has not been as common as Morton curves due to the computational complexity associated with the Hilbert ordering. In this work, we present an alternative algorithm for computing the Hilbert ordering that can be implemented almost as efficiently as the Morton ordering. Additionally, being based on the concept of the nearest common ancestor—a fundamental property of all SFCs—the algorithm can be applied to all SFCs. We also present comparisons of Morton and Hilbert curve based partitions for adaptive meshes using two partitioning schemes demonstrating the superiority of Hilbert ordering.

Practical Floating-Point Divergence Detection

Author: Wei-Fan Chiang (University of Utah)

Reducing floating-point precision allocation in HPC programs is of considerable interest from the point of view of obtaining higher performance. However, this can lead to unexpected behavioral deviations from the programmer's intent. In this poster, we focus on the problem of divergence detection: when a given floating-point program exhibits different control flow (or differs in terms of other discrete outputs) with respect to the same program interpreted under reals. This problem has remained open even for everyday programs such as those that compute convex-hulls. We propose a classification of the divergent behaviors exhibited by programs, and propose efficient heuristics to generate inputs causing divergence. Our experimental results demonstrate that our input generation heuristics are far more efficient than random input generation for divergence detection, and can exhibit divergence even for programs with thousands of inputs.

Non-Blocking Preconditioned Conjugate Gradient Methods for Extreme-Scale Computing

Author: Paul Eller (University of Illinois at Urbana-Champaign)

To achieve the best performance on extreme-scale systems we need to develop more scalable methods. For the preconditioned conjugate gradient method (PCG), dot products limit scalability because they are a synchronization point. Non-blocking methods have the potential to hide most of the cost of the allreduce and avoid the synchronization cost due to performance variation across cores.

We study three scalable methods that rearrange PCG to reduce communication latency by using a single allreduce (L56PCG, PIPECG) and/or overlap communication and computation using a non-blocking allreduce (NBPCG, PIPECG). Tests on up to 32k cores of Blue Waters show that current non-blocking solver implementations cannot efficiently enough overlap communication and computation to overcome the increased vector operations cost. However performance models show potential for non-blocking solvers to be more scalable than PCG. Performance models show that the ability to minimize the impact of noise throughout PCG may be a key benefit.

I/O Performance Analysis Framework on Measurement Data from Scientific Clusters

Author: Michelle V. Koo (University of California, Berkeley)

Due to increasing scales in data volumes, number of machines, and complex workflow in science applications, it has become more challenging to diagnose job performance and system performance on scientific clusters. This project is motivated by the observations that I/O performance analyses can be conducted from monitored performance measurement data from scientific clusters. Studies of I/O performance behavior have been conducted on the Palomar Transient Factory (PTF) application, an automated wide-field survey of the sky to detect variable and transient objects, by analyzing measurement data collected on NERSC Edison. Visualization tools were developed to aid in identifying I/O performance bottlenecks in the PTF data analysis pipeline. This work led to building an interactive I/O performance analysis framework for measurement data from scientific clusters to further identify performance characteristics and bottlenecks in scientific applications.

Modeling the Impact of Thread Configuration on Power and Performance of GPUs

Author: Tiffany A. Connors (Texas State University)

Because graphics processing units (GPUs) are a low-cost option for achieving high computational power, they have become widely used in high-performance computing. However, GPUs can consume large amounts of power. Due to the associated energy costs, improving energy-efficiency has become a growing concern.

By evaluating the impact of thread configuration on performance and power trade-off, energy-efficient solutions can be identified. Using machine learning, the effect of applying a thread configuration to a program can be accurately predicted in terms of the relative change in performance/power trade-off of a GPU kernel. This enables us to establish which dynamic program features are used to predict the impact of a thread configuration on a program's performance.

Efficient Multiscale Platelets Modeling Using Supercomputers

Author: Na Zhang (Stony Brook University)

This work focuses on developing multiscale models and efficient numerical algorithms for simulating platelets on supercomputers. More specifically, the development of multiple time-stepping algorithms can be applied to optimally use computing resources to model platelet structures at multiple scales, enabling the study of flow-induced platelet-mediated thrombogenicity. In order to achieve this, sophisticated parallel computing algorithms are developed and detailed performance analysis has been conducted. The performance results manifest the possibility of simulating the millisecond-scale hematology at resolutions of nanoscale platelets and meso-scale bio-flows using millions of particles. The computational methodology using multiscale models and algorithms on supercomputers will enable efficient predictive simulations for initial thrombogenicity study and may provide a useful guide for exploring mechanisms of other complex biomedical problems at disparate spatiotemporal scales. This poster will cover the multiscale model, a two-fold speedup strategy, i.e., combined algorithmic multiple time-stepping and hardware GPGPU acceleration, and performance analysis.

Optimization Strategies for Materials Science Applications on Cori: An Intel Knights Landing, Many Integrated Core Architecture

Author: Luther D. Martin (National Energy Research Scientific Computing Center)

NERSC is preparing for the arrival of its Cray XC40 machine dubbed Cori. Cori is built on Intel's Knights-Landing Architecture. Each compute node will have 72 physical cores and 4 hardware threads per core. This is 6x the number of physical cores and 10x the number of virtual cores than the Cray XC30 machine. Cori also comes with a larger hardware vector unit, 512 bits, and high-bandwidth, on-package memory. While most of the current applications currently running on NERSC's XC30 machine will be able to execute on Cori with little to no code refactoring, they will not be optimized and may suffer performance loss. This poster recounts the effectiveness of three optimization strategies on the materials science application VASP: (1) Increasing on node parallelism by adding OpenMP where applicable; (2) refactoring code to allow compilers to vectorize loops; and (3) identifying candidate arrays for the high-bandwidth, on-package memory.

Exploring the Trade-Off Space of Hierarchical Scheduling for Very Large HPC Centers**Authors:** *Stephen Herbein (University of Delaware)*

Hierarchical batch scheduling has many advantages, but a lack of trade-off studies against traditional scheduling precludes the emergence of effective solutions for very large HPC centers. Under hierarchical scheduling, any job can instantiate its own scheduler to schedule the sub-jobs with a distinct policy. While such scheduling parallelism and specialization are attractive for the expanding scale and resource diversity of the centers, hierarchical resource partitioning can lead to poor resource utilization. In this poster, we explore the trade-off space between scheduling complexity and resource utilization under hierarchical scheduling. Our approach includes novel techniques to create a hierarchical workload automatically from a traditional HPC workload. Our preliminary results show that one additional scheduling level can reduce scheduling complexity by a factor of 3.58 while decreasing resource utilization by up to 24%. Our study, therefore, suggests that poor utilization can happen under hierarchical scheduling and motivates dynamic scheduling as a complementary technique.

High Order Automatic Differentiation with MPI**Author:** *Mu Wang (Purdue University)*

Automatic Differentiation (AD) is a technology for analytically (thus accurately within machine precision) computing the derivatives of a function encoded as a computer program. AD has two major modes: forward and reverse. We designed the first high order reverse mode algorithm which can efficiently evaluate derivatives up to any order for a function specified by an MPI program. We begin with identifying an invariant in the first order reverse mode of AD by taking a data-flow perspective. Then extending the invariant to higher orders, we get a serial version of the high order reverse mode algorithm. In order to get the parallel version of the algorithm, we compare the semantics in AD and MPI. We find the connections between declaration of dependent/independent variables in AD and Sending/Receiving variables in MPI. The parallel reverse mode algorithm is derived by exploiting that connection.

Scientific Visualization & Data Analytics Showcase

Wednesday, November 18

Scientific Visualization & Data Analytics Showcase

Chair: Jean M. Favre (Swiss National Supercomputing Center)

10:30am-12pm

Room: Ballroom E

Gasoline Compression Ignition: Optimizing Start of Injection Time

Authors: Joseph A. Insley, Janardhan Kodavasal (Argonne National Laboratory); Xiaochuan Chai (Convergent Science Incorporated); Kevin Harms, Marta Garcia, Sibendu Som (Argonne National Laboratory)

We present visualization of a high-fidelity internal combustion engine computational fluid dynamics (CFD) simulation. This engine operates in an advanced combustion mode called Gasoline Compression Ignition (GCI), where gasoline is used as a fuel in a diesel engine without the use of a spark plug, to combine the high efficiency of a diesel engine with low soot emissions of gasoline fuel. Further, combustion is a result of sequential autoignition without propagating flames, resulting in low temperature combustion, which in turn significantly reduces nitrogen oxides emissions. Fuel injection timing is a critical parameter determining the ignitability of the gasoline-air mixture and engine operation stability. Four different start of injection (SOI) timings were evaluated through CFD simulation.

The simulations confirmed experimental findings of a “sweet spot” in SOI timing, where the most stable combustion was achieved. The engine experiments were unable to explain the reason for the non-monotonic behavior of stability with relation to SOI timing since the combustion chamber of the engine is not optically accessible. However, the visualization of these simulations was critical in explaining this behavior. The visualizations showed that earlier SOI timings resulted in fuel being

directed into the “squish” region of the engine combustion chamber, resulting in greater heat losses and lower reactivity and stability. Later SOI timings, however, did not provide enough time for the gasoline to chemically react and ignite in a stable fashion. This insight is critical in terms of determining optimum fuel injection strategies to enable stable operation of gasoline fuel in GCI mode.

Visualization of a Tornado-Producing Thunderstorm: A Study of Visual Representation

Authors: David Bock (University of Illinois at Urbana-Champaign), Leigh Orf (University of Wisconsin), Robert R. Sisneros (University of Illinois at Urbana-Champaign)

We present a study of visual representations for large-scale scientific data. While non-photorealistic rendering techniques encapsulate significant potential for both research and deployment, photorealistic visualizations remain king with regard to accessibility and thereby acceptability. Therefore in this work we endeavor to understand this phenomenon via the creation of the analogous visualizations in a setting that is readily comparable. Specifically, we employ identical non-color parameters for two visualizations of scientific data, CM1 simulation output data. In the first, data elements are colored according to a natural diagrammatic scale, and in the other physically-based techniques are deployed. The visualizations are then presented to a domain scientist to illicit feedback toward classifying the effectiveness of the different design types.

Visualization Of Airflow Through The Human Respiratory System: The Sniff

Authors: Fernando M. Cucchiatti, Guillermo Marin, Hadrien Calmet, Guillaume Houzeaux, Mariano Vazquez (Barcelona Supercomputing Center)

We created a visualization of a CFD simulation and transport of Lagrangian particles in the airflow through extensive and realistic upper human airways during a rapid and short inhalation (a sniff). The simulation was performed using HPC resources

with the Alya system. The output data was stored in a distributed database connected with an in-house plugin to ParaView for post-processing analysis, and then converted into a proprietary Maya 2014 particle cache. The rendering was performed using HPC resources with Maya, using a multihued color scale to represent the velocity and a ghost trail to indicate direction.

Visualization of Ocean Currents and Eddies in a High-Resolution Ocean Model

Authors: *Francesca Samsel (The University of Texas at Austin), Mark Petersen (Los Alamos National Laboratory); Terece Turton, Gregory Abram (The University of Texas at Austin); James Ahrens, David Rogers (Los Alamos National Laboratory)*

Climate change research relies on models to better understand and predict the complex, interdependent processes that affect the atmosphere, ocean, and land. These models are computationally intensive and produce terabytes to petabytes of data. Visualization and analysis is increasingly difficult, yet is critical to gain scientific insights from large simulations. The recently-developed Model for Prediction Across Scales-Ocean (MPAS-Ocean) is designed to investigate climate change at global high-resolution (5 to 10 km grid cells) on high performance computing platforms. In the accompanying video, we use state-of-the-art visualization techniques to explore the physical processes in the ocean relevant to climate change. These include heat transport, turbulence and eddies, weakening of the meridional overturning circulation, and interaction between a warming ocean and Antarctic ice shelves. The project exemplifies the benefits of tight collaboration among scientists, artists, computer scientists, and visualization specialists. Video link - <https://vimeo.com/136971670>

Extreme Multi-Resolution Visualization: A Challenge On Many Levels

Authors: *Joanna Balme, Eric Brown-Dymkoski, Victor Guerrero, Stephen Jones, Andre Kessler, Adam Lichtl, Kevin Lung, William Moses, Ken Museth (Space Exploration Technologies Corp.); Tom Fogel (NVIDIA Corporation)*

Accurate simulation of turbulent flows presents significant challenges for both computation and visualization, with length and time scales often spanning 6 orders of magnitude in applied cases. Multi-resolution wavelet analysis is an emerging method for dynamically adaptive simulation, offering compression rates greater than 95% with less than 1% error. However the extreme levels of detail still present challenges for visualization.

These challenges require rendering of multi-resolution data directly in order to avoid an explosion in computation and memory cost. Unfortunately, a wavelet grid is ill-suited to direct visualization. By identifying the opportunity to exploit topological similarities between wavelet grids and octrees, using a multi-resolution data structure known as a VDB-tree, it is possible to adapt the simulation-optimal grid such that it may directly be visualized.

To demonstrate this technique, we set up a series of shear-flow simulations which develop turbulent Kelvin-Helmholtz instabilities. Accurate simulation of these instabilities requires extremely fine resolution as the smallest scales couple strongly with the large. We show how direct multi-resolution rendering enables visualization pipelines for simulations with between 14 and 17 levels of detail at near-interactive frame rates. This represents a scale-factor of over 100,000 between the smallest and largest structures with individual frames being generated in less than 1 second.

The use of VDB structures is common in professional and production-quality rendering engines such as Houdini, RenderMan and OptiX, and we show how one of these (OptiX) is enabled by this technique to perform real-time ray tracing on multi-resolution data.

Chemical Visualization of Human Pathogens: The Retroviral Capsids

Authors: *Juan R. Perilla, Boon Chong Goh, John Stone, Klaus Schulten (University of Illinois at Urbana-Champaign)*

Retroviruses are pathogens characterized by their ability to incorporate viral DNA into a host cell's genome. Retroviruses like Rous Sarcoma Virus (RSV) infect cells during mitosis, when the chromatin is exposed to the cytoplasm. Conversely, the genus of lentiviruses, like the human immunodeficiency virus (HIV), has evolved to infect non-dividing cells. Despite infecting cells at different stages of their life cycles, RSV and HIV share a similar late stage replication cycle that is highly dependent on the group actin polyprotein precursor (Gag), which contains the matrix (MA), capsid (CA) and nucleocapsid (NC) proteins. Both HIV's CA and Gag are considered unexploited targets for pharmaceutical intervention. We describe the techniques that were used to build, simulate, analyze and visualize the structures of both Gag and CA. We discuss scientific visualization needs that spurred development of an interactive GPU-accelerated ray tracing engine and the use of remote visualization technologies.

Student Programs

Monday, November 16

Experiencing HPC for Undergraduates

Chair: Alan Sussman (University of Maryland)

3pm-5pm

Room: Hilton 404

Orientation

This session will provide an introduction to HPC for the participants in the HPC for Undergraduates Program. Topics will include an introduction to MPI, shared memory programming, domain decomposition and the typical structure of scientific programming.

Tuesday, November 17

Experiencing HPC for Undergraduates

Chair: Alan Sussman (University of Maryland)

10:30am-12pm

Room: Hilton 404

Introduction to HPC Research

A panel of leading practitioners in HPC will introduce the various aspects of HPC, including architecture, applications, programming models and tools.

Wednesday, November 18

Experiencing HPC for Undergraduates

Chair: Alan Sussman (University of Maryland)

10:30am-12pm

Room: Hilton 404

Graduate Student Perspective

This session will be held as a panel discussion. Current graduate students, some of whom are candidates for the Best Student Paper Award in the Technical Papers program at SC15, will discuss their experiences in being a graduate student in an HPC discipline. They will also talk about the process of writing their award-nominated paper.

Thursday, November 19

Experiencing HPC for Undergraduates

Chair: Alan Sussman (University of Maryland)

10:30am-12pm

Room: Hilton 404

Careers in HPC

This panel will feature a number of distinguished members of the HPC research community discussing their varied career paths. The panel includes representatives from industry, government labs and universities. The session will include ample time for questions from the audience.

Students@SC

Monday, November 16

Mentor-Protégé Mixer**Chair:** *Christine Sweeney (Los Alamos National Laboratory)***3:30pm-5pm****Room: 9ABC**

The Mentor-Protégé Mixer is part of the Mentor-Protégé Program offered by the Students@SC Program. Student protégés are assigned mentors who are experienced conference attendees. The Mixer provides an opportunity for assigned mentor-protégé pairs to meet and discuss topics such as orienting to the SC15 conference and HPC research ideas. There will also be an inspirational speaker. The event will conclude with a general mixing session where all can mingle equally to give students more networking opportunities. Refreshments will be served. Students participating in this program are from the Student Volunteers Program, HPC for Undergraduates, the Student Cluster Competition, Doctoral Showcase and the ACM Student Research Competition. More information on the program can be found at <http://sc15.supercomputing.org/conference-program/studentssc/mentor-protégé-program>.

Student Cluster Competition Kickoff**5pm-11:55pm****Room: Palazzo**

Come to the Palazzo on the first floor to help kick off the Student Cluster Competition and cheer on 9 student teams competing in this real-time, non-stop, 48-hour challenge to assemble a computational cluster at SC15 and demonstrate the greatest sustained performance across a series of applications. Teams selected for this year's competition come from universities all over the United States, Germany, Taiwan, China, Colombia, and Australia. The teams are using big iron hardware with off-the-shelf (or off-the-wall!) solutions for staying within a 26 amp power limit to survive zombie invasions, discover exotic particles, forecast storms, and sequence DNA!

Wednesday, November 18

Student-Postdoc Job & Opportunity Fair**10am-3pm****Room: 9ABC**

This face-to-face event will be open to all students and postdocs participating in the SC15 conference, giving them an opportunity to meet with potential employers. There will be employers from research labs, academic institutions, recruiting agencies and private industry who will meet with students and postdocs to discuss graduate fellowships, internships, summer jobs, co-op programs, graduate school assistant positions and/or permanent employment.

Sunday, November 15

Architecting, Implementing, and Supporting Multi-Level Security Eco-System in HPC, ISR, Big Data Analysis and Other Environments

8:30am-12pm
Room: 18A

Presenters: Joe Swartz, Joshua Koontz, Sarah Storms (Lockheed Martin Corporation); Nathan Rutman (Seagate Technology LLC), Shawn Wells (Red Hat, Inc.), Chuck White (Semper Fortis Solutions, LLC), Carl Smith (Altair Engineering, Inc.), Enoch Long (Splunk Inc.)

Historically cyber security in HPC has been limited to detecting intrusions rather than designing security from the beginning in a holistic, layered approach to protect the system. SELinux has provided the needed framework to address cyber security issues for a decade, but the lack of an HPC and data analysis eco-system based on SELinux and the perception that the resulting configuration is “hard” to use has prevented SELinux configurations from being widely accepted. This tutorial discusses the eco-system that has been developed and certified, debunk the “hard” perception, and illustrate approaches for both government and commercial applications. The tutorial includes discussions on: • SELinux architecture and features • Scale-out Lustre Storage • Applications Performance on SELinux (Vectorization and Parallelization) • Big Data Analysis (Accumulo and Hadoop) • Relational Databases • Batch Queuing • Security Functions (Auditing and other Security Administration actions). The tutorial is based on currently existing, certified and operational SELinux HPC eco-systems and the Department of Energy (DoE) Los Alamos National Labs (LANL) and DoD High Performance Computing Modernization Office (HPCMO) are working through evaluations with the intention of implementing in their systems.

Benchmarking Platforms for Large-Scale Graph Processing and RDF Data Management: The LDBC Approach

8:30am-12pm
Room: 19B

Presenters: Alexandru Iosup (Delft University of Technology), Ana Lucia Varbanescu (University of Amsterdam), Josep Larriba-Pey (Polytechnic University of Catalonia), Peter Boncz (CWI Amsterdam)

Graphs model social networks, human knowledge, and other vital information for business, governance, and academic practice.

Although both industry and academia are developing and tuning many graph-processing algorithms and platforms, the performance of graph-processing platforms has never been explored or compared in-depth. Moreover, graph processing exposes new bottlenecks in traditional HPC systems (see differences in Top500 and Graph500 rankings).

Complementing Graph500, the Linked Data Benchmark Council (LDBC) consortium, which involves over 25 academic and industry partners, focuses on the development of benchmarks that will spur research and industry progress in large-scale graph and RDF data management.

This tutorial introduces the SC audience to the latest LDBC benchmarks. Attendees will learn about and experiment with methods and tools for performance evaluation and optimization for graph processing platforms, and become able to explain the performance dependency Platform-Algorithm-Dataset. The attendees will understand in-depth three workloads, interactive social networking, semi-online business intelligence based on analytical queries, and batch full-graph analytics. The tutorial presents real-world experiences with commonly used systems, from graph databases such as Virtuoso and Neo4j, to parallel GPU-based systems such as Totem and Medusa, to

distributed graph-processing platforms such as Giraph and GraphX. The tutorial also includes significant hands-on and Q&A components.

Data-Intensive Applications on HPC Using Hadoop, Spark and RADICAL-Cybertools

8:30am-12pm

Room: 18C

Presenters: *Shantenu Jha (Rutgers University), Andre Luckow (Clemson University)*

HPC environments have traditionally been designed to meet the compute demands of scientific applications; data has only been a second order concern. With science moving toward data-driven discoveries relying on correlations and patterns in data to form scientific hypotheses, the limitations of HPC approaches become apparent: Low-level abstractions and architectural paradigms, such as the separation of storage and compute, are not optimal for data-intensive applications. While there are powerful computational kernels and libraries available for traditional HPC, there is an apparent lack of functional completeness of analytical libraries. In contrast, the Apache Hadoop ecosystem has grown to be rich with analytical libraries, e.g. Spark MLlib. Bringing the richness of the Hadoop ecosystem to traditional HPC environments will help address some gaps.

In this tutorial, we explore a light-weight and extensible way to provide the best of both: We utilize the Pilot-Abstraction to execute a diverse set of data-intensive and analytics workloads using Hadoop MapReduce and Spark as well as traditional HPC workloads. The audience will learn how to efficiently use Spark and Hadoop on HPC to carry out advanced analytics tasks, e.g. KMeans and graph analytics, and will understand deployment/performance trade-offs for these tools on HPC.

Large Scale Visualization with ParaView

8:30am-12pm

Room: 13AB

Presenters: *Kenneth Moreland, W. Alan Scott (Sandia National Laboratories); David E. DeMarle (Kitware, Inc.), Ollie Lo (Los Alamos National Laboratory), Joe Insley (Argonne National Laboratory), Rich Cook (Lawrence Livermore National Laboratory)*

ParaView is a powerful open-source turnkey application for analyzing and visualizing large data sets in parallel. Designed to be configurable, extendible, and scalable, ParaView is built upon the Visualization Toolkit (VTK) to allow rapid deployment of visualization components. This tutorial presents the architecture of ParaView and the fundamentals of parallel

visualization. Attendees will learn the basics of using ParaView for scientific visualization with hands-on lessons. The tutorial features detailed guidance in visualizing the massive simulations run on today's supercomputers and an introduction to scripting and extending ParaView. Attendees should bring laptops to install ParaView and follow along with the demonstrations.

MPI+X - Hybrid Programming on Modern Compute Clusters with Multicore Processors and Accelerators

8:30am-12pm

Room: 16AB

Presenters: *Rolf Rabenseifner (High Performance Computing Center Stuttgart), Georg Hager (Erlangen Regional Computing Center)*

Most HPC systems are clusters of shared memory nodes. Such SMP nodes can be small multi-core CPUs up to large many-core CPUs. Parallel programming may combine the distributed memory parallelization on the node interconnect (e.g., with MPI) with the shared memory parallelization inside of each node (e.g., with OpenMP or MPI-3.0 shared memory). This tutorial analyzes the strengths and weaknesses of several parallel programming models on clusters of SMP nodes. Multi-socket-multi-core systems in highly parallel environments are given special consideration. MPI-3.0 introduced a new shared memory programming interface, which can be combined with inter-node MPI communication. It can be used for direct neighbor accesses similar to OpenMP or for direct halo copies, and enables new hybrid programming models. These models are compared with various hybrid MPI+OpenMP approaches and pure MPI. This tutorial also includes a discussion on OpenMP support for accelerators. Benchmark results are presented for modern platforms such as Intel Xeon Phi and Cray XC30. Numerous case studies and micro-benchmarks demonstrate the performance-related aspects of hybrid programming. The various programming schemes and their technical and performance implications are compared. Tools for hybrid programming such as thread/process placement support and performance analysis are presented in a "how-to" section.

Tutorial on Automatic Performance and Energy Tuning with the Periscope Tuning Framework

8:30am-12pm

Room: 19A

Presenters: *Michael Gerndt (Technical University of Munich), Renato Miceli (SENAI CIMATEC)*

A new era in performance engineering of HPC applications is upcoming where traditional manual tuning methods will be challenged by complex performance interactions, dynamic

runtime execution, multi-objective optimizations and the never-ending evolution of hardware architectures. A promising strategy is to apply auto-tuning methodologies enhanced by performance analysis capabilities, sophisticated search techniques and code generation methods.

The AutoTune consortium and its partners propose a half-day tutorial on parallel application performance and energy-efficiency auto-tuning, featuring the Periscope Tuning Framework (PTF), an open-source extensible framework for automatic optimization of HPC applications. The tutorial builds on a long track of tutorials on automatic performance analysis and tuning at conferences such as ISC, EuroMPI, CGO as well as on a series of end-user training workshops.

The target audience are HPC application developers striving for higher performance of their applications as well as performance tool developers seeking for comprehensive tool infrastructures. Based on case studies with real-world applications participants will be empowered to utilize PTF to achieve gains in performance and energy-efficiency for their applications. In the second half we will walk through on how to write your own auto-tuning plugin for PTF enabling tool developers and advanced users to quickly prototype tailored auto-tuners.

A “Hands-On” Introduction to OpenMP

8:30am-5pm

Room: 17B

Presenters: *Tim Mattson (Intel Corporation), J. Mark Bull (Edinburgh Parallel Computing Centre), Mike Pearce (Intel Corporation)*

OpenMP is the de facto standard for writing parallel applications for shared memory computers. With multi-core processors in everything from tablets to high-end servers, the need for multithreaded applications is growing and OpenMP is one of the most straightforward ways to write such programs. In this tutorial, we will cover the core features of the OpenMP 4.0 standard.

We expect students to use their own laptops (with Windows, Linux, or OS/X). We will have access to systems with OpenMP (a remote SMP server), but the best option is for students to load an OpenMP compiler onto their laptops before the tutorial. Information about OpenMP compilers is available at www.openmp.org.

Efficient Parallel Debugging for MPI, Threads, and Beyond

8:30am-5pm

Room: 14

Presenters: *Matthias S. Mueller (Aachen University), David Lecomber (Allinea Software), Bronis R. de Supinski (Lawrence Livermore National Laboratory), Tobias Hilbrich (Technical University of Dresden), Ganesh Gopalakrishnan (University of Utah), Mark O'Connor (Allinea Software)*

Parallel software enables modern simulations on high performance computing systems. Defects—or commonly bugs—in these simulations can have dramatic consequences on matters of utmost importance. This is especially true if defects remain unnoticed and silently corrupt results. The correctness of parallel software is a key challenge for simulations in which we can trust. At the same time, the parallelism that enables these simulations also challenges their developers, since it gives rise to new sources of defects.

This tutorial addresses the efficient removal of software defects. The tutorial provides systematic debugging techniques that are tailored to defects that revolve around parallelism. We present leading-edge tools that aid developers in pinpointing, understanding, and removing defects in MPI, OpenMP, MPI-OpenMP, and further parallel programming paradigms. The tutorial tools include the widely used parallel debugger Allinea DDT, the data race detector Intel Inspector XE, the highly-scalable stack analysis tool STAT, and the two MPI correctness tools MUST and ISP. Hands-on examples—make sure to bring a laptop computer—will guide attendees in the use of these tools. We conclude with a discussion and provide pointers for debugging with paradigms such as CUDA or on architectures such as Xeon Phi.

Fault-Tolerance for HPC: Theory and Practice

8:30am-5pm

Room: 18D

Presenters: *George Bosilca, Aurélien Bouteiller, Thomas Herault (University of Tennessee, Knoxville; Yves Robert (ENS Lyon and University of Tennessee, Knoxville)*

Resilience becomes a critical issue for large-scale platforms. This tutorial provides a comprehensive survey of fault-tolerant techniques for high-performance computing, with a fair balance between practice and theory. It is organized along four main topics: (i) An overview of failure types (software/hardware, transient/fail-stop), and typical probability distributions (Exponential, Weibull, Log-Normal); (ii) General-purpose techniques, which include several checkpoint and rollback recovery protocols, replication, prediction and silent error detection; (iii)

Application-specific techniques, such as ABFT for grid-based algorithms or fixed-point convergence for iterative applications, user-level checkpointing in memory; and (iv) Practical deployment of fault tolerant techniques with User Level Fault Mitigation (a proposed MPI standard extension). Relevant examples based on ubiquitous computational solver routines will be protected with a mix of checkpoint-restart and advanced recovery techniques in a hands-on session.

This tutorial is open to all SC15 attendees who are interested in the current status and expected promise of fault-tolerant approaches for scientific applications. There are no audience prerequisites, but basic knowledge of MPI will be helpful for the hands-on session. Background will be provided for all protocols and probabilistic models.

Linear Algebra Libraries for High-Performance Computing: Scientific Computing with Multicore and Accelerators

8:30am-5pm

Room: 15

Presenters: Jack Dongarra, Jakub Kurzak (University of Tennessee, Knoxville); Michael Heroux (Sandia National Laboratories), James Demmel (University of California, Berkeley)

Today, a desktop with a multicore processor and a GPU accelerator can already provide a TeraFlop/s of performance, while the performance of the high-end systems, based on multicores and accelerators, is already measured in tens of PetaFlop/s. This tremendous computational power can only be fully utilized with the appropriate software infrastructure, both at the low end (desktop, server) and at the high end (supercomputer installation). Most often a major part of the computational effort in scientific and engineering computing goes in solving linear algebra subproblems. After providing a historical overview of legacy software packages, the tutorial surveys the current state-of-the-art numerical libraries for solving problems in linear algebra, both dense and sparse. MAGMA, (D) PLASMA and Trilinos software packages are discussed in detail. The tutorial also highlights recent advances in algorithms that minimize communication, i.e. data motion, which is much more expensive than arithmetic.

Massively Parallel Task-Based Programming with HPX

8:30am-5pm

Room: 17A

Presenters: Hartmut Kaiser, Steven Brandt (Louisiana State University); Alice Koniges (Lawrence Berkeley National Laboratory), Thomas Heller (Frederich-Alexander University), Sameer Shende (University of Oregon), Martin Stumpf (Frederich-Alexander Universität), Dominic Marcello, Zach Byerly, Alireza

Kheirkhahan, Zahra Khatami (Louisiana State University); Zhaoyi Meng (Lawrence Berkeley National Laboratory)

The C++11/14 standard brings significant new parallel programming capability to the language with the introduction of a unified interface for asynchronous programming using futures. This style of programming enables fine-grained constraint-based parallelism, and avoids many load-balancing issues. HPX is a system which builds upon the C++11/14 standard, extending it to distributed operations and increasing its composability. By conforming to the standard, students will learn parallel concepts in a seamless and familiar environment. In this tutorial, students will learn first-hand the capabilities of these revolutionary tools. Firstly, we introduce participants to modern C++11/14 parallel programming, and then we show how to adapt C++ programs to modern massively parallel environments using HPX. Then through presentations, hands-on examples, and tool demonstrations we show how this emerging runtime and language, HPX, allows users to use asynchronous concepts and active messaging techniques to gain performance and productivity in a variety of applications.

Parallel Computing 101

8:30am-5pm

Room: 18B

Presenters: Quentin F. Stout, Christiane Jablonowski (University of Michigan)

This tutorial provides a comprehensive overview of parallel computing, emphasizing those aspects most relevant to the user. It is suitable for new users, managers, students and anyone seeking an overview of parallel computing. It discusses software and hardware/software interaction, with an emphasis on standards, portability, and systems that are widely available.

The tutorial surveys basic parallel computing concepts, using examples selected from multiple engineering and scientific problems. These examples illustrate using MPI on distributed memory systems, OpenMP on shared memory systems, MPI+OpenMP on hybrid systems, GPU and accelerator programming, and Hadoop on big data. It discusses numerous parallelization and load balancing approaches, and software engineering and performance improvement aspects, including the use of state-of-the-art tools.

The tutorial helps attendees make intelligent decisions by covering the options that are available, explaining how they are used and what they are most suitable for. Extensive pointers to the literature and web-based resources are provided to facilitate follow-up studies.

Parallel I/O in Practice

8:30am-5pm

Room: 12A

Presenters: Robert Latham, Robert Ross (Argonne National Laboratory); Brent Welch (Google), Katie Antypas (Lawrence Berkeley National Laboratory)

I/O on HPC systems is a black art. This tutorial sheds light on the state-of-the-art in parallel I/O and provides the knowledge necessary for attendees to best leverage I/O resources available to them. We cover the entire I/O software stack including storage and parallel file systems at the lowest layer, the role of burst buffers (nvram), intermediate layers (such as MPI-IO), and high-level I/O libraries (such as HDF-5). We emphasize ways to use these interfaces that result in high performance and tools for generating insight into these stacks. Benchmarks on real systems are used throughout to show real-world results.

In the first half of the tutorial we cover the fundamentals of parallel I/O. We discuss storage technologies, both present and near-future. Our parallel file systems material covers general concepts and examines three examples: GPFS, Lustre, and PanFS. Our second half takes a more application-oriented focus. We examine the upper library layers of the I/O stack, covering MPI-IO, Parallel netCDF, and HDF5. We discuss interface features, show code examples, and describe how application calls translate into PFS operations. Finally, we discuss tools for capturing and understanding I/O behavior.

Parallel Programming in Modern Fortran

8:30am-5pm

Room: 12B

Presenters: Karla Morris (Sandia National Laboratories), Damian Rouson (Sourcery, Inc.), Salvatore Filippone (University of Rome Tor Vergata), Fernanda S. Foertter (Oak Ridge National Laboratory)

Fortran remains widely used in high-performance computing (HPC) [1], but most users describe their programming skills as self-taught and most use older language versions. The increasing compiler support for modern Fortran makes the time ripe to teach new language features that target HPC. We will teach single-program, multiple-data (SPMD) programming with Fortran 2008 coarrays. We also introduce Fortran's loop concurrency and pure procedure features and demonstrate their use in asynchronous expression evaluation for partial differential equation (PDE) solvers. We incorporate other language features, including object-oriented (OO) programming, when they support parallel programming pedagogy. In particular, we demonstrate OO design patterns for hybrid CPU/GPU calculations in the Parallel Sparse Basic Linear Algebra Subroutines

(PSBLAS) library. Attendees will use the GCC 5.1.0 Fortran compiler and the OpenCoarrays library [2] to compile parallel executables inside virtual machines. Those interested in GPU computing will have access to Oak Ridge's Cray supercomputer Titan.

Effective HPC Visualization and Data Analysis using VisIt

1:30pm-5pm

Room: 18C

Presenters: Cyrus Harrison (Lawrence Livermore National Laboratory), David Pugmire (Oak Ridge National Laboratory), Hank Childs (University of Oregon), Robert Sisneros (University of Illinois at Urbana-Champaign), Amit Chourasia (San Diego Supercomputer Center)

Visualization and data analysis are essential components of the scientific discovery process. Scientists and analysts running HPC simulations rely on visualization and analysis tools for data exploration, quantitative analysis, visual debugging, and communication of results. This half-day tutorial will provide SC15 attendees with a practical introduction to mesh-based HPC visualization and analysis using VisIt, an open source parallel scientific visualization and data analysis platform. We will provide a foundation in basic HPC visualization practices and couple this with hands-on experience using VisIt.

This tutorial includes: 1) An introduction to visualization techniques for mesh-based simulations. 2) A guided tour of VisIt. 3) Hands-on demonstrations visualizing a CFD Blood Flow simulation. 4) Publishing results to the SeedMe platform.

This tutorial builds on the past success of VisIt tutorials, updated for SC15 with new content to help users easily share results and details on recent developments in HPC visualization community. Attendees will gain practical knowledge and recipes to help them effectively use VisIt to analyze data from their own simulations.

Insightful Automatic Performance Modeling

1:30pm-5pm

Room: 18A

Presenters: Alexandru Calotoiu, Felix Wolf (Technical University of Darmstadt); Torsten Hoefler (ETH Zurich), Martin Schulz (Lawrence Livermore National Laboratory)

Many applications suffer from latent performance limitations that may cause them to consume too many resources under certain conditions. Examples include the input-dependent growth of the execution time. Solving this problem requires the ability to properly model the performance of a program to understand its optimization potential in different scenarios. In this tutorial, we will present a method to automatically gener-

ate such models for individual parts of a program from a small set of measurements. We will further introduce a tool that implements our method and teach how to use it in practice. The learning objective of this tutorial is to familiarize the attendees with the ideas behind our modeling approach and to enable them to repeat experiments at home.

Live Programming: Bringing the HPC Development Workflow to Life

1:30pm-5pm

Room: 16AB

Presenters: Ben Swift, Andrew Sorensen, Henry Gardner (Australian National University); Viktor K. Decyk (University of California, Los Angeles)

This tutorial is for any HPC application programmer who has ever made a change to their code and been frustrated at how long it takes to see whether their change worked. We provide an introduction to tools for bringing the near-instant feedback of live programming to the HPC application development workflow. Through worked examples in the Extempore programming environment (<http://github.com/digego/extempore>), this hands-on tutorial will guide participants through the process of taking a scientific code (in C/C++ or Fortran), and running it “live”---so that parameters/subroutines can be examined and even modified with real-time feedback. This opens up a new development workflow for HPC application developers; instead of waiting hours for batch jobs to finish before receiving feedback on changes to the code, incremental modifications can be just-in-time compiled (through LLVM’s efficient JIT compiler) and hot-swapped into the running process. Application developers will discover first-hand the opportunities (and challenges) of a more interactive development workflow, culminating in participating in and interactive programming a “live cluster” running across their laptops by the end of the tutorial.

Power Aware High Performance Computing: Challenges and Opportunities for Application Developers

1:30pm-5pm

Room: 19A

Presenters: Martin Schulz (Lawrence Livermore National Laboratory), Dieter Kranzlmüller (Ludwig-Maximilians-University Munich), Barry Rountree (Lawrence Livermore National Laboratory), David Lowenthal (University of Arizona)

Power and energy consumption are critical design factors for any next generation large-scale HPC system. The costs for energy are shifting the budgets from investment to operating costs, and more and more often the size of systems will be determined by its power needs.

As a consequence, the US Department of Energy (DOE) has set the ambitious limit of 20MW for the power consumption of their first exascale system, and many other funding agencies around the world have expressed similar goals. Yet, with today’s HPC architectures and systems, this is still impossible and far out of reach: the goal will only be achievable through a complex set of mechanisms and approaches at all levels of the hardware and software stack, which will additionally and directly impact the application developer. On future HPC systems, running a code efficiently (as opposed to purely with high performance) will be a major requirement for the user.

In this tutorial, we will discuss the challenges caused by power and energy constraints, review available approaches in hardware and software, highlight impacts on HPC center design and operations, and ultimately show how this change in paradigm from cycle awareness to power awareness will impact application development.

Productive Programming in Chapel: A Computation-Driven Introduction

1:30pm-5pm

Room: 19B

Presenters: Bradford L. Chamberlain, Greg Titus, Michael Ferguson, Lydia Duncan, (Cray Inc.)

Chapel (<http://chapel.cray.com>) is an emerging open-source language whose goal is to vastly improve the programmability of parallel systems while also enhancing generality and portability compared to conventional techniques. Considered by many to be the most promising of recent parallel languages, Chapel is seeing growing levels of interest not only among HPC users, but also in the data analytic, academic, and mainstream communities. Chapel’s design and implementation are portable and open-source, supporting a wide spectrum of platforms from desktops (Mac, Linux, and Windows) to commodity clusters, the cloud, and large-scale systems developed by Cray and other vendors.

This tutorial will provide an in-depth introduction to Chapel’s features using a computation-driven approach: rather than simply lecturing on individual language features, we motivate each Chapel concept by illustrating its use in real computations taken from motivating benchmarks and proxy applications. A pair of hands-on segments will let participants write, compile, and execute parallel Chapel programs, either directly on their laptops (gcc should be pre-installed) or by logging onto remote systems. We’ll end the tutorial by providing an overview of Chapel project status and activities, and by soliciting feedback from participants with the goal of improving Chapel’s utility for their parallel computing needs.

Towards Comprehensive System Comparison: Using the SPEC HPG Benchmarks for Better Analysis, Evaluation, and Procurement of Next-Generation HPC Systems

1:30pm-5pm

Room: 13AB

Presenters: Sandra Wienke (RWTH Aachen University), Sunita Chandrasekaran (University of Houston), Robert Henschel (Indiana University), Ke Wang (University of Virginia), Guido Juckeland (Technical University of Dresden), Oscar Hernandez (Oak Ridge National Laboratory)

The High Performance Group (HPG) of the Standard Performance Evaluation Corporation (SPEC) is a forum for discussing and developing benchmark methodologies for High Performance Computing (HPC) systems. At the same time, the group released production quality benchmark suites like SPEC MPI2007, SPEC OMP2012, and SPEC ACCEL, that can evaluate all dimensions of parallelism. These benchmark suites are used in academia and industry to conduct research in HPC systems and facilitate procurement, testing, and tuning of HPC systems. In this tutorial we first provide an overview of the SPEC HPG benchmark suites and their philosophy. Then we focus in-depth on how to use the benchmarks and how to interpret the results. Participants will learn how to install, compile and run the benchmark and how to submit results to the SPEC website for publication. We will analyze use cases of how the benchmarks can be used to access compiler performance, tune system parameters, evaluate application scalability, compare systems, and monitor power consumption. During the hands-on sessions, participants can gain experience with the benchmarks on a standard HPC cluster and a hybrid Cray XE6/XK7 system. An SSH-capable laptop is required for the hands-on sessions.

Tutorials

Monday, November 16

From Description to Code Generation: Building High-Performance Tools in Python

8:30am-12pm

Room: 18A

Presenters: Andreas Kloeckner (University of Illinois at Urbana-Champaign)

High-performance software is stretched between three ‘tent poles’: high-performance, parallel implementation, asymptotically optimal algorithms, and often highly technical application domains. This tension contributes considerably to making HPC software difficult to write and hard to maintain.

Building abstractions and tools to help maintain separation of concerns through domain-specific mini-languages (‘DSLs’), code generators, and tools is a proven, though difficult, way to help manage the complexity. This tutorial presents a set of tools rooted in the Python language that help with all parts of this process: First, with the design and implementation of DSLs, second, with transformation and rewriting of DSLs into various intermediate representations, third, with the transformation of DSL code into high-performance parallel code on GPUs and heterogeneous architectures, and fourth, with the creation and management of a runtime system for such code. The tutorial includes a brief introduction to Python for those unfamiliar, and it is centered around frequent, short, and interactive practice problems.

Getting Started with In Situ Analysis and Visualization Using ParaView Catalyst

8:30am-12pm

Room: 16AB

Presenters: Andrew C. Bauer (Kitware, Inc.), David H. Rogers (Los Alamos National Laboratory), Jeffrey A. Mauldin (Sandia National Laboratories)

As supercomputing moves towards exascale, scientists, engineers and medical researchers will look for efficient and cost effective ways to enable data analysis and visualization for the products of their computational efforts. The ‘exa’ metric prefix stands for quintillion, and the proposed exascale computers would approximately perform as many operations per second as 50 million laptops. Clearly, typical spatial and temporal data reduction techniques employed for post-processing will not yield desirable results where reductions of 10e3 to 10e6 may still produce petabytes to terabytes of data to transfer or store.

Since transferring or storing data may no longer be viable for many simulation applications, data analysis and visualization must now be performed in situ. This tutorial presents the fundamentals of in situ data analysis and visualization focusing on production use. Attendees will learn the basics of in situ analysis and visualization and be exposed to advanced functionality such as interactive steering and exploratory analysis. We demonstrate the technology using ParaView Catalyst. ParaView Catalyst is a parallel, open-source data analysis and visualization library designed to work with extremely large datasets. It aims to reduce IO by tightly coupling simulation, data analysis and visualization codes.

Practical Hybrid Parallel Application Performance Engineering

8:30am-12pm

Room: 14

Presenters: Markus Geimer (Juelich Research Center), Sameer Shende (University of Oregon), Bert Wesarg (Dresden University of Technology), Brian Wylie (Juelich Research Center)

This tutorial presents state-of-the-art performance tools for leading-edge HPC systems founded on the community-developed Score-P instrumentation and measurement infrastructure, demonstrating how they can be used for performance engineering of effective scientific applications based on standard MPI, OpenMP, hybrid combination of both, and increasingly common usage of accelerators. Parallel performance tools from the Virtual Institute – High Productivity Supercomputing (VI-HPS) are introduced and featured in demonstrations with Scalasca, Vampir, and TAU. We present the complete workflow of performance engineering, including instrumentation, measurement (profiling and tracing, timing and PAPI hardware counters), data storage, analysis, and visualization. Emphasis is placed on how tools are used in combination for identifying performance problems and investigating optimization alternatives. An HPC Linux [<http://www.hpclinux.org>] OVA image containing all of the necessary tools will be available to the participants for use on their own notebook computers (running within a virtual machine). The knowledge gained in this tutorial will help participants to locate and diagnose performance bottlenecks in their own parallel programs.

InfiniBand and High-Speed Ethernet for Dummies

8:30am-12pm

Room: 16AB

Presenters: Dhableswar K. (DK) Panda, Hari Subramoni, Khaled Hamidouche (Ohio State University)

InfiniBand (IB) and High-speed Ethernet (HSE) technologies are generating a lot of excitement towards building next genera-

tion High-End Computing (HEC) systems including clusters, datacenters, file systems, storage, cloud computing and Big Data (Hadoop, Spark, HBase and Memcached) environments. RDMA over Converged Enhanced Ethernet (RoCE) technology is also emerging. This tutorial will provide an overview of these emerging technologies, their offered architectural features, their current market standing, and their suitability for designing HEC systems. It will start with a brief overview of IB and HSE. In-depth overview of the architectural features of IB and HSE (including iWARP and RoCE), their similarities and differences, and the associated protocols will be presented. An overview of the emerging Omni-Path architecture will be provided. Next, an overview of the OpenFabrics stack which encapsulates IB, HSE and RoCE (v1/v2) in a unified manner will be presented. An overview of libfabrics stack will also be provided. Hardware/software solutions and the market trends behind IB, HSE and RoCE will be highlighted. Finally, sample performance numbers of these technologies and protocols for different environments will be presented.

Managing Data Throughout the Research Lifecycle Using Globus Software-as-a-Service

8:30am-12pm

Room: 15

Presenters: Steve Tuecke, Vas Vasiliadis (University of Chicago)

Over the past four years, Globus has become a preferred service for moving and sharing research data on a wide variety of HPC and campus computing resources. With the recent release of data publication and discovery capabilities, Globus now provides useful tools for managing data at every stage of the research lifecycle. While usage across the R&E ecosystem continues to grow, there are many institutions and investigators who are either not aware of the capabilities and benefits Globus can provide, or have limited-scope deployments that they would like to expand.

In this session, participants will learn about the features of the Globus service, and how to use it for delivering robust research data management services that span campus systems, national cyberinfrastructure, and public cloud resources. Globus is installed at most national supercomputing resource providers, and we will draw on experiences from research computing centers (e.g. Michigan, Purdue, Colorado) and HPC facilities (e.g. NCSA Blue Waters, SDSC) to highlight the challenges such facilities face in delivering scalable research data management services. Attendees will be introduced to Globus and will have multiple opportunities for hands-on interaction with the service.

MCDRAM (High Bandwidth Memory) on Knights Landing - Analysis Methods/Tools

8:30am-12pm

Room: 18A

Presenters: Christopher Cantalupo, Karthik Raman, Ruchira Sasanka (Intel Corporation)

Intel's next generation Xeon Phi™ processor family x200 product (Knights Landing) brings in new memory technology, a high bandwidth on package memory called Multi-Channel DRAM (MCDRAM) in addition to the traditional DDR4. MCDRAM is a high bandwidth (~4x more than DDR4), low capacity (up to 16GB) memory, packaged with the Knights Landing Silicon. MCDRAM can be configured as a third level cache (memory side cache) or as a distinct NUMA node (allocatable memory) or somewhere in between. With the different memory modes by which the system can be booted, it becomes very challenging from a software perspective to understand the best mode suitable for an application. At the same time, it is also very essential to utilize the available memory bandwidth in MCDRAM efficiently without leaving any performance on the table. Our tutorial will cover some methods/tools users can exploit to analyze the suitable memory mode for an application. In addition it will also cover the use the “memkind” library interface which is basically a user extensible heap manager built on top of jemalloc. This enables the users to change their application memory allocations to the high bandwidth MCDRAM as opposed to the standard DDR4.

Practical Fault Tolerance on Today's HPC Systems

8:30am-12pm

Room: 18D

Presenters: Kathryn Mohror (Lawrence Livermore National Laboratory), Nathan DeBardeleben (Los Alamos National Laboratory), Laxmikant V. Kale (University of Illinois at Urbana-Champaign), Eric Roman (Lawrence Berkeley National Laboratory)

The failure rates on high performance computing systems are increasing with increasing component count. Applications running on these systems currently experience failures on the order of days; however, on future systems, predictions of failure rates range from minutes to hours. Developers need to defend their application runs from losing valuable data by using fault tolerant techniques. These techniques range from changing algorithms, to checkpoint and restart, to programming model-based approaches. In this tutorial, we will present introductory material for developers who wish to learn fault tolerant techniques available on today's systems. We will give background information on the kinds of faults occurring on today's systems and trends we expect going forward. Following

this, we will give detailed information on several fault tolerant approaches and how to incorporate them into applications. Our focus will be on scalable checkpoint and restart mechanisms and programming model-based approaches.

Advanced MPI Programming

8:30am-5pm

Room: 18B

Presenters: Pavan Balaji (Argonne National Laboratory), William Gropp (University of Illinois at Urbana-Champaign), Torsten Hoefler (ETH Zurich), Rajeev Thakur (Argonne National Laboratory)

The vast majority of production parallel-scientific applications today use MPI and run successfully on the largest systems in the world. For example, several MPI applications are running at full scale on the Sequoia system (on ~1.6 million cores) and achieving 12 to 14 petaflops/s of sustained performance (60 to 70% of peak). At the same time, the MPI standard itself is evolving to address the needs and challenges of future extreme-scale platforms as well as applications. This tutorial will cover several advanced features of MPI, including new MPI-3 features, that can help users program modern systems effectively. Using code examples based on scenarios found in real applications, we will cover several topics including efficient ways of doing 2D and 3D stencil computation, derived datatypes, one-sided communication, hybrid (MPI + shared memory) programming, topologies and topology mapping, and neighborhood and nonblocking collectives. Attendees will leave the tutorial with an understanding of how to use these advanced features of MPI and guidelines on how they might perform on different platforms and architectures.

Advanced OpenMP: Performance and 4.1 Features

8:30am-5pm

Room: 13AB

Presenters: Christian Terboven (RWTH Aachen University), Michael Klemm (Intel Corporation), Ruud van der Pas (Oracle Corporation), Eric J. Stotzer (Texas Instruments Incorporated), Bronis R. de Supinski (Lawrence Livermore National Laboratory)

With the increasing prevalence of multicore processors, shared-memory programming models are essential. OpenMP is a popular, portable, widely supported and easy-to-use shared-memory model. Developers usually find OpenMP easy to learn. However, they are often disappointed with the performance and scalability of the resulting code. This disappointment stems not from shortcomings of OpenMP but rather with the lack of depth with which it is employed. Our “Advanced OpenMP Programming” tutorial addresses this

critical need by exploring the implications of possible OpenMP parallelization strategies, both in terms of correctness and performance. While we quickly review the basics of OpenMP programming, we assume attendees understand basic parallelization concepts. We focus on performance aspects, such as data and thread locality on NUMA architectures, false sharing, and exploitation of vector units. We discuss language features in-depth, with emphasis on advanced features like tasking or cancellation. We close with the presentation of the directives for attached compute accelerators. Throughout all topics we present the new additions of OpenMP 4.1, which will be released during SC15.

How to Analyze the Performance of Parallel Codes 101

8:30am-5pm

Room: 19A

Presenters: *Martin Schulz (Lawrence Livermore National Laboratory); Jim Galarowicz, Don Maghrak (Krell Institute); Mahesh Rajan (Sandia National Laboratories), Matthew LeGendre (Lawrence Livermore National Laboratory), Jennifer Green (Los Alamos National Laboratory)*

Performance analysis is an essential step in the development of HPC codes. It will even gain in importance with the rising complexity of machines and applications that we are seeing today. Many tools exist to help with this analysis, but the user is too often left alone with interpreting the results.

We will provide a practical road map for the performance analysis of HPC codes and will provide users step by step advice on how to detect and optimize common performance problems, covering both on-node performance and communication optimization as well as issues on threaded and accelerator-based architectures. Throughout this tutorial, we will show live demos using Open|SpeedShop, a comprehensive and easy to use tool set. Additionally, we will provide hands-on exercises for attendees to try out the new techniques learned. All techniques will, however, apply broadly to any tool and we will point out alternative tools where useful.

Node-Level Performance Engineering

8:30am-5pm

Room: 18C

Presenters: *Georg Hager, Jan Eitzinger (Erlangen Regional Computing Center); Gerhard Wellein (University of Erlangen-Nuremberg)*

The advent of multi- and many-core chips has led to a further opening of the gap between peak and application performance for many scientific codes. This trend is accelerating as we move from petascale to exascale. Paradoxically, bad

node-level performance helps to “efficiently” scale to massive parallelism, but at the price of increased overall time to solution. If the user cares about time to solution on any scale, optimal performance on the node level is often the key factor. We convey the architectural features of current processor chips, multiprocessor nodes, and accelerators, as far as they are relevant for the practitioner. Peculiarities like SIMD vectorization, shared vs. separate caches, bandwidth bottlenecks, and ccNUMA characteristics are introduced, and the influence of system topology and affinity on the performance of typical parallel programming constructs is demonstrated. Performance engineering and performance patterns are suggested as powerful tools that help the user understand the bottlenecks at hand and to assess the impact of possible code optimizations. A cornerstone of these concepts is the roofline model, which is described in detail, including useful case studies, limits of its applicability, and possible refinements.

OpenACC Programming For Accelerators

8:30am-5pm

Room: 12A

Presenters: *John Urbanic, Tom Maiden (Carnegie Mellon University-Pittsburgh Supercomputing Center); Jeff Larkin, Michael Wolfe (NVIDIA Corporation)*

This tutorial is an upgraded version of the very popular XSEDE OpenACC workshop which has been successfully delivered to over 1200 students at 38 plus institutions. We use exercises to introduce the standard, directive-based methods of programming accelerators: OpenACC and OpenMP 4.0. The techniques learned will be applicable to both GPU and Intel Phi platforms. The course is intense, but only assumes a working knowledge of serial C or Fortran. Experienced CUDA programmers will also find the capabilities of this higher level approach worthwhile. The goal of this tutorial is to allow attendees to walk away with a working knowledge of accelerator programming. Because it is hands-on and exercise driven, attendees will be able to immediately apply these techniques to their own codes.

Portable Programs for Heterogeneous Computing: A Hands-On Introduction

8:30am-5pm

Room: 17B

Presenters: *Tim Mattson (Intel Corporation), Alice Koniges (Lawrence Berkeley National Laboratory), Simon McIntosh-Smith (University of Bristol)*

Heterogeneous computing involves programs running on systems composed of some combination of CPUs, GPUs and other processors critical for high performance computing (HPC).

To succeed as a long-term architecture for HPC, however, it is vital that the center of attention shift away from proprietary programming models to standards-based approaches.

This tutorial is a unique blend of a high-level survey and a hands-on introduction to writing portable software for heterogeneous platforms. We will only work with portable APIs and languages supported by vendor-neutral standards. We believe such an approach is critical if the goal is to build a body of software for use across multiple hardware generations and platforms from competing vendors.

We'll start with OpenCL using the Python interface. By using Python we avoid many of the complexities associated with the host programming API thereby freeing-up more time to focus on the kernel programs that run on the OpenCL devices (such as a GPU or the Intel® Xeon Phi™ coprocessor). Then in the second half of the tutorial, we will shift gears to the OpenMP 4.1 target directive (and associated constructs).

Programming the Xeon Phi

8:30am-5pm

Room: 19B

Presenters: *Jerome Vienne, Victor Eijkhout, Kent Milfeld, Si Liu (The University of Texas at Austin)*

The Intel Xeon Phi co-processor also known as the MIC is becoming more popular in HPC. Current HPC clusters like Tianhe-2, Stampede and Cascade are currently using this technology, and upcoming clusters like Cori and the Stampede upgrade will be comprised of the next generation of MIC known as the Knights Landing (KNL). However, the MIC architecture has significant features that are different from that of current x86 CPUs. It is important for developers to know these differences to attain optimal performance. This tutorial is designed to introduce attendees to the MIC architecture in a practical manner and to prepare them for the new generation of the co-processor (KNL). Experienced C/C++ and Fortran programmers will be introduced to techniques essential for utilizing the MIC architecture efficiently. Multiple lectures and exercises will be used to acquaint attendees with the MIC platform and to explore the different execution modes as well as parallelization and optimization through example testing and reports. All exercises will be executed on the Stampede system at the Texas Advanced Computing Center (TACC). Stampede features more than 2PF of performance using 100,000 Intel Xeon E5 cores and an additional 7+ PF of performance from more than 6,400 Xeon Phi.

Accelerating Big Data Processing with Hadoop, Spark, and Memcached on Modern Clusters

1:30pm-5pm

Room: 16AB

Presenters: *Dhabaleswar K. (DK) Panda, Xiaoyi Lu, Hari Subramoni (Ohio State University)*

Apache Hadoop and Spark are gaining prominence in handling Big Data and analytics. Similarly, Memcached in Web-2.0 environment is becoming important for large-scale query processing. Recent studies have shown default Hadoop, Spark, and Memcached can not leverage the features of modern high-performance computing clusters efficiently, like Remote Direct Memory Access (RDMA) enabled high-performance interconnects, high-throughput and large-capacity parallel storage systems (e.g. Lustre). These middleware are traditionally written with sockets and do not deliver best performance on modern high-performance networks. In this tutorial, we will provide an in-depth overview of the architecture of Hadoop components (HDFS, MapReduce, RPC, HBase, etc.), Spark and Memcached. We will examine the challenges in re-designing networking and I/O components of these middleware with modern interconnects, protocols (such as InfiniBand, iWARP, RoCE, and RSocket) with RDMA and storage architectures. Using the publicly available software packages in the High-Performance Big Data (HiBD, <http://hibd.cse.ohio-state.edu>) project, we will provide case studies of the new designs for several Hadoop/Spark/Memcached components and their associated benefits. Through these case studies, we will also examine the interplay between high-performance interconnects, storage systems (HDD and SSD), and multi-core platforms to achieve the best solutions for these components and Big Data applications on modern HPC clusters.

An Introduction to the OpenFabrics Interface API

1:30pm-5pm

Room: 12B

Presenters: *Sean Hefty (Intel Corporation), Jeffrey M. Squyres (Cisco Systems, Inc.), Robert D. Russell (University of New Hampshire), Howard Pritchard (Los Alamos National Laboratory), Paul Grun (Cray Inc.)*

The OpenFabrics Alliance has been the focal point for an open source project known as the OpenFabrics Interface (OFI) or libfabric. The desire for a new application-centric approach to developing fabric APIs was first expressed at a BoF held during SC13; this tutorial describes the API that the community of OFI developers designed to fulfill the challenges outlined at that BoF. The new library contains a set of fabric APIs that build upon and expand the goals and objectives of the original Verbs API but with a special emphasis on scalability, availability across a range of RDMA-capable networks and on serving

HPC applications among others. Under active development by a broad coalition of industry, academic and national labs partners for two years, the new OFI API has matured to a point where it is ready for use by the SC15 community. This tutorial strengthens the engagement with the community of consumers who will benefit most directly from the new API by providing a high-level introduction, a detailed look at the architecture and hands on experience with running a simple applications over the OFI API.

Data Management, Analysis and Visualization Tools for Data-Intensive Science

1:30pm-5pm

Room: 17A

Presenters: *Norbert Podhorszki, Qing Gary Liu, Scott Klasky, George Ostrouchov, Dave Pugmire, Qing Gary Liu (Oak Ridge National Laboratory)*

This tutorial synergizes three SC14 tutorials: ADIOS, pbdR and VisIt. The goal is to teach users how to achieve high performance (I/O, analytics, and visualization) using a complete software ecosystem. As complexities continue to increase on high-end machines, experimental and observational devices, managing, analyzing, and visualizing “Big Data” efficiently becomes challenging. This tutorial will focus on a software ecosystem (part I) to efficiently write, read, process, and visualize data in-motion and at-rest. Part II of this tutorial introduces parallel I/O (ADIOS, HDF5), file systems, and discusses how to use these efficiently on high-end systems. We will have a hands-on-session using the ADIOS I/O framework to efficiently write and read big data from simulations. Part III will continue with data analysis, covering the pbdR system of packages, which enable parallel processing with R. We also include an approach to this with Python. Finally in Part IV, we will show users how to visualize data at rest and in motion, using VisIt and ParaView.

Debugging and Performance Tools for MPI and OpenMP 4.0 Applications for CPU and Accelerators/Coprocessors

1:30pm-5pm

Room: 14

Presenters: *Sandra Wienke (RWTH Aachen University), Mike Ashworth (STFC Daresbury), Damian Alvarez (Juelich Research Center), Woo-Sun Yang (Lawrence Berkeley National Laboratory); Chris Gottbrath, Nikolay Piskun (Rogue Wave Software, Inc.)*

Scientific developers face challenges adapting software to leverage increasingly heterogeneous architectures. Many systems feature nodes that couple multi-core processors with GPU-based computational accelerators, like the NVIDIA Kepler,

or many-core coprocessors, like the Intel Xeon Phi. In order to effectively utilize these systems, application developers need to demonstrate an extremely high level of parallelism while also coping with the complexities of multiple programming paradigms, including MPI, OpenMP, CUDA, and OpenACC.

This tutorial provides in-depth exploration of parallel debugging and optimization focused on techniques that can be used with accelerators and coprocessors. We cover debugging techniques such as grouping, advanced breakpoints and barriers, and MPI message queue graphing. We discuss optimization techniques like profiling, tracing, and cache memory optimization with tools such as Vampir, Scalasca, Tau, CrayPAT, Vtune and the NVIDIA Visual Profiler. Participants have the opportunity to do hands-on GPU and Intel Xeon Phi debugging and profiling. Additionally, up-to-date capabilities in accelerator and coprocessing computing (e.g. OpenMP 4.0 device constructs, CUDA Unified Memory, CUDA core file debugging) and their peculiarities with respect to error finding and optimization will be discussed. For the hands-on sessions SSH and NX clients have to be installed in the attendees laptops.

Getting Started with Vector Programming using AVX-512 on Multicore and Many-Core Platforms

1:30pm-5pm

Room: 15

Presenters: *Jesus Corbal, Milind Girkar, Shuo Li, John Pennycook, ElMoustapha Ould-Ahmed-Vall, David Mackay, Bob Valentine, Xinmin Tian (Intel Corporation)*

With the recent announcement of AVX-512, the biggest extension to Intel Instruction Set Architecture (ISA), the next generation of Intel’s multicore and many-core product lines will be built around its features such as wider SIMD ALU, more vector registers, new masking architecture for prediction, embedded broadcast and rounding capabilities and the new integer/floating-point instructions. AVX-512 ushers in a new era of converged ISA computing in which the HPC application developer needs to utilize these hardware features through programming tools for maximum performance. This tutorial is the first to bring the AVX-512 ISA to the Supercomputing Community. The first part covers AVX-512 architecture, design philosophy, key features and its performance characteristics. The second part covers the programming tools such as compilers, libraries and the profilers that support the new ISA in a parallel programming framework to guide the developers step-by-step to turn their scalar serial applications into vector parallel applications. Central to the second part is the explicit vector programming methodology under the new industry standard, OpenMP* 4.0 and 4.1. We will present many examples that illustrate how the power of the compiler can be harnessed with minimal user effort to enable SIMD parallelism with AVX-512 instructions from high-level language constructs.

Kokkos: Enabling Manycore Performance Portability for C++ Applications and Domain Specific Libraries/ Languages

1:30pm-5pm

Room: 17A

Presenters: *H. Carter Edwards (Sandia National Laboratories), Jeff Amelang (Harvey Mudd College); Christian Trott, Mark Hoemmen (Sandia National Laboratories)*

The Kokkos library enables applications and domain specific libraries/libraries to implement intra-node thread scalable algorithms that are performance portable across diverse manycore architectures. Kokkos uses C++ template meta-programming, as opposed to compiler extensions or source-to-source translators, to map user code onto architecture-targeted mechanisms such as OpenMP, Pthreads, and CUDA. Kokkos' execution mapping inserts users' parallel code bodies into well-defined parallel patterns and then uses an architecture-appropriate scheduling to execute the computation. Kokkos' data mapping implements polymorphic layout multidimensional arrays that are allocated in architecture-abstracted memory spaces with a layout (e.g., row-major, column-major, tiled) appropriate for that architecture. By integrating execution and data mapping into a single programming model Kokkos eliminates the contemporary array-of-structures versus structure-of-arrays dilemma from user code. Kokkos' programming model consists of the following extensible abstractions: execution spaces

where computations execute, execution policies for scheduling computations, parallel patterns, memory spaces where data is allocated, array layouts mapping multi-indices onto memory, and data access intent traits to portably map data accesses to architecture-specific mechanisms such as GPU texture cache. Tutorial participants will learn Kokkos' programming model through lectures, hands on exercises, and presented examples.

Measuring the Power/Energy of Modern Hardware

1:30pm-5pm

Room: 18A

Presenters: *Victor W. Lee (Intel Corporation), Jee Choi (Georgia Institute of Technology), Kenneth Czechowski (Georgia Institute of Technology)*

This tutorial provides a comprehensive overview of techniques for empirically measuring the power and energy efficiency of modern hardware. In particular, we will focus on experiment methodologies for studying how workload characteristics and system utilization affect power consumption. This course will begin with a basic introduction to the core concepts regarding energy consumption in digital electronics. Next, we will demonstrate step-by-step instructions for measuring power of the full system as well as individual components, such as CPU, DRAM, GPU, and MIC. This will also consist of a "best practices" guide for avoiding common pitfalls. Finally, we will conclude with a brief discussion of related research on power models.

Workshops

Sunday, November 15

Computational and Data Challenges in Genomic Sequencing

9am-12:30pm

Room: Hilton 406

Organizers: Patricia Kovatch (Icahn School of Medicine at Mount Sinai), Dan Stanzione (The University of Texas at Austin), Jaroslaw Zola (University at Buffalo), Ravi Madduri (Argonne National Laboratory), Shane Canon (Lawrence Berkeley National Laboratory), Eric Stahlberg (National Cancer Institute), Steve Bailey (National Institutes of Health), Matthew Vaughn (The University of Texas at Austin), Kjersten Fagnan (Lawrence Berkeley National Laboratory)

As genomic sequencing has proliferated, so has the quantity and complexity of the data generated. To explore how to efficiently and effectively synthesize this data into knowledge, we will bring together an interdisciplinary community of researchers to share state-of-the-art approaches. The interactive workshop will include international speakers from industry, academia, clinics, institutes, funding agencies and laboratories. Moderated panels and lightning talks will cover academic advances and experiences from the field and will be guided by real-time audience feedback. Topics include Strategies for Accelerating Large Scale Genomic Analysis, Approaches for Big Data Genomics Management and Collaborating across Disciplines. The target audience is genomics researchers and those interested in supporting and using big data in the life sciences. We will summarize participant's insights and observations to produce and publicize a report. This workshop is linked to the HPC and Cancer Workshop the following day. (More at: <http://sc15compgenome.hpc.mssm.edu>)

First International Workshop on Heterogeneous Computing with Reconfigurable Logic

9am-12:30pm

Room: Salon D

Organizers: Kevin Skadron (University of Virginia), Torsten Hoefler (ETH Zurich), Michael Lysaght (Irish Centre for High-End Computing), Jason Bakos (University of South Carolina), Michaela Blott (Xilinx Inc.)

The advent of high-level synthesis (HLS) creates exciting new opportunities for using FPGAs in HPC. HLS allows programs written in OpenCL, C, etc. to be mapped directly and effectively to FPGAs, without the need for low-level RTL design. At the same time, FPGA-based acceleration presents the opportunity for dramatic improvements in performance and energy-efficiency for a variety of HPC applications. This workshop will bring together application experts, FPGA experts, and researchers in heterogeneous computing to present cutting-edge research and explore opportunities and needs for future research in this area. (More at: <http://h2rc.cse.sc.edu/>)

Portability Among HPC Architectures for Scientific Applications

9am-12:30pm

Room: Salon G

Organizers: Timothy J. Williams, Tjerk P. Straatsma (Oak Ridge National Laboratory); Katie B. Antypas (Lawrence Berkeley National Laboratory)

This workshop will explore how the large-scale scientific HPC developer community can and should engineer our applications for portability across the distinct HPC architectures of today, the next generation of large systems, and on toward exascale. How should HPC centers advise computational scientists to develop codes that scale well, have good aggregate and single-node performance, yet are as source-code portable as possible across the system architectures. We invite develop-

ers who have already faced this challenge in running across distinct present-day architectures to present their lessons learned and best practices, and developers who are targeting the next generation HPC systems, such as those coming in the 2016-2018 timeframe. How do we handle threads, vectorization, and memory hierarchies in as general a way as possible? Do we need to revise our data structures when they may need to be different to map efficiently on different architectures? (More at: <http://hpcport.alcf.anl.gov>)

DISCS2015: International Workshop on Data Intensive Scalable Computing Systems

9am-5:30pm

Room: Hilton Salon J

Organizers: Philip C. Roth (Oak Ridge National Laboratory), Weikuan Yu (Auburn University), R. Shane Canon (Lawrence Berkeley National Laboratory), Yong Chen (Texas Tech University)

Traditional HPC systems were designed from a compute-centric perspective, with an emphasis on high floating point performance. As scientific and analytics applications become more data intensive, there is a need to rethink HPC system architectures, programming models, runtime systems, and tools with a focus on data intensive computing. Industry approaches supporting data intensive applications have been highly successful, leading many in the HPC community to explore how to apply them. Conversely, the HPC community's expertise in designing, deploying, and using high performance systems is attractive to those in industry. The 2015 International Workshop on Data Intensive Scalable Computing Systems (DISCS-2015) provides a forum for researchers to discuss recent results and the future challenges of running data intensive applications on traditional HPC systems and the latest data-centric computing systems. The workshop includes a keynote address and presentation of peer-reviewed research papers, with ample opportunity for informal discussion. (More at: <http://ft.ornl.gov/discs-2015>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC and free after that to SIGHPC members).

E2SC2015: Energy Efficient Super Computing

9am-5:30pm

Room: Hilton Salon F

Organizers: Darren J. Kerbyson (Pacific Northwest National Laboratory), Kirk W. Cameron (Virginia Polytechnic Institute and State University), Adolfo Hoisie (Pacific Northwest National Laboratory), David K. Lowenthal (University of Arizona), Dimitrios S. Nikolopoulos (Queen's University Belfast), Sudha Yalamanchili (Georgia Institute of Technology)

With exascale systems on the horizon, we have ushered in an era with power and energy consumption as the primary concerns for scalable computing. To achieve a viable Exaflop high performance computing capability, revolutionary methods are required with a stronger integration among hardware features, system software and applications. Equally important are the capabilities for fine-grained spatial and temporal measurement and control to facilitate these layers for energy efficient computing across all layers. Current approaches for energy efficient computing rely heavily on power efficient hardware in isolation. However, it is pivotal for hardware to expose mechanisms for energy efficiency to optimize power and energy consumption for various workloads. At the same time, high fidelity measurement techniques, typically ignored in data-center level measurement, are of high importance for scalable and energy efficient inter-play in different layers of application, system software and hardware. (More at: <http://hpc.pnnl.gov/conf/e2sc/2015/>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

ESPM2: First International Workshop on Extreme Scale Programming Models and Middleware

9am-5:30pm

Room: Hilton 410

Organizers: Dhabaleswar K. (DK) Panda, Khaled Hamidouche, Hari Subramoni (Ohio State University); Karl W. Schulz (Intel Corporation)

Next generation architectures and systems being deployed are characterized by high concurrency, low memory per-core, and multi-levels of hierarchy and heterogeneity. These characteristics bring out new challenges in energy efficiency, fault-tolerance and scalability. It is commonly believed that software has the biggest share of the responsibility to tackle these challenges. In other words, this responsibility is delegated to the next generation programming models and their associated middleware/runtimes. This workshop focuses on different aspects of programming models such as Task-based

parallelism, PGAS (OpenSHMEM, UPC, CAF, Chappel), Directive-based languages (OpenMP, OpenACC), Hybrid MPI+X, etc. It also focuses on their associated middleware (unified runtimes, interoperability for hybrid programming, tight integration with accelerators) for next generation systems and architectures. Objective of ESPM2 workshop is to serve as a forum that brings together researchers from academia and industry working on the areas of programming models, runtime systems, compilation and languages, and application developers. (More at: <http://web.cse.ohio-state.edu/~hamidouc/ESPM2/>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

IA³ 2015: Fifth Workshop on Irregular Applications: Architecture and Algorithms

9am-5:30pm

Room: Hilton Salon A

Organizers: Antonino Tumeo (*Pacific Northwest National Laboratory*), John Feo (*Context Relevant*), Oreste Villa (*NVIDIA Corporation*)

Many data intensive applications are naturally irregular. They may present irregular data structures, control flow or communication. Current supercomputing systems are organized around components optimized for data locality and regular computation. Developing irregular applications on current machines demands a substantial effort, and often leads to poor performance. However, solving these applications efficiently is a key requirement for next generation systems. The solutions needed to address these challenges can only come by considering the problem from all perspectives: from micro- to system-architectures, from compilers to languages, from libraries to runtimes, from algorithm design to data characteristics. Only collaborative efforts among researchers with different expertise, including end users, domain experts, and computer scientists, could lead to significant breakthroughs. This workshop aims at bringing together scientists with all these different backgrounds to discuss, define and design methods and technologies for efficiently supporting irregular applications on current and future architectures. (More at: <http://hpc.pnl.gov/IA3>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

LLVM-HPC2015: Second Workshop on the LLVM Compiler Infrastructure in HPC

9am-5:30pm

Room: Hilton Salon E

Organizer: Hal J. Finkel (*Argonne National Laboratory*)

LLVM, winner of the 2012 ACM Software System Award, has become an integral part of the software-development ecosystem for optimizing compilers, dynamic-language execution engines, source-code analysis and transformation tools, debuggers and linkers, and a whole host of programming-language and toolchain-related components. Now heavily used in both academia and industry, where it allows for rapid development of production-quality tools, LLVM is increasingly used in work targeted at high-performance computing. Research in, and implementation of, programming-language analysis, compilation, execution and profiling has clearly benefited from the availability of a high-quality, freely-available infrastructure on which to build. This second annual workshop will focus on recent developments, from both academia and industry, that build on LLVM to advance the state of the art in high-performance computing. (More at: <http://llvm-hpc2-workshop.github.io/>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

PMBS2015: Sixth International Workshop on Performance Modeling, Benchmarking, and Simulation of High Performance Computing Systems

9am-5:30pm

Room: Hilton Salon K

Organizers: Simon Hammond (*Sandia National Laboratories*); Steven Wright, Stephen A. Jarvis (*University of Warwick*)

This workshop is concerned with the comparison of high-performance computing systems through performance modeling, benchmarking or through the use of tools such as simulators. We are particularly interested in research that reports the ability to measure and make tradeoffs in software/hardware co-design to improve sustained application performance. We are also keen to capture the assessment of future systems.

The aim of this workshop is to bring together researchers, from industry and academia, concerned with the qualitative and quantitative evaluation and modeling of high-performance computing systems. Authors are invited to submit novel research in all areas of performance modeling, benchmarking and simulation, and we welcome research that brings together current theory and practice. We recognize that the

coverage of the term ‘performance’ has broadened to include power/energy consumption and reliability, and that performance modeling is practiced through analytical methods and approaches based on software tools and simulators. (More at: www.pmbsworkshop.org)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC and free after that to SIGHPC members).

PyHPC2015: Fifth Workshop on Python for High-Performance and Scientific Computing

9am-5:30pm

Room: Hilton Salon C

Organizers: *Andreas Schreiber (German Aerospace Center), William Sculling (Argonne National Laboratory), Bill Spatz (Sandia National Laboratories), Andy R. Terrel (Continuum Analytics)*

Python is an established, high-level programming language with a large community in academia and industry. It is a general-purpose language adopted by many scientific applications such as computational fluid dynamics, biomolecular simulation, machine learning, data analysis, scientific visualization etc. The use of Python for scientific, high performance parallel, and distributed computing is increasing since years. Traditionally, system administrators are using Python a lot for automating tasks. Since Python is extremely easy to learn with a very clean syntax, it's well-suited for education in scientific computing. Programmers are much more productive by using Python.

The workshop will bring together researchers and practitioners using Python in all aspects of high performance and scientific computing. The goal is to present Python applications from mathematics, science, and engineering, to discuss general topics regarding the use of Python, and to share experience using Python in scientific computing education. (More information: www.dlr.de/sc/pyhpc2015)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC and free after that to SIGHPC members).

The Sixth International Workshop on Data-Intensive Computing in the Clouds

9am-5:30pm

Room: Hilton 412

Organizers: *Yong Zhao (University of Electronic Science and Technology of China), Wei Tang (Argonne National Laboratory)*

Applications and experiments in all areas of science are becoming increasingly complex and more demanding in terms of their computational and data requirements. Some applications generate data volumes reaching hundreds of terabytes and even petabytes. As scientific applications become more data intensive, the management of data resources and dataflow between the storage and compute resources is becoming the main bottleneck. Analyzing, visualizing, and disseminating these large data sets has become a major challenge and data intensive computing is now considered as the “fourth paradigm” in scientific discovery after theoretical, experimental, and computational science. (More at: datasys.cs.iit.edu/events/DataCloud2015/)

VISTech Workshop 2015: Visualization Infrastructure & Systems Technology

9am-5:30pm

Room: Hilton Salon B

Organizers: *Kelly Gaither (The University of Texas at Austin), Bill Sherman (Indiana University), Madhu Srinivasan (King Abdullah University of Science and Technology)*

Human perception is centered on the ability to process information contained in visible light, and our visual interface is a tremendously powerful data processor. For many types of computational research, the field of visualization is the only viable means of extracting information and developing understanding from this data. Integrating our visual capacity with technological capabilities has tremendous potential for transformational science. We seek to explore the intersection between human perception and large-scale visual analysis through the study of visualization interfaces and interactive displays. This rich intersection includes: virtual reality systems, visualization through augmented reality, large scale visualization systems, novel visualization interfaces, high-resolution interfaces, mobile displays, and visualization display middleware. The VISTech (Visualization Infrastructure & Systems Technology) workshop will provide a space in which experts in the large-scale visualization technology field and users can come together to discuss state-of-the art technologies for visualization and visualization laboratories. (More at: www.tacc.utexas.edu/sc15/workshops/vistech)

WORKS2015: Tenth Workshop on Workflows in Support of Large-Scale Science

9am-5:30pm

Room: Hilton 408

Organizers: *Johan Montagnat (French National Center for Scientific Research), Ian Taylor (Cardiff University)*

Data Intensive Workflows (a.k.a. scientific workflows) are routinely used in most scientific disciplines today, especially in the context of parallel and distributed computing. Workflows provide a systematic way of describing the analysis and rely on workflow management systems to execute the complex analyses on a variety of distributed resources. This workshop focuses on the many facets of data-intensive workflow management systems, ranging from job execution to service management and the coordination of data, service and job dependencies. The workshop therefore covers a broad range of issues in the scientific workflow lifecycle that include: data intensive workflows representation and enactment; designing workflow composition interfaces; workflow mapping techniques that may optimize the execution of the workflow; workflow enactment engines that need to deal with failures in the application and execution environment; and a number of computer science problems related to scientific workflows such as semantic technologies, compiler methods, fault detection and tolerance. (More at: <http://works.cs.cf.ac.uk/>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC and free after that to SIGHPC members).

Computational Approaches for Cancer

2pm-5:30pm

Room: Hilton 406

Organizers: *Eric Stahlberg (Frederick National Laboratory for Cancer Research), Thomas J. Barr (Research Institute at Nationwide Children's Hospital), Patricia Kovatch (Icahn School of Medicine at Mount Sinai)*

As the size, number, variety and complexity of cancer datasets have grown in recent years, new computational challenges and opportunities are emerging within the cancer research and clinical application areas. The workshop focuses on bringing together interested individuals ranging from clinicians, mathematicians, data scientists, computational scientists, hardware experts, engineers, developers, leaders and others with an interest in advancing the use of computation at all levels to better understand, diagnose, treat and prevent cancer. With an interdisciplinary focus, the workshop will provide opportunities for participants to learn about how computation

is employed across multiple areas including imaging, genomics, analytics, modeling, pathology and drug discovery. The forward focus of the workshop looks at challenges and opportunities for large scale HPC, including potential for exascale applications involving cancer. (More at: <http://sc15compcancer.hpc.mssm.edu/>)

Many-Task Computing on Clouds, Grids, and Supercomputers

2pm-5:30pm

Room: Hilton Salon G

Organizers: *Justin M. Wozniak (Argonne National Laboratory), Ioan Raicu (Illinois Institute of Technology), Yong Zhao (University of Electronic Science and Technology of China), Ian Foster (Argonne National Laboratory)*

The 8th workshop on Many-Task Computing on Clouds, Grids, and Supercomputers (MTAGS) will provide the scientific community a dedicated forum for presenting new research, development, and deployment efforts of large-scale many-task computing (MTC) applications on large scale clusters, clouds, grids, and supercomputers. MTC, the theme of the workshop encompasses loosely coupled applications, which are generally composed of many tasks to achieve some larger application goal. This workshop will cover challenges that can hamper efficiency and utilization in running applications on large-scale systems, such as local resource manager scalability and granularity, efficient utilization of raw hardware, parallel file-system contention and scalability, data management, I/O management, reliability at scale, and application scalability. We welcome paper submissions in theoretical, simulations, and systems topics with special consideration to papers addressing the intersection of petascale/exascale challenges with large-scale cloud computing. We invite the submission of original research work of 6 pages. (More at: <http://datasys.cs.iit.edu/events/MTAGS15>)

MLHPC2015: Machine Learning in HPC Environments

2pm-5:30pm

Room: Hilton Salon D

Organizers: *Robert M. Patton, Thomas E. Potok (Oak Ridge National Laboratory); Barry Y. Chen (Lawrence Livermore National Laboratory), Lawrence Carin (Duke University)*

The intent of this workshop is to bring together researchers, practitioners, and scientific communities to discuss methods that utilize extreme scale systems for machine learning. This workshop will focus on the greatest challenges in utilizing HPC for machine learning and methods for exploiting data parallelism, model parallelism, ensembles, and parameter search. We invite researchers and practitioners to participate in this

workshop to discuss the challenges in using HPC for machine learning and to share the wide range of applications that would benefit from HPC powered machine learning. Topics will include but are not limited to: machine learning models, including deep learning, for extreme scale systems; feature engineering; learning large models/optimizing hyper parameters (e.g. deep learning, representation learning); facilitating very large ensembles; applications of machine learning utilizing HPC; and future research challenges for machine learning at large scale. (More at: <http://ornlcta.github.io/MLHPC2015/>)

Refereed proceedings for this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

Monday, November 16

INDIS-15: The Second Workshop on Innovating the Network for Data Intensive Science

9am-12:30pm

Room: Hilton 410

Organizers: Brian L. Tierney (Lawrence Berkeley National Laboratory), Cees de Laat (University of Amsterdam), Matthew J. Zekauskas (Internet2)

Wide area networks are now an integral and essential part of the data driven Supercomputing ecosystem connecting information-sources, processing, simulation, visualization and user communities together. Every year SCinet develops and implements the network for the SC conference. This network is state of the art, connects many demonstrators of big science data processing infrastructures at the highest line speeds and newest technologies available, and demonstrates novel functionality. The show floor network connects to many laboratories and universities worldwide using high-bandwidth connections. This workshop brings together the network researchers and innovators to present challenges and novel ideas that stretch SCinet and network research even further. We invite papers that propose new and novel techniques regarding capacity and functionality of networks, its control and its architecture to be demonstrated at current and future supercomputing conferences. (More at: <https://scinet.supercomputing.org/workshop/>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

Second SC Workshop on Best Practices for HPC Training

9am-12:30pm

Room: Hilton Salon K

Organizers: Fernanda Foertter (Oak Ridge National Laboratory), Rebecca Hartman-Baker (National Energy Research Scientific Computing Center), Scott Lathrop (University of Illinois at Urbana-Champaign), Richard Gerber (National Energy Research Scientific Computing Center), Vincent Betro (University of Tennessee, Knoxville), Robert Whitten (University of Tennessee, Knoxville), Steve Gordon (Ohio Supercomputer Center), Henry Neeman (University of Oklahoma), Barbara Chapman (University of Houston), Nia Alexandrov (Barcelona Supercomputing Center), Jay Alameda (University of Illinois at Urbana-Champaign), Maria-Grazia Giuffreda (Swiss National Supercomputing Center)

Post-petascale systems and future exascale computers are expected to continue hierarchical architectures with nodes of many-core processors and accelerators. Existing programming paradigms and algorithms will have to be adapted to fit wider parallel architectures. The increase in heterogeneity of architectures points to a need in cross-training and collaboration among trainers across centers. Therefore, novel and portable methods of scalable algorithms that better manage the locality of widely distributed data, having fault resilient properties, will become necessary. The critical need to educate the users to develop on the next generation HPC has been widely recognized both in USA and Europe. The aim of this workshop is to bring together the educators and scientists from the key HPC and Supercomputer Centers from USA, Europe, Australia and worldwide to share expertise and best practices for in-person, web-based and asynchronous HPC training tackling the HPC skills gap. (More at: <https://sites.google.com/a/lbl.gov/hpc-training-best-practices/home>)

Ultravis '15: The Tenth Workshop on Ultrascale Visualization

9am-12:30pm

Room: Hilton Salon D

Organizers: Kwan-Liu Ma (University of California, Davis), Venkatram Vishwanath (Argonne National Laboratory), Hongfeng Yu (University of Nebraska-Lincoln)

The output from leading-edge scientific simulations is so voluminous and complex that advanced visualization techniques are necessary to interpret the computed data. Even though visualization technology has progressed significantly in recent years, we are barely capable of exploiting petascale data to its full extent, and exascale datasets are on the horizon. The Ultravis workshop, in its tenth year at SC, aims at addressing this pressing issue by fostering communication between visualization researchers and the users of visualization. Attendees will

be introduced to the latest and greatest research innovations in large data visualization, and also learn how these innovations impact scientific supercomputing and the discovery process. (More at: <http://vis.cs.ucdavis.edu/Ultravis15>)

ATIP Workshop on Chinese HPC Research Toward New Platforms and Real Applications

9am-5:30pm

Room: Hilton 412

Organizers: David Kahaner (Asian Technology Information Program), Depei Qian (Beihang University)

Over the past several years, China has emerged as a significant player in High-Performance Computing (HPC). At least one and possibly up to three very large HPC systems are expected to become operational during 2015. China is also developing component technologies. In addition, while not as advanced, the development of applications is proceeding. Chinese vendors have been making their appearance outside of China, and Chinese scientists are becoming more active in a variety of HPC research areas. ATIP's Workshop will provide a balanced picture of key Chinese developments in HPC. Sessions will cover systems, centers, and government plans; progress in components, system software, and applications will also be included, as well as human resource and ecosystem development. Such a balanced program will enable the US research community to gain a more complete and holistic understanding of the current situation and better evaluate potential opportunities for collaboration. (More at: <http://atip.org/upcoming-events/china-hpc-workshop-at-sc15.html>)

Co-HPC2015: Second International Workshop on Hardware-Software Co-Design for High Performance Computing

9am-5:30pm

Room: Hilton Salon C

Organizers: Shirley Moore (University of Texas at El Paso), Laura Carrington (San Diego Supercomputer Center), Richard Vuduc (Georgia Institute of Technology), Gregory Peterson (University of Tennessee, Knoxville), Theresa Windus (Iowa State University)

Hardware-software co-design involves the concurrent design of hardware and software components of complex computer systems, whereby application requirements influence architecture design and hardware constraints influence design of algorithms and software. Concurrent design of hardware and software has been used for the past two decades for embedded systems to optimize for design constraints such as performance, power, and cost. HPC is facing a similar challenge as we move towards the exascale era, with the necessity of designing systems that run large-scale simulations with high

performance while meeting cost and energy consumption constraints. This workshop will focus on co-design for advanced architectures, including new low-power processors and new memory technologies. (More at: <http://codaash.org/co-hpc/>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC and free after that to SIGHPC members).

ESPT2015: Extreme-Scale Programming Tools

9am-5:30pm

Room: Hilton Salon B

Organizers: Andreas Knuepfer (Dresden University of Technology), Martin Schulz (Lawrence Livermore National Laboratory), Felix Wolf (German Research School for Simulation Sciences), Brian Wylie (Juelich Research Center)

The architectural complexity in HPC is growing and this brings various challenges such as tight power budgets, variability in CPU clock frequencies, load balancing in heterogeneous systems, hierarchical memories and shrinking I/O bandwidths. This is especially prominent on the path to exascale. Therefore, tool support for debugging and performance optimization becomes more necessary than ever. However, the challenges mentioned above also apply to tools development and, in particular, raise the importance of topics such as automatic tuning and methodologies for tools-aided application development. This workshop will serve as a forum for HPC application developers, system designers, and tools researchers to discuss the requirements for exascale-enabled tools and the roadblocks that need to be addressed on the way. We also highly encourage application developers to share their experiences using existing tools. The event will serve as a community forum for all interested in interoperable tool-sets ready for an exascale software stack. (More at: www.vi-hps.org/symposia/Other/espt-sc15.html)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

ExaMPI15: Workshop on Exascale MPI**9am-5:30pm****Room: Hilton Salon F**

Organizers: *Stefano Markidis (KTH Royal Institute of Technology), Erwin Laure (KTH Royal Institute of Technology), William D. Gropp (University of Illinois at Urbana-Champaign), Jesper Larsson Träff (Vienna University of Technology), Masamichi Takagi (RIKEN), Roberto Gioiosa (Pacific Northwest National Laboratory), Mirko Rahn (Fraunhofer Society), Mark Bull (Edinburgh Parallel Computing Centre), Daniel Holmes (Edinburgh Parallel Computing Centre)*

MPI is currently the de-facto standard for HPC systems and applications. On the road to exascale, the workshop investigates if there is a need for re-examination of the Message Passing (MP) model and for exploring new innovative and potentially disruptive concepts and algorithms in MPI. The aim of workshop is to bring together researchers and developers to present and discuss innovative algorithms and concepts in the MP programming model and to create a forum for open and potentially controversial discussions on the future of MPI in the exascale era. Possible workshop topics include innovative algorithms for collective operations, extensions to MPI, including data centric models such as active messages, scheduling/routing to avoid network congestion, “fault-tolerant” communication, interoperability of MP and PGAS models, integration of task-parallel models in MPI and use of MPI in large scale applications. (More at: <https://www.pdc.kth.se/exampi15>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

PDSW10 2015: Tenth Workshop on Parallel Data Storage**9am-5:30pm****Room: Hilton Salon G**

Organizers: *Dries Kimpe (Argonne National Laboratory), Garth Gibson (Carnegie Mellon University), Robert Ross (Argonne National Laboratory), Dean Hildebrand (IBM Corporation)*

Peta- and exascale computing infrastructures make unprecedented demands on storage capacity, performance, concurrency, reliability, availability, and manageability. This one-day workshop focuses on the data storage and management problems and emerging solutions found in peta- and exascale scientific computing environments, with special attention to issues in which community collaboration can be crucial for problem identification, workload capture, solution interoperability, standards with community buy-in, and shared tools. Addressing storage media ranging from tape, HDD, and SSD, to new media like NVRAM, the workshop seeks contributions

on relevant topics, including but not limited to performance and benchmarking, failure tolerance problems and solutions, APIs for high performance features, parallel file systems, high bandwidth storage architectures, support for high velocity or complex data, metadata intensive workloads, autonomies for HPC storage, virtualization for storage systems, archival storage advances, resource management innovations, and incorporation of emerging storage technologies. (More at: <http://www.pdsw.org>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

Runtime Systems for Extreme Scale Programming Models and Architectures (RESPA)**9am-5:30pm****Room: Hilton 400-402**

Organizers: *Siegfried Benkner (University of Vienna), Vivek Sarkar (Rice University)*

Extreme-scale and exascale systems impose new requirements on application developers and programming systems to target platforms with hundreds of homogeneous and heterogeneous cores, as well as energy, data movement and resiliency constraints within and across nodes. Runtime systems can play a critical role in enabling future programming models, execution models and hardware architectures to address these challenges, and in reducing the widening gap between peak performance and the performance achieved by real applications.

The goal of this workshop is to attract leading international researchers to share their latest results involving runtime approaches to address these extreme-scale and exascale software challenges. The scope of the workshop includes (but is not limited to) runtime system support for: high-level programming models and domain-specific languages; scalable intra-node and inter-node scheduling; memory management across coherence domains and vertical hierarchies of volatile/non-volatile storage; optimized locality and data movement; energy management and optimization; performance tuning; and, resilience. (More at: <http://respa15.rice.edu>)

SCalA15: Sixth Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems

9am-5:30pm

Room: Hilton Salon E

Organizers: Vassil Alexandrov (Barcelona Supercomputing Center), Jack Dongarra (University of Tennessee, Knoxville), Al Geist (Oak Ridge National Laboratory), Christian Engelmann (Oak Ridge National Laboratory)

Novel scalable scientific algorithms are needed to enable key science applications to exploit the computational power of large-scale systems, in particular, the current tier of leading petascale machines and the road to exascale as HPC systems continue to scale up in compute node and processor core count. These extreme-scale systems require novel scientific algorithms to hide network and memory latency, have very high computation/communication overlap, have minimal communication, have no synchronization points. With the advent of heterogeneous compute nodes employing standard processors and GPGPUs, scientific algorithms need to match these architectures to extract the most performance. Additionally, with the advent of Big Data, the requirement of key science applications employing such scalable mathematical methods and algorithms that are able to handle compute intensive applications and applications with Big Data at scale and are able to address fault-tolerance and resilience challenges of current- and future-generation extreme-scale HPC systems becomes tremendously important. (More at: www.csm.ornl.gov/srt/conferences/Scala/2015/)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

Sixth Annual Workshop for the Energy Efficient HPC Working Group (EE HPC WG)

9am-5:30pm

Room: Hilton Salon A

Organizers: Natalie Bates (Energy Efficient HPC Working Group), Anna Maria Bailey (Lawrence Livermore National Laboratory), Stephen Poole (Department of Defense), Herbert Huber (Leibniz Supercomputing Center), James H. Rogers (Oak Ridge National Laboratory), Susan Coghlan (Argonne National Laboratory), James Laros (Sandia National Laboratories), Daniel Hackenberg (Dresden University of Technology), Francis Belot (French Alternative Energies and Atomic Energy Commission), Josip Loncaric (Los Alamos National Laboratory), David Martinez (Sandia National Laboratories), Thomas Durbin (University of Illinois at Urbana-Champaign), Steve Martin (Cray Inc.), Ramkumar Nagappan (Intel Corporation), Nicolas Dube (Hewlett-Packard Development Company, L.P.), Ingmar Meijer (IBM Corporation), Marriann Silveira (Lawrence Livermore National Laboratory), Andres Marquez (Pacific Northwest National Laboratory)

This annual workshop is organized by the Energy Efficient HPC Working Group (<http://eehpcwg.lbl.gov/>). It provides a strong blended focus that includes both the facilities and system perspectives; from architecture through design and implementation. The topics reflect the activities and interests of the EE HPC WG, which is a group with over 500 members from ~20 different countries. (More at: https://eehpcwg.llnl.gov/pages/conf_sc15.htm)

Sixth SC Workshop on Big Data Analytics: Challenges and Opportunities (BDAC-15)

9am-5:30pm

Room: Hilton Salon J

Organizers: Ranga Raju Vatsavai (North Carolina State University), Scott Klasky (Oak Ridge National Laboratory), Manish Parashar (Rutgers University)

The recent decade has witnessed data explosion, and petabyte sized data archives are not uncommon any more. It is estimated that organizations with high end computing (HEC) infrastructures and data centers are doubling the amount of data that they are archiving every year. On the other hand computing infrastructures are becoming more heterogeneous. The first six workshops held with SC10 to SC14 are a great success. Continuing on this success, in addition to the cloud focus, we propose to broaden the topic of this workshop with an emphasis on middleware infrastructure that facilitate efficient data analytics on big data. The proposed workshop intends to bring together researchers, developers, and practitioners from academia, government, and industry to discuss new and

emerging trends in high end computing platforms, programming models, middleware and software services, and outline the data mining and knowledge discovery approaches that can efficiently exploit this modern computing infrastructure. (More at: <http://research.csc.ncsu.edu/stac/conferences/BDAC-SC15/>)

WACCPD: Workshop on Accelerator Programming Using Directives

9am-5:30pm

Room: Hilton 406

Organizers: Sunita Chandrasekaran (University of Houston), Fernanda Foertter (Oak Ridge National Laboratory)

Directive-based programming models offer scientific applications a path on to HPC platforms without undue loss of portability or programmer productivity. Using directives, application developers can port their codes to the accelerators incrementally while minimizing code changes. Challenges remain because the directives models need to support a rapidly evolving array of hardware with diverse memory subsystems, which may or may not be unified. The programming model will need to adapt to such developments, make improvements to raise its performance portability that will make accelerators as first-class citizens for HPC. Such improvements are being continuously discussed within the standard committees such as OpenMP and OpenACC. This workshop aims to capture the assessment of the improved feature set, their implementations and experiences with their deployment in HPC applications. The workshop aims at bringing together the user and tools community to share their knowledge and experiences of using directives to program accelerators. (More at: http://openacc.org/WACCPD/call_for_papers)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

WOLFHPC15: Fifth International Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing

9am-5:30pm

Room: Hilton 408

Organizers: Sriram Krishnamoorthy (Pacific Northwest National Laboratory), Jagannathan Ramanujam (Louisiana State University), P. Sadayappan (Ohio State University)

Multi-level heterogeneous parallelism and deep memory hierarchies in current and emerging computer systems makes their programming very difficult. Domain-specific languages

(DSLs) and high-level frameworks (HLFs) provide convenient abstractions, shielding application developers from much of the complexity of explicit parallel programming in standard programming languages like C/C++/Fortran. However, achieving scalability and performance portability with DSLs and HLFs is a significant challenge. This workshop seeks to bring together developers and users of DSLs and HLFs to identify challenges and discuss solution approaches for their effective implementation and use on massively parallel systems. More at: <http://hpc.pnl.gov/conf/wolfhpc/2015/>

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

EduHPC2015: Workshop on Education for High Performance Computing

2pm-5:30pm

Room: Hilton Salon K

Organizers: Sushil K. Prasad (Georgia State University), Anshul Gupta (IBM Corporation), Arnold Rosenberg (Northeastern University), Alan Sussman (University of Maryland), Charles Weems (University of Massachusetts), Almadena Chitchelekanova (National Science Foundation)

The EduHPC Workshop is devoted to the development and assessment of educational resources for undergraduate education in High Performance Computing (HPC) and Parallel and Distributed Computing (PDC). Both PDC and HPC now permeate the world of computing to a degree that makes it imperative for even entry-level computer professionals to incorporate these computing modalities into their computing kitbags, no matter what aspect of computing they work on. This workshop focuses on the state of the art in HPC and PDC education, by means of both contributed and invited papers from academia, industry, and other educational and research institutions. Topics of interest include all topics pertaining to the teaching of PDC and HPC within Computer Science and Engineering, Computational Science, and Domain Science and Engineering curricula. The emphasis of the workshop is undergraduate education, but fundamental issues related to graduate education are also welcome. (More at: <http://cs.gsu.edu/~tcpp/curriculum/?q=edupdhpc>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

ISAV2015: First Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization

2pm-5:30pm

Room: Hilton Salon D

Organizers: E. Wes Bethel (Lawrence Berkeley National Laboratory), Venkatram Vishwanath (Argonne National Laboratory), Gunther H. Weber (Lawrence Berkeley National Laboratory), Matthew Wolf (Georgia Institute of Technology)

The considerable interest in the HPC community regarding in situ analysis and visualization is due to several factors. First is an I/O cost savings, where data is analyzed/visualized while being generated, without first storing to a filesystem. Second is the potential for increased accuracy, where fine temporal sampling of transient analysis might expose some complex behavior missed in coarse temporal sampling. Third is the ability to use all available resources, CPU's and accelerators, in the computation of analysis products.

The workshop brings together researchers, developers and practitioners from industry, academia, and government laboratories using in situ methods in extreme-scale, high performance computing. The goal is to present existing in-situ infrastructures, reference examples in a range of science and engineering applications, to discuss topics like opportunities presented by new architectures; existing infrastructure needs, requirements, and gaps; and experiences to foster and enable in situ analysis and visualization. (More at: <http://vis.lbl.gov/Events/ISAV-2015/>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

NDM15: Fifth International Workshop on Network-Aware Data Management

2pm-5:30pm

Room: Hilton 410

Organizers: Mehmet Balman (VMware, Inc and Lawrence Berkeley National Laboratory), Surendra Byna (Lawrence Berkeley National Laboratory), Brian L. Tierney (Lawrence Berkeley National Laboratory)

Today's data centric environment both in industry and scientific domains depends on the underlying network infrastructure and its performance, to deliver highly distributed extreme-scale application workloads. As current technology enables faster storage devices and larger interconnect bandwidth, there is a substantial need for novel system design and middleware architecture to address increasing latency and scalability

requirements. Furthermore, limitations in end-system architecture and system software design play an important role in many-core platforms. Traditional network and data management techniques are unlikely to scale to meet the needs of future data-intensive systems. We require new collaborations between data management and networking communities to develop intelligent networking middleware and efficient data management infrastructure. This workshop seeks contributions from academia, government, and industry to discuss future design principles of network-aware data management. We focus on emerging trends in resource coordination, data-aware scheduling, storage technologies, end-to-end performance, network-aware workflow management, high-performance networking, and network virtualization. (More at: <http://2015.ndm-meeting.org>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

Friday, November 20

HUST2015: Second International Workshop on HPC User Support Tools

8:30am-12pm

Room: Hilton Salon B

Organizers: Ralph Christopher Bording (Pawsey Supercomputing Center), Gamblin (Lawrence Livermore National Laboratory), Vera Hansper (Victorian Life Sciences Computation Initiative)

Supercomputing centers exist to drive scientific discovery by supporting researchers in computational science fields. To make users more productive in the complex HPC environment, HPC centers employ user support teams. These teams serve many roles, from setting up accounts, to consulting on math libraries and code optimization, to managing HPC software stacks. Often, support teams struggle to adequately support scientists. HPC environments are extremely complex, and combined with the complexity of multi-user installations, exotic hardware, and maintaining research software, supporting HPC users can be extremely demanding.

With the second HUST workshop, we will continue to provide a necessary forum for system administrators, user support team members, tool developers, policy makers and end users. We will provide fora to discuss support issues and we will provide a publication venue for current support developments. Best practices, user support tools, and any ideas to streamline user support at supercomputing centers are in scope. (More at: <http://hust15.github.io>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

NRE2015: Numerical Reproducibility at Exascale

8:30am-12pm

Room: Hilton 400-402

Organizers: Michael Mascagni (Florida State University and National Institute for Standards and Technology), Walid Keyrouz (National Institute of Standards and Technology)

A cornerstone of the scientific method is experimental reproducibility. As computation has grown into a powerful tool for scientific inquiry, the assumption of computational reproducibility has been at the heart of numerical analysis in support of scientific computing. With ordinary CPUs, supporting a single, serial, computation, the ability to document a numerical result has been a straight-forward process. However, as computer hardware continues to develop, it is becoming harder to ensure computational reproducibility, or to even completely document a given computation. This workshop will explore the current state of computational reproducibility in HPC, and will seek to organize solutions at different levels. The workshop will conclude with a panel discussion aimed at defining the current state of computational reproducibility for the Exascale. We seek contributions in the areas of computational reproducibility in HPC. (More at: <http://www.nist.gov/itl/ssd/is/numreprod2015.cfm>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

Producing High Performance and Sustainable Software for Molecular Simulation

8:30am-12pm

Room: Hilton 408

Organizers: Arno Proeme (Edinburgh Parallel Computing Centre), Lorna Smith (Edinburgh Parallel Computing Centre), Mario Antonioletti (Edinburgh Parallel Computing Centre), Neil Chue Hong (Edinburgh Parallel Computing Centre), Weronika Filinger (Edinburgh Parallel Computing Centre), Teresa Head-Gordon (University of California, Berkeley), Jay Ponder (Washington University), Jonathan Essex (University of Southampton)

Molecular simulation software continues to account for a large percentage of HPC resource utilisation by the scientific community. It has taken significant effort over the past decade to equip these codes with strategies to exploit the parallelism on offer and continue delivering cutting-edge discoveries. As we

head towards exascale this workshop brings together developers and representatives of major molecular simulation software efforts as well as high-performance computing experts and researchers of relevant numerical methods and algorithms in order to identify key current and future performance bottlenecks and discuss the challenges faced in creating and maintaining sustainable high-performance molecular simulation software. The workshop will provide a forum for an exchange of ideas about how best to achieve high performance and what to aim for in the coming decade. It will also encourage discussion of software development practices that can support these goals as well as benchmarking, testing and performance comparisons. (More at: <http://www.apes-soft.org/sc15>)

SE-HPCCSE2015: Third International Workshop on Software Engineering for High Performance Computing in Computational Science and Engineering

8:30am-12pm

Room: Hilton Salon A

Organizers: Jeffrey Carver (University of Alabama), Neil Chue Hong (University of Edinburgh), Selim Ciraci (Microsoft Corporation)

Researchers are increasingly using high performance computing (HPC), including GPGPUs and computing clusters, for computational science & engineering (CSE) applications. Unfortunately, when developing HPC software, developers must solve reliability, availability, and maintainability problems in extreme scales, understand domain specific constraints, deal with uncertainties inherent in scientific exploration, and develop algorithms that use computing resources efficiently. Software engineering (SE) researchers have developed tools and practices to support development tasks, including: validation & verification, design, requirements management and maintenance. HPC CSE software requires appropriately tailored SE tools/methods. The SE-HPCCSE workshop addresses this need by bringing together members of the SE and HPC CSE communities to share perspectives, present findings from research and practice, and generating an agenda to improve tools and practices for developing HPC CSE software. This workshop builds on the success of the 2013 and 2014 editions. (More at: <http://se4science.org/workshops/sehpcse15/>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

Software Defined Networking (SDN) for Scientific Networking

8:30am-12pm

Room: Hilton Salon D

Organizers: Nick Buraglio (*Energy Sciences Network*), Dale Carder (*University of Wisconsin-Madison*), Anita Nikolich (*National Science Foundation and Kentucky State University*)

High-speed networking is a critical component of global scientific collaborations. However, high speed links in and of themselves are not enough to ensure successful scientific discovery. As scientific experiments create more complex workflows over distributed high capacity networks, campuses, labs and supercomputing centers are faced with the challenge of traffic optimization for scientific data. Software Defined Networking (SDN) is a promising solution to this challenge. SDN has gone beyond the experimental phase and into production at many sites. This workshop will explore interesting uses of SDN for network optimization, address security challenges with SDN and discuss emerging SDN architectures such as the Software Defined Exchange.

SDN for Scientific Networking will use reviewed short papers, keynote speakers, panels to provide a forum for discussing these challenges, including both positions and experiences. All material and discussions will be archived for continued discussion. (More at: <https://es.net/network-r-and-d/workshops>.)

VPA2015: Second International Workshop on Visual Performance Analysis

8:30am-12pm

Room: Hilton 410

Organizers: Peer-Timo Bremer (*Lawrence Livermore National Laboratory*), Bernd Mohr (*Juelich Research Center*), Valerio Pascucci (*University of Utah*), Martin Schulz (*Lawrence Livermore National Laboratory*)

Over the last decades incredible amounts of resources have been devoted to building ever more powerful supercomputers. However, exploiting the full capabilities of these machines is becoming exponentially more difficult with each generation of hardware. To understand and optimize the behavior of massively parallel simulations the performance analysis community has created a wide range of tools to collect performance data, such as flop counts or network traffic at the largest scale. However, this success has created a new challenge, as the resulting data is far too large and too complex to be analyzed in a straightforward manner. Therefore, new automatic analysis and visualization approaches must be created to allow application developers to intuitively understand the multiple, interdependent effects that their algorithmic choices have on the final performance. This workshop will bring together researchers

from performance analysis and visualization to discuss new approaches of combining both areas to understand large-scale applications. (More at: <http://cedmav.com/events/vpa-2015.html>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

WHPCF2015: Eighth Workshop on High Performance Computational Finance

8:30am-12pm

Room: Hilton 406

Organizers: Jose E. Moreira (*IBM Corporation*), Matthew F. Dixon (*University of San Francisco*), Mohammad Zubair (*Old Dominion University*)

The purpose of this workshop is to bring together practitioners, researchers, vendors, and scholars from the complementary fields of computational finance and high performance computing, in order to promote an exchange of ideas, develop common benchmarks and methodologies, discuss future collaborations and develop new research directions. Financial companies increasingly rely on high performance computers to analyze high volumes of financial data, automatically execute trades, and manage risk.

Recent years have seen the dramatic increase in compute capabilities across a variety of parallel systems. The systems have also become more complex with trends towards heterogeneous systems consisting of general-purpose cores and acceleration devices. The workshop will enable the dissemination of recent advances and findings in the application of high performance computing to computational finance among researchers, scholars, vendors and practitioners, and will encourage and highlight collaborations between these groups in addressing high performance computing research challenges. (More at: <http://ewh.ieee.org/conf/whpcf/>)

Refereed proceedings from this workshop are available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members).

Women in HPC: Changing the Face of HPC**8:30am-12pm****Room: Hilton 412**

Organizers: Toni Collis (Edinburgh Parallel Computing Centre), Barbara Chapman (University of Houston), Daniel Holmes (Edinburgh Parallel Computing Centre), Lorna Smith (Edinburgh Parallel Computing Centre), Alison Kennedy (Edinburgh Parallel Computing Centre), Adrian Jackson (Edinburgh Parallel Computing Centre), Julia Andrys (Murdoch University), Jesmin Jahan Tithi (Stony Brook University), Rapela Regina Maphanga (University of Limpopo), Sunita Chandrasekaran (University of Houston), Rebecca Hartman-Baker (Lawrence Berkeley National Laboratory)

Building on the success of the first international workshop for Women in HPC at SC14 this workshop will bring together leading women working in HPC and female early career researchers. The workshop will showcase HPC research by early career women on a broad range of topics from any area of study in HPC, as well as discussing the importance of gender equality.

The workshop will open with an introduction on current research by the Women in HPC network including up-to-date demographics. We will also host three panel discussions:

1. Women in Industry. Are there specific barriers to women working in the industry in HPC?
2. Career progression. Are women more likely to move into project management roles than men?
3. Are women the problem or do we also have 'the wrong kind of men'?

More at: <http://www.womeninhpc.org.uk/SC15>

Acknowledgements

SC Planning Committee

Conference Management

Conference Chair

Jackie Kern, University of Illinois at Urbana-Champaign

Vice Chair

Becky Verastegui, Oak Ridge National Laboratory

Executive Director

Janet M. McCord, The University of Texas at Austin

Deputy Chair

John West, The University of Texas at Austin

Assistant to the Chair

Clea Marples, Lawrence Livermore National Laboratory
Karl Kern
Maisie Kern
Sherry Sollers
Ashleigh Swanson
Jon Swanson

HPC Matters

Brian Ban, Ban Communications
Trish Damkroger, Lawrence Livermore National Laboratory
Wilfred Pinfold, Concurrent Systems LLC
Lauren Rotman, Lawrence Berkeley National Laboratory

Committee Suite

Mary Amiot, Northstar Event Management

Executive

Conference Chair

Jackie Kern, University of Illinois at Urbana-Champaign

Vice Chair

Becky Verastegui, Oak Ridge National Laboratory

Executive Director

Janet M. McCord, The University of Texas at Austin

Deputy Chair

John West, The University of Texas at Austin

Society Participants

Donna Cappel, Association for Computing Machinery
Ashley Cozzi, Association for Computing Machinery
Anne Marie Kelly, IEEE Computer Society
Carmen Saliba, CMP, IEEE Computer Society

Keynote Chair

William Kramer, University of Illinois at Urbana-Champaign

Donation Coordinator

Brent Gorda, Intel Corporation

Technical Program Chair

Jeffrey S. Vetter, Oak Ridge National Laboratory and Georgia Institute of Technology

Communications Chair

Christine E. Cuicchi

Finance Chair

Eric Sills, North Carolina State University

Exhibits Chair

Trey Breckenridge, Mississippi State University

Infrastructure Chair

Jamie Van Randwyk, Lawrence Livermore National Laboratory
Tim Yeager

Student Programs Chair

Jeanine Cook, Sandia National Laboratories

SCinet Chair

Davey Wheeler, University of Illinois at Urbana-Champaign

Education Coordinator

Scott Lathrop, Shodor and University of Illinois at Urbana-Champaign

Communications

Communications Chair

Christine E. Cuicchi

Communications Deputy

Jon Bashor, Lawrence Berkeley National Laboratory

Deputy Communications Chair

Leah Glick, Linklings, LLC

Conference Publications Coordinator

Vivian M. Benton, Pittsburgh Supercomputing Center

Program Submissions Editor

Janet Brown, Pittsburgh Supercomputing Center

Web Content Manager

Mary Hester, Lawrence Berkeley National Laboratory

Graphics and Website Design Contractor

Carlton Bruett, Carlton Bruett Design, LLC
Chris Hamb, Carlton Bruett Design, LLC

Media Partnerships

Mike Bernhardt, Intel Corporation

Media Relations Consultant

Brian Ban, Ban Communications

Newsletter Editor

Leah Glick, Linklings, LLC

Writer

John West, The University of Texas at Austin

Info Booth/Wayfinding

Douglas Fuller, Red Hat

Social Media Coordinator

Linda Vu, Lawrence Berkeley National Laboratory

Social Media Liaisons

Treshea N. Wade, Carnegie Mellon University
Shandra Williams, Pittsburgh Supercomputing Center

Mobile App Coordinator

Ian MacConnell, Battelle

hpcmatters Liaison

Lauren Rotman, Lawrence Berkeley National Laboratory

Visual Production

Ralph Sulser, Bill Thomas Associates
Franki Thomas, TCG Visual

Exhibits**Exhibits Chair**

Trey Breckenridge, Mississippi State University

Exhibits Deputy Chair

Matt Link, Indiana University

Exhibits Management

Pete Erickson, Hall-Erickson, Inc.
Paul Graller, Hall-Erickson, Inc.
Ann Loporini, Hall-Erickson, Inc.
Chrissy Petracek, Hall-Erickson, Inc.
Mike Weil, Hall-Erickson, Inc.

Event Services

Stephen E. Hagstette Jr, Freeman Company
Regina Martin, Freeman Company
Darryl Monahan, Freeman Company

Exhibitor Forum Chair

John Cazes, The University of Texas at Austin
Paul Domagala, Argonne National Laboratory

HPC Impact Showcase Chair

Christy Adkinson, Cray Inc.
David Martin, Argonne National Laboratory

Exhibits Presence Coordinator

Kelly Gaither, The University of Texas at Austin

Financial Management**Finance Chair**

Eric Sills, North Carolina State University

Finance Advisor

Sandra Huskamp

Meetings and Local Arrangements Chair

Janet M. McCord, The University of Texas at Austin

Catering and Events**Catering/Events Chair**

Barbara Horner-Miller, BHM Consulting LLC

Catering/Events Contractor

Carole Garner, MeetGreen
Rebecca Mebane, MeetGreen
Nicole Morris-Judd, MeetGreen
Linda Snyder, MeetGreen

Finance**Finance Contractor**

Danielle Bertrand, Talbot, Korvola & Warwick, LLP
Donna Martin, Talbot, Korvola & Warwick, LLP
Anne Nottingham, Talbot, Korvola & Warwick, LLP
Brad Rafish, Talbot, Korvola & Warwick, LLP

Housing**Housing Chair**

Cherri M. Pancake, Oregon State University

Housing Contractor

Mike Farr, Orchid Event Solutions

Deputy Housing Chair

Bronis R. de Supinski, Lawrence Livermore National Laboratory

Registration/Store/ Merchandise**Registration/Store/ Merchandise Chair**

Michele Bianchini-Gunn, Lawrence Livermore National Laboratory

Registration Contractor

Karen Shipe, J. Spargo & Associates

Merchandise Contractor

Linda Sheeran, Linda Sheeran Promotions

Infrastructure**Infrastructure Chair**

Jamie Van Randwyk, Lawrence Livermore National Laboratory
Tim Yeager

AV/PC**AV/PC Chair**

James W. Ferguson, University of Tennessee, Knoxville

AV Concepts Contractor

Gabe Blakeney, AV Concepts
Danielle Martinez, AV Concepts
Thomas Nichols, AV Concepts
Richard Steinau, AV Concepts

Committee Office**Office Staff**

Valori Archuleta, The University of Texas at Austin
Katie Cohen, Texas Advanced Computing Center
Beth McKown, University of Illinois at Urbana-Champaign
Carolyn Peters, Argonne National Laboratory

Electrical**Electrical Chair**

Gary New, National Center for Atmospheric Research

Information Technology**Information Technology Contractor**

Luke Montague, Linklings LLC
Mark Montague, Linklings LLC

Security**Security Chair**

Philip C. Roth, Oak Ridge National Laboratory

Security Member(s)

Peggy Cler, BodyWork Associates, Inc.
Eric Greenwade, Microsoft Corporation
Dustin Leverman, Oak Ridge National Laboratory
Matt Link, Indiana University
Kate Mace, Clemson University

Security / Communications Liaison

Lauren Rotman, Lawrence Berkeley National Laboratory

Coat Check/Lost & Found**Coat Check/Lost & Found Mgr**

Celine Andre, RA Consulting
Laura Hubbard, RA Consulting

Security**Security Contractor**

Peter Alexan, RA Consulting
Celine Andre, RA Consulting
Wendi Ekblade, RA Consulting
Sam Ridley, RA Consulting
Jerome Williams, RA Consulting

Security SCinet Contractor

Guy Julian, RA Consulting
Aaron Parker, RA Consulting

Signage**Signage Chair**

Kevin Wohlever, OpenFPGA

Airport Signage Lead

Dustin Leverman, Oak Ridge National Laboratory

Space**Space Chair**

Jay Cliburn, SMS

SCinet**SCinet Chair**

Davey Wheeler, University of Illinois at Urbana-Champaign

SCinet Deputy Chair

Corby Schmitz, Argonne National Laboratory

SCinet Executive Director

Linda Winkler, Argonne National Laboratory

SCinet Vice Chair

Brandon George, DataDirect Networks

SCinet Finance Chair

Ralph A. McElDowney, Department of Defense
HPC Modernization Program

Architecture**Architecture Chair**

Linda Winkler, Argonne National Laboratory

Architecture Team Member(s)

Jon Dugan, ESnet / LBNL

Commodity Network**Commodity Network Co-Chair**

Kevin Hayden, Argonne National Laboratory
Jeffrey R. Schwab, Purdue University

Commodity Network Team**Member(s)**

Ryan Braithwaite, Los Alamos National Laboratory
Parks Fields, Los Alamos National Laboratory
Mary Hester, Lawrence Berkeley National Laboratory
Amy Liebowitz, University of Michigan
Dave Modl, Los Alamos National Laboratory

Communications**Communications Chair**

Mary Hester, Lawrence Berkeley National Laboratory

Communications Team Member(s)

Sylvia Kuijpers, SURFnet

Fiber**Fiber Co-Chair**

Kevin Hayden, Argonne National Laboratory
Annette E. Kitajima, Sandia National Laboratories

Fiber Team Member(s)

Joshua D. Alexander, University of Oklahoma
Chris Cavallo, Lawrence Berkeley National Laboratory and Energy Sciences Network
Jason Charcalla, University of Tennessee, Knoxville
DA Fye, Oak Ridge National Laboratory
Kent M. Goldsmith, CABLExpress Corporation
Rebecca Hutchinson, SCinet
Camille Magno, JDSU
Sarah Marie Neuwirth, University of Heidelberg
Scott Richmond, Energy Sciences Network

Help Desk**Help Desk Co-Chair**

Virginia Bedford, ERDC/ITL
Karmen Goddard, Self

Help Desk Team Member(s)

Timothy W. Dunaway, U.S. Army Engineer Research and Development Center

IT/Measurement**IT/Measurement Co-Chair**

Ezra Kissel, Indiana University
Jason Zurawski, Energy Sciences Network

IT/Measurement Team Member(s)

Sana Bellamine, CENIC
Aaron Brown, Internet2
Scott Chevalier, Indiana University
David Ediger, Georgia Institute of Technology
Greg Goddard, Radware
Trevor Goodyear, Georgia Tech Research Institute
Benjamin Grover, Lawrence Livermore National Laboratory
Neil H. McKee, InMon Corporation
Ken Miller, Pennsylvania State University
David Mitchell, Energy Sciences Network
Robert Stoy, DFN-Verein
Greg Veldman, Purdue University
Kevin Walsh, University of California, San Diego
Miao Zhang, Indiana University
Brett Zimmerman, University of Oklahoma

ITS**IT Services Member(s)**

Benjamin Grover, Lawrence Livermore National Laboratory
Mark Keele, Indiana University
Greg Veldman, Purdue University
Miao Zhang, Indiana University
Brett Zimmerman, University of Oklahoma

Measurement**Measurement Team Member(s)**

Sana Bellamine, CENIC
Aaron Brown, Internet2
Scott Chevalier, Indiana University
David Ediger, Georgia Institute of Technology
Greg Goddard, Radware
Trevor Goodyear, Georgia Tech Research Institute
Neil H. McKee, InMon Corporation
Ken Miller, Pennsylvania State University
David Mitchell, Energy Sciences Network
Robert Stoy, DFN-Verein
Kevin Walsh, University of California, San Diego

Interconnect**Interconnect Co-Chair**

Lance V. Hutchinson, Sandia National Laboratories

Interconnect Team Members

Scott Richmond, Energy Sciences Network

Logistics/Equipment**Logistics/Equipment Co-Chair**

Ken Brice, Army Research Laboratory
James H. Rogers, Oak Ridge National Laboratory

Logistics Team Members

Frank Indiviglio, National Oceanic and Atmospheric Administration

Network Research Exhibition**Network Research Exhibition****Co-Chair**

Cees de Laat, University of Amsterdam
Brian L. Tierney, Lawrence Berkeley National Laboratory
Matthew J. Zekauskas, Internet2

Network Research Exhibition**Team Members**

Nick Buraglio, Energy Sciences Network
Mary Hester, Lawrence Berkeley National Laboratory
Kevin Walsh, University of California, San Diego

Network Security**Network Security Co-Chair**

Jeff W. Boote, Sandia National Laboratories
Carrie Gates, Dell Research

Network Security Team Members

Ryan Birmingham, Sandia National Laboratories
Alan Commike, Reservoir Labs
Michael Dopheide, Energy Sciences Network
Clark Gaylord, Virginia Polytechnic Institute and State University
Thomas Kroeger, Sandia National Laboratories
Patrick Storm, The University of Texas at Austin
Vince Urias, Sandia National Laboratories
Kyeongseon West, Indiana University of Pennsylvania
Brian Wright, Sandia National Laboratories

Physical Security**Physical Security Co-Chair**

Jeff Graham, Air Force Research Laboratory

Power**Power Co-Chair**

Jim Pepin, Clemson University

Power Team Members

Jay Harris, Clemson University
Fred Keller, University of Oklahoma

Routing**Routing Co-Chair**

Conan Moore, University of Colorado Boulder
JP Velders, University of Amsterdam

Routing Team Members

Nick Buraglio, Energy Sciences Network
Evangelos Chaniotakis, Energy Sciences Network
Jamie Curtis, Research and Education Advanced Network New Zealand
Pieter de Boer, SURFnet
Ben Earl, University of Pittsburgh
Debbie Fligor, University of Illinois at Urbana-Champaign
Jackson Gor, Lawrence Berkeley National Laboratory
Thomas Hutton, University of California, San Diego
Indira Kassymkhanova, Lawrence Berkeley National Laboratory
Nathan Miller, Indiana University
Conan Moore, University of Colorado Boulder
Jared Schlemmer, Indiana University
Michael Smitasin, Lawrence Berkeley National Laboratory
JP Velders, University of Amsterdam
Alan Verlo, Freelance
Paul Wefel, University of Illinois at Urbana-Champaign
Routing Vendor
Marc Lonnais, Ciena Corporation

Student Volunteers**Student Volunteers Co-Chair**

Ralph A. McElDowney, Department of Defense HPC Modernization Program

WAN Transport**WAN Transport Co-Chair**

Akbar Kara, Lonestar Education and Research Network
Jim Stewart, Utah Education Network

WAN Transport Members

Eric Boomer, FLR
John Caffery, Louisiana State University
Tom Edmonson, Lonestar Education and Research Network
Kurt Freiburger, Lonestar Education and Research Network

Tan Geok Lian, A*STAR Computational Resource Centre
Byron Hicks, Lonestar Education and Research Network
Bill Jensen, University of Wisconsin-Madison
Patrick Keenan, Louisiana Optical Network Initiative
Lonnie Leger, Louisiana State University
E. Paul Love, Internet Consulting of Vermont
Gary Mumphy, Louisiana Optical Network Initiative
Kevin Nicholson, Lonestar Education and Research Network
Dave Pokorney, Florida LambdaRail
Kevin Quire, Utah Education Network
Chris Stowe, FLR
Chris Tracy, Energy Sciences Network
Wayne Wedemeyer, University of Texas
Tim Woodbridge, Lonestar Education and Research Network
Matthew J. Zekauskas, Internet2

WAN Transport Members/Vendor

Chad Dennis, Infinera Corporation
Fred Finlay, Infinera Corporation
Doug Hogg, Ciena Corporation
Chris Hunt, Alcatel-Lucent
Matt Jary, Infinera Corporation
Andis Kakeli, Infinera
Marc Lonnais, Ciena Corporation
Michael McAfee, CenturyLink
Carlos Pagan, Ciena Corporation
Rod Wilson, Ciena Corporation
Peng Zhu, Alcatel-Lucent

Wireless**Wireless Co-Chair**

Mark Mitchell, Sandia National Laboratories
Matt Smith, National Oceanic and Atmospheric Administration

Wireless Team Members

Matt Chrosniak, Cisco Systems, Inc.
Steven Phillips, Cisco Systems, Inc.
Megan Sorensen, Idaho State University
Benny Sparks, University of Tennessee, Knoxville

Student Programs**Student Programs Chair**

Jeanine Cook, Sandia National Laboratories
Student Programs Deputy Chair
Dorian C. Arnold, University of New Mexico

Arrangements**Student Programs Finance Chair**

Sandra Huskamp

Student Programs Communications Chair

Jon Bashor, Lawrence Berkeley National Laboratory

Student Programs Infrastructure Chair

Yuho Jin, New Mexico State University

HPC for Undergraduates**HPC for Undergraduates Chair**

Alan Sussman, University of Maryland

Mentor/Protégé**Mentor/Protégé Chair**

Christine Sweeney, Los Alamos National Laboratory

Student Activities**Student Activities Chair**

Dorian C. Arnold, University of New Mexico

Student Cluster Competition**Student Cluster Competition Chair**

Hai Ah Nam, Los Alamos National Laboratory

SCC Deputy Chair

Stephen Lien Harrell, Purdue University
SCC Voice of Reason, Sound Advisor & Webmeister
Sam Coleman, Retired

SCC Applications Lead

Scott Michael, Indiana University

SCC Technical Lead

Jason Kincl, Oak Ridge National Laboratory

SCC Social Media

Christopher Bross, Friedrich-Alexander-Universität Erlangen-Nürnberg

Student Job Fair**Student Job Fair Chair**

Richard Murphy, Micron Technology, Inc.

Student Job Fair Member(s)

Beth McKown, University of Illinois at Urbana-Champaign
Kyle B. Wheeler, University of Notre Dame and Sandia National Laboratories

Student Outreach**Student Outreach Chair**

Tony Baylis, Lawrence Livermore National Laboratory

Student Programs**Coordination****Student Programs Coordination Chair**

Verónica Vergara L., Oak Ridge National Laboratory

Student Volunteers**Student Volunteers Chair**

Yashema Mack, University of Tennessee, Knoxville

Student Volunteers Deputy Chair

Sally Ellingson, University of Kentucky

Student Volunteers Administration

Jason Grant, University of Notre Dame
Student Volunteers Event Planning and Social Media
Christine E. Harvey, MITRE Corporation

Student Volunteers Recruiting

Kathryn Traxler, Louisiana State University

Student Volunteers SCinet Liaison

Ralph A. McEldowney, Department of Defense HPC Modernization Program

Student Volunteer Application**Reviewer**

Sally Ellingson, University of Kentucky
Vikram Gazula, University of Kentucky
Jason Grant, University of Notre Dame
Christine E. Harvey, MITRE Corporation
Yuho Jin, New Mexico State University
Kate Mace, Clemson University
Yashema Mack, University of Tennessee, Knoxville
Ralph A. McEldowney, Department of Defense HPC Modernization Program
Damian Rouson, Sourcery, Inc.
Kathryn Traxler, Louisiana State University
Verónica Vergara L., Oak Ridge National Laboratory
Bobby Whitten, University of Tennessee

Technical Program**Technical Program Chair**

Jeffrey S. Vetter, Oak Ridge National Laboratory and Georgia Institute of Technology

Technical Program Deputy Chair

Lori Diachin, Lawrence Livermore National Laboratory

Assistant to the Technical Program Chair

Liz Hebert, Oak Ridge National Laboratory

Archive**Archive Chair**

Janet Brown, Carnegie Mellon University

Awards**Awards Co-Chair**

Franck Cappello, Argonne National Laboratory
Padma Raghavan, Pennsylvania State University

Birds-of-a-Feather**Birds-of-a-Feather Chair**

Karen L. Karavanic, Portland State University

Birds-of-a-Feather Vice Chair

Karl Fuerlinger, Ludwig Maximilian University of Munich

BOF Viewer

Manish Parashar, Rutgers University

Data Center Operations**BOF Data Center Operations Chair**

Rebecca Hartman-Baker, National Energy Research Scientific Computing Center and Lawrence Berkeley National Laboratory

BOF Data Center Operations Committee Member(s)

Fernanda Foertter, Oak Ridge National Laboratory
Henry J. Neeman, University of Oklahoma
Bobby Whitten, University of Tennessee

Diversity**BOF Diversity Chair**

Rashawn L. Knapp, Intel Corporation
BOF Diversity Committee Member(s)
Rashawn L. Knapp, Intel Corporation

Education**BOF Education Chair**

David P. Bunde, Knox College

BOF Education Committee Member(s)

Carsten Trinitis, Technical University of Munich
Large Scale Data Analysis
BOF Large Scale Data Analysis Chair
Yong Chen, Texas Tech University

BOF Large Scale Data Analysis Committee Member(s)

Fabrizio Petrini, IBM Corporation

Meetings

OS & Runtime Systems

BOF OS & Runtime Systems Chair

Ron Brightwell, Sandia National Laboratories

BOF OS & Runtime Systems Committee Member(s)

David P. Bunde, Knox College
Kurt B. Ferreira, Sandia National Laboratories
Richard Vuduc, Georgia Institute of Technology

BOF Other Committee Member(s)

David P. Bunde, Knox College
Karl Fuerlinger, Ludwig Maximilian University of Munich
Carsten Trinitis, Technical University of Munich Performance

BOF Performance Chair

Nicholas J. Wright, National Energy Research Scientific Computing Center

BOF Performance Committee Member(s)

Tapasya Patki, University of Arizona
Richard Vuduc, Georgia Institute of Technology
Nicholas J. Wright, National Energy Research Scientific Computing Center
Programming Languages

BOF Programming Languages Chair

Ali Jannesari, Technical University of Darmstadt

BOF Programming Languages Committee Member(s)

David Boehme, Lawrence Livermore National Laboratory
Marc-Andre Hermanns, RWTH Aachen University
Daniel Lorenz, Darmstadt University of Technology
Olga Pearce, Lawrence Livermore National Laboratory
Yukinori Sato, Tokyo Institute of Technology
Christian Terboven, RWTH Aachen University
Josef Weidendorfer, Technical University of Munich
Resilience

BOF Resilience Chair

Christian Engelmann, Oak Ridge National Laboratory

BOF Resilience Committee Member(s)

Wesley Bland, Intel Corporation
Franck Cappello, Argonne National Laboratory
Zizhong Chen, University of California, Riverside
Kurt B. Ferreira, Sandia National Laboratories
Qiang Guan, Los Alamos National Laboratory
Tanzima Z. Islam, Lawrence Livermore National Laboratory
Larry Kaplan, Cray Inc.

Sriram Krishnamoorthy, Pacific Northwest National Laboratory
Ignacio Laguna, Lawrence Livermore National Laboratory
Keita Teranishi, Sandia National Laboratories
Storage

BOF Storage Chair

Evan Felix, Pacific Northwest National Laboratory

BOF Storage Committee Member(s)

John Bent, EMC Corporation
David ML Brown, Pacific Northwest National Laboratory
Mark Gary, Lawrence Livermore National Laboratory
Gary Grider, Los Alamos National Laboratory
Quincey Koziol, HDF Group
Paul Nowoczynski, DataDirect Networks
Ron Oldfield, Sandia National Laboratories

Doctoral Showcase

Doctoral Showcase Chair

Melissa C. Smith, Clemson University

Doctoral Showcase Vice Chair

Volodymyr Kindratenko, University of Illinois at Urbana-Champaign

Doctoral Showcase Committee Members

Sunita Chandrasekaran, University of Delaware
Arun Chauhan, Indiana University and Google
Michael Gerndt, Technical University of Munich
Miaoqing Huang, University of Arkansas
Kamil Iskra, Argonne National Laboratory
Alice Koniges, Lawrence Berkeley National Laboratory
Oleksiy Koshulko, Institute of Cybernetics of VM Glushkov National Academy of Sciences of Ukraine
Elizabeth Leake, STEM-Trek
Joshua A. Levine, Clemson University
Sally A. McKee, Chalmers University of Technology
Catherine Olschanowsky, Colorado State University
Vivek K. Pallipuram, University of Delaware
Virginia W. Ross, Air Force Research Laboratory
Martin Schulz, Lawrence Livermore National Laboratory
Ziliang Zong, Texas State University
Early Career Activities
Early Career Activities Chair
Jeffrey K. Hollingsworth, University of Maryland

Early Career Activities Member(s)

Dorian C. Arnold, University of New Mexico
Jon Bashor, Lawrence Berkeley National Laboratory

Emerging Technologies

Emerging Technologies Chair

Sadaf R. Alam, Swiss National Supercomputing Center

Emerging Technologies Vice Chair

Simon McIntosh-Smith, University of Bristol

Emerging Technologies Member(s)

Sadaf R. Alam, Swiss National Supercomputing Center
James A. Ang, Sandia National Laboratories
Wu Feng, Virginia Polytechnic Institute and State University
Wayne Gaudin, AWE
Andy Herdman, AWE
Kurt Keville, Massachusetts Institute of Technology
Jaejin Lee, Seoul National University
Simon McIntosh-Smith, University of Bristol
Dhabaleswar K. (DK) Panda, Ohio State University
Felix Schürmann, Swiss Federal Institute of Technology in Lausanne
John Shalf, Lawrence Berkeley National Laboratory
Galen Shipman, Los Alamos National Laboratory
Neil Stringfellow, Pawsey Supercomputing Center

Invited Speakers

Invited Talks Chair

Bernd Mohr, Juelich Supercomputing Center

Invited Talks Committee Members

Jean-Yves Berthou, French National Research Agency
Taisuke Boku, University of Tsukuba
Mattan Erez, The University of Texas at Austin
David E. Keyes, King Abdullah University of Science and Technology
Irene Qualters, National Science Foundation
Mary Wheeler, University of Texas

Panels

Panels Chair

Torsten Hoefler, ETH Zurich

Panels Vice Chair

Hatem Ltaief, King Abdullah University of Science and Technology

Panels Committee Member(s)

David A. Bader, Georgia Institute of Technology
 Jed Brown, Argonne National Laboratory
 Kirk Cameron, Virginia Polytechnic Institute and State University
 William Harrod, Department of Energy Office of Advanced Scientific Computing Research
 Laxmikant V. Kale, University of Illinois at Urbana-Champaign
 Jakub Kurzak, University of Tennessee
 Michael M. Resch, High Performance Computing Center, Stuttgart
 Olaf Schenk, University of Italian Switzerland
 Rajeev Thakur, Argonne National Laboratory
 Mateo Valero, Barcelona Supercomputing Center

Posters

Posters Chair

Manish Parashar, Rutgers University

Posters Vice Chair

Dorian C. Arnold, University of New Mexico
 Michela Becchi, University of Missouri

Posters Committee Members

David Abramson, University of Queensland
 Ilkay Altintas, San Diego Supercomputer Center
 Amy Apon, Clemson University
 Dorian C. Arnold, University of New Mexico
 Susan R. Atlas, University of New Mexico
 Jason Bakos, University of South Carolina
 Michela Becchi, University of Missouri
 Janine C. Bennett, Sandia National Laboratories
 George Bosilca, University of Tennessee, Knoxville
 Ivona Brandic, Vienna University of Technology
 Patrick Bridges, University of New Mexico
 Vetrica Byrd, Clemson University
 Claris Castillo, RENCi
 John Cavazos, University of Delaware
 Dilma Da Silva, Texas A & M University
 Tony Drummond, Lawrence Berkeley National Laboratory
 Anshu Dubey, Argonne National Laboratory
 Erika Fuentes, University of Washington
 Bothell

Judit Gimenez, Barcelona Supercomputing Center
 Jennifer Green, Los Alamos National Laboratory
 Dean Hildebrand, IBM Corporation
 Scott Klasky, Oak Ridge National Laboratory
 Tamara Kolda, Sandia National Laboratories
 Alice Koniges, Lawrence Berkeley National Laboratory
 John Lange, University of Pittsburgh
 Miriam Leeser, Northeastern University
 Charles Lively, IBM Corporation
 Kathryn Mohror, Lawrence Livermore National Laboratory
 Christine Morin, French Institute for Research in Computer Science and Automation
 Manish Parashar, Rutgers University
 Olga Pearce, Lawrence Livermore National Laboratory
 Andres Quiroz, Xerox Corporation
 Ioan Raicu, Illinois Institute of Technology
 Raghunath Raja Chandrasekar, Cray Inc.
 Vignesh Ravi, A2ZLogix
 Ivan Rodero, Rutgers University
 Damian Rouson, Sourcery, Inc.
 Kitrick Sheets, Cray Inc.
 Michelle Mills Strout, Colorado State University
 Zehra Sura, IBM Corporation
 Christine Sweeney, Los Alamos National Laboratory
 Keita Teranishi, Sandia National Laboratories
 Matthew Turk, University of Illinois at Urbana-Champaign
 Sathish S. Vadiyar, Indian Institute of Science
 Carlos Varela, Rensselaer Polytechnic Institute
 Lizhe Wang, Chinese Academy of Sciences
 Patrick Widener, Sandia National Laboratories
 Timothy Wood, George Washington University
 Carole-Jean Wu, Arizona State University
 Esma Yildirim, Rutgers University
 Zheng Zhang, Rutgers University
 Jaroslaw Zola, University at Buffalo
 Proceedings, Tech

Tech Proceedings Chair

Cherri M. Pancake, Oregon State University

Scientific Visualization and Data

Analytics Showcase

Scientific Visualization and Data Analytics Showcase Chair

Jean M. Favre, Swiss National Supercomputing Center

Scientific Visualization and Data Analytics Showcase Vice Chair

Berk Geveci, Kitware, Inc.

Scientific Visualization and Data Analytics Showcase Committee Member(s)

Utkarsh Ayachit, Kitware, Inc.
 David Camp, Lawrence Berkeley National Laboratory
 Thierry Carrard, French Alternative Energies and Atomic Energy Commission
 Amit Chourasia, University of California, San Diego
 Kelly Gaither, The University of Texas at Austin
 Christoph Garth, University of Kaiserslautern
 Aaron Knoll, University of Utah
 Burlen Loring, Lawrence Berkeley National Laboratory
 Paul A. Navratil, The University of Texas at Austin
 Kenji Ono, RIKEN
 Tom Peterka, Argonne National Laboratory
 David Pugmire, Oak Ridge National Laboratory
 Paul Rosen, University of South Florida
 David Semeraro, University of Illinois
 Robert R. Sisneros, University of Illinois at Urbana-Champaign

Tech Program Liaisons

Liaison to SC Student Programs

Verónica Vergara L., Oak Ridge National Laboratory
 Liaison to ACM Gordon Bell Award
 Satoshi Matsuoka, Tokyo Institute of Technology

Technical Papers

Technical Papers Co-Chair

Ewa Deelman, University of Southern California
 California Information Sciences Institute
 José Moreira, IBM Corporation

Algorithms

Area Co-Chair, Algorithms

Umit Catalyurek, Ohio State University
 Karen Devine, Sandia National Laboratories

Member(s), Algorithms

Srinivas Aluru, Georgia Institute of Technology
 Cevdet Aykanat, Bilkent University
 Wolfgang Bangerth, Texas A & M University
 Ioana Banicescu, Mississippi State University
 Olivier Beaumont, French Institute for Research in Computer Science and Automation
 Costas Bekas, IBM Corporation
 Anne Benoit, ENS Lyon
 Sanjukta Bhowmick, University of Nebraska Omaha
 Laura Grigori, French Institute for Research in Computer Science and Automation
 John A. Gunnels, IBM Corporation
 Judith C. Hill, Oak Ridge National Laboratory

Kamer Kaya, Sabancı University
 Sarah Knepper, Intel Corporation
 X. Sherry Li, Lawrence Berkeley National Laboratory
 Kamesh Madduri, Pennsylvania State University
 Fredrik Manne, University of Bergen
 Kengo Nakajima, University of Tokyo
 Jacopo Pantaleoni, NVIDIA Corporation
 Cynthia Phillips, Sandia National Laboratories
 Ali Pinar, Sandia National Laboratories
 Sivasankaran Rajamanickam, Sandia National Laboratories
 Erik Saule, University of North Carolina at Charlotte
 Didem Unat, Koç University
 Frédéric Vivien, French Institute for Research in Computer Science and Automation
 Weichung Wang, National Taiwan University
 Carol Woodward, Lawrence Livermore National Laboratory

Applications

Area Co-Chair, Applications

Gabrielle Allen, University of Illinois at Urbana-Champaign
 Henry Tufo, University of Colorado

Member(s), Applications

H. Metin Aktulga, Michigan State University
 Mihai Anitescu, Argonne National Laboratory
 Takayuki Aoki, Tokyo Institute of Technology
 Alan Calder, Stony Brook University
 Daniela di Serafino, Second University of Naples
 Omar Ghattas, The University of Texas at Austin
 Marc Hamilton, NVIDIA Corporation
 Kirk E. Jordan, IBM Corporation
 Daniel S. Katz, University of Chicago and Argonne National Laboratory
 John Levesque, Cray Inc.
 Abani Patra, University at Buffalo
 Alex Pothen, Purdue University
 Ulrich J. Ruede, University of Erlangen-Nuremberg
 Karl W. Schulz, Intel Corporation
 Gilad Shainer, HPC Advisory Council
 Spencer Sherwin, Imperial College London
 Suzanne Shontz, University of Kansas
 William Tang, Princeton University
 Ray S. Tuminaro, Sandia National Laboratories
 Theresa Windus, Iowa State University
 Pelin Yilmaz, Max Planck Institute for Marine Microbiology
 Huai Zhang, University of Chinese Academy of Sciences

Architectures and Networks

Area Co-Chair, Architectures and Networks

Ilya Baldin, RENCi
 John Kim, Korea Advanced Institute of Science and Technology
 Member(s), Architectures and Networks
 Jung Ho Ahn, Seoul National University
 Ilya Baldin, RENCi
 Keren Bergman, Columbia University
 David Black-Schaffer, Uppsala University
 Sangyeun Cho, Samsung
 Reetu Das, University of Michigan
 Mattan Erez, The University of Texas at Austin
 Admela Jukan, Braunschweig University of Technology
 Jangwoo Kim, Pohang University of Science and Technology
 Eren Kursun, JPM
 Yongkee Kwon, SK hynix Inc.
 Xiaolin Andy Li, University of Florida
 Timothy Miller, Binghamton University
 Inder Monga, Lawrence Berkeley National Laboratory
 José Moreira, IBM Corporation
 Mike O'Connor, NVIDIA Corporation
 Seung-Jong Park, Louisiana State University
 Steve Scott, Google
 Brian Towles, D. E. Shaw Research
 Keith Underwood, Intel Corporation
 Eric Van Hensbergen, ARM Research Labs, LLC
 Malathi Veeraraghavan, University of Virginia

Clouds and Distributed Computing

Area Co-Chair, Clouds and Distributed Computing

Rosa M. Badia, Barcelona Supercomputing Center
 Satoshi Sekiguchi, National Institute of Advanced Industrial Science and Technology

Member(s), Clouds and Distributed Computing

David Abramson, University of Queensland
 Kento Aida, National Institute of Informatics
 Ignacio Blanquer, Polytechnic University of Valencia and Center for Energy, Environment and Technology
 Claris Castillo, RENCi
 Vijay Gadepally, Massachusetts Institute of Technology
 Antonella Galizia, Institute of Applied Mathematics and Information Technologies
 Satomi Hasegawa, Hitachi Data Systems Corporation
 Weicheng Huang, National Center for High-Performance Computing, Taiwan
 Shantenu Jha, Rutgers University
 Bastian Koller, High Performance Computing Center, Stuttgart
 Craig Lee, Aerospace Corporation

Jysoo Lee, Korea Institute of Science and Technology Information
 Marta Mattoso, Federal University of Rio de Janeiro
 Jelena Pjesivac-Grbovic, Google
 Lavanya Ramakrishnan, Lawrence Berkeley National Laboratory
 Tin Wee Tan, National Supercomputing Center Singapore
 Fang Zheng, IBM Corporation

Data Analytics, Visualization and Storage

Area Co-Chair, Data Analytics, Visualization and Storage

Manoj Kumar, IBM Corporation
 Zheng Zhang, Rutgers University

Member(s), Data Analytics, Visualization and Storage

Daniel Acevedo-Feliz, King Abdullah University of Science and Technology
 Bill Bolosky, Microsoft Corporation
 Daniel G. Chavarría-Miranda, Pacific Northwest National Laboratory
 Yong Chen, Texas Tech University
 Hank Childs, University of Oregon and Lawrence Berkeley National Laboratory
 Toni Cortes, Barcelona Supercomputing Center
 David Daly, MongoDB, Inc.
 David Ebert, Purdue University
 Thomas Ertl, University of Stuttgart
 Kelly Gaither, The University of Texas at Austin
 Mahmut Kandemir, Pennsylvania State University
 Dries Kimpe, KCG Holdings, Inc.
 Manoj Kumar, IBM Corporation
 Qing Gary Liu, Oak Ridge National Laboratory
 Mausam Mausam, Indian Institute of Technology Delhi and University of Washington
 Suzanne McIntosh, Cloudera, Inc. and New York University
 Priya Nagpurkar, IBM Corporation
 Alma Riska, NetApp
 Steven Smith, Massachusetts Institute of Technology
 Madhu Srinivasan, King Abdullah University of Science and Technology
 Michela Taufer, University of Delaware
 Brent Welch, Google
 Mohammad Zubair, Old Dominion University

Performance

Area Co-Chair, Performance

Kirk Cameron, Virginia Polytechnic Institute and State University
 Matthias S. Mueller, RWTH Aachen University

Member(s), Performance

Christos Antonopoulos, University of Thessaly, Greece

Filip Blagojevic, HGST, Inc.
 Ali R. Butt, Virginia Polytechnic Institute and State University
 Suren Byna, Lawrence Berkeley National Laboratory
 John Cavazos, University of Delaware
 Sunita Chandrasekaran, University of Delaware
 Dilma Da Silva, Texas A & M University
 Bronis R. de Supinski, Lawrence Livermore National Laboratory
 Chen Ding, University of Rochester
 Alejandro Duran, Intel Corporation
 Rong Ge, Clemson University
 Markus Geimer, Juelich Supercomputing Center
 Lizy K. John, The University of Texas at Austin
 Karen L. Karavanic, Portland State University
 Darren Kerbyson, Pacific Northwest National Laboratory
 Kalyan Kumaran, Argonne National Laboratory
 Alexey Lastovetsky, University College Dublin
 Dong Li, University of California, Merced
 Min Li, IBM Corporation
 David Lowenthal, University of Arizona
 Xiaosong Ma, Qatar Computing Research Institute
 Shirley Moore, University of Texas at El Paso
 Matthias S. Mueller, RWTH Aachen University
 Wolfgang E. Nagel, Technical University of Dresden
 Dimitrios Nikolopoulos, Queen's University Belfast
 Fabrizio Petrini, IBM Corporation
 Amanda Randles, Duke University
 Xipeng Shen, North Carolina State University
 Shuaiwen Leon Song, Pacific Northwest National Laboratory
 Xian-He Sun, Illinois Institute of Technology
 Gabriel Tanase, IBM Corporation
 Jun Wang, University of Central Florida
 Zhiwei Xu, Institute of Computing Technology, Chinese Academy of Sciences

Programming Systems

Area Co-Chair, Programming Systems

Barbara Chapman, University of Houston
 Thomas Fahringer, University of Innsbruck

Member(s), Programming Systems

George Almasi, IBM Corporation
 Pavan Balaji, Argonne National Laboratory
 Siegfried Benkner, University of Vienna
 Joao Cardoso, University of Porto
 Laura C. Carrington, University of California, San Diego and EP Analytics
 Brad Chamberlain, Cray Inc.
 Evelyn Duesterwald, IBM Corporation
 Maria Garzan, University of Illinois at Urbana-Champaign
 Ganesh Gopalakrishnan, University of Utah

Sebastian Hack, Saarland University
 Herbert Jordan, University of Innsbruck
 Erwin Laure, KTH Royal Institute of Technology
 I-Ting A. Lee, Washington University in St. Louis

Chunhua Liao, Lawrence Livermore National Laboratory
 Tomas Margalef, Autonomous University of Barcelona
 Naoya Maruyama, RIKEN
 Angeles Navarro, Universidad of Malaga, Spain
 Catherine Olschanowsky, Colorado State University
 Swaroop Pophale, Mellanox Technologies
 Ioan Raicu, Illinois Institute of Technology
 Toyotaro Suzumura, IBM Corporation
 Christian Terboven, RWTH Aachen University
 Jesper Larsson Träff, Vienna University of Technology
 Ana Lucia Varbanescu, University of Amsterdam
 Yonghong Yan, Oakland University
 Antonia Zhai, University of Minnesota

State of the Practice

Area Co-Chair, State of the Practice

Neil Chue Hong, University of Edinburgh
 Robin J. Goldstone, Lawrence Livermore National Laboratory

Member(s), State of the Practice

Sadaf R. Alam, Swiss National Supercomputing Center
 Jaydeep Bardhan, Northeastern University
 Neil Chue Hong, University of Edinburgh
 Giri Chukkapalli, Broadcom Corporation
 Michael A. Clark, NVIDIA Corporation
 Susan Coghlan, Argonne National Laboratory
 Kim Cupps, Lawrence Livermore National Laboratory
 Wu Feng, Virginia Polytechnic Institute and State University
 Robin J. Goldstone, Lawrence Livermore National Laboratory
 Alison Kennedy, Edinburgh Parallel Computing Centre
 Jennifer M. Schopf, Indiana University
 Devesh Tiwari, Oak Ridge National Laboratory

System Software

Area Co-Chair, System Software

Frank Mueller, North Carolina State University
 Peng Wu, Huawei Technologies Co., Ltd.

Member(s), System Software

Gabriel Antoniu, French Institute for Research in Computer Science and Automation
 Greg Bronevetsky, Lawrence Livermore National Laboratory
 Wenguang Chen, Tsinghua University

Zizhong Chen, University of California, Riverside
 Andrew Chien, University of Chicago
 Kaoutar El Maghraoui, IBM Corporation
 Christian Engelmann, Oak Ridge National Laboratory
 Kurt B. Ferreira, Sandia National Laboratories
 Ada Gavrilovska, Georgia Institute of Technology
 Ziang Hu, Huawei Technologies Co., Ltd.
 Larry Kaplan, Cray Inc.
 Satoshi Matsuoka, Tokyo Institute of Technology
 Mori Ohara, IBM Corporation
 Rafael Asenjo Plaza, University of Málaga
 Robert B. Ross, Argonne National Laboratory
 Michael Spear, Lehigh University
 Yongpeng Zhang, Stone Ridge Technology

Test of Time Award

Test of Time Award Chair

Jack Dongarra, University of Tennessee, Knoxville

Test of Time Award Vice Chair

Mary Hall, University of Utah

Test of Time Award Committee Member(s)

David Abramson, University of Queensland
 William Douglas Gropp, University of Illinois at Urbana-Champaign
 Michael A. Heroux, Sandia National Laboratories
 Elizabeth Jessup, University of Colorado Boulder
 Ewing Lusk, Argonne National Laboratory
 Leonid Oliker, Lawrence Berkeley National Laboratory
 Padma Raghavan, Pennsylvania State University
 Yves Robert, ENS Lyon
 Valerie Taylor, Texas A & M University
 Mateo Valero, Barcelona Supercomputing Center

Tutorials

Tutorials Chair

Richard Vuduc, Georgia Institute of Technology

Tutorials Vice Chair

Chris J. Newburn, Intel Corporation

Tutorials Committee Member(s)

Wonsun Ahn, University of Pittsburgh
 Francois Bodin, University of Rennes 1
 Sung-Eun Choi, Cray Inc.
 Guillaume Colin de Verdiere, French Alternative Energies and Atomic Energy Commission
 Guojing Cong, IBM Corporation
 Luiz DeRose, Cray Inc.

Joe Dummy, TryIt Inc.
 Alejandro Duran, Intel Corporation
 Fernanda Foertter, Oak Ridge National Laboratory
 Jeff R. Hammond, Intel Corporation
 Andy Herdman, AWE
 Elizabeth Jessup, University of Colorado Boulder
 Takahiro Katagiri, University of Tokyo
 Ivan Kuzmin, Intel Corporation
 Jesus Labarta, Barcelona Supercomputing Center
 David Lecomber, Allinea Software
 Jonathan Lifflander, University of Illinois at Urbana-Champaign
 Andrew Lumsdaine, Indiana University
 Peter Messmer, ETH Zurich
 Manoj Nambiar, Tata Consultancy Services
 Stephen L. Olivier, Sandia National Laboratories
 Govindarajan Ramaswamy, Indian Institute of Science
 Carlos Rosales, The University of Texas at Austin
 Damian Rouson, Sourcery, Inc.
 Karl W. Schulz, Intel Corporation
 John Shalf, Lawrence Berkeley National Laboratory
 Happy Sithole, Centre for High Performance Computing
 Lauren L. Smith, National Security Agency
 Mahidhar Tatineni, San Diego Supercomputer Center
 Didem Unat, Koç University
 Sandra Wienke, RWTH Aachen University
 Rio Yokota, Tokyo Institute of Technology

Workshops

Workshops Conflict Chair

Seetharami Seelam, IBM Corporation

Workshops Chair

Michela Taufer, University of Delaware

Workshops Vice-Chair

Trilce Estrada, University of New Mexico

Workshop Committee Member(s)

Dong H. Ahn, Lawrence Livermore National Laboratory
 Pavan Balaji, Argonne National Laboratory
 Umit Catalyurek, Ohio State University
 Brad Chamberlain, Cray Inc.
 Almadena Chtchelkanova, National Science Foundation
 Pietro Cicotti, San Diego Supercomputer Center
 Frederic Desprez, French Institute for Research in Computer Science and Automation
 Alexandru Iosup, Delft University of Technology
 Bahman Javedi, University of Western Sydney
 Daniel S. Katz, University of Chicago and Argonne National Laboratory
 Naoya Maruyama, RIKEN
 Tim Mattson, Intel Corporation
 Scott Pakin, Los Alamos National Laboratory
 Song J. Park, Army Research Laboratory
 María S. Pérez, UPM
 Seetharami Seelam, IBM Corporation
 Francesco Silvestri, IT University of Copenhagen
 Alan Sussman, University of Maryland
 Kenjiro Taura, University of Tokyo
 Patricia J. Teller, University of Texas at El Paso
 Douglas Thain, University of Notre Dame
 Antonino Tumeo, Pacific Northwest National Laboratory

Society Representation

Society Participants

Donna Cappel, Association for Computing Machinery
 Ashley Cozzi, Association for Computing Machinery
 Anne Marie Kelly, IEEE Computer Society
 Carmen Saliba, CMP, IEEE Computer Society

The International Conference for High Performance Computing, Networking, Storage and Analysis



SC16

Salt Lake City, Utah | hpc matters.

Call for Participation

<http://sc16.supercomputing.org/>

Conference Dates: Nov. 13-18, 2016 • Exhibition Dates: Nov. 14-17, 2016 • Salt Lake City, UT

SC16: HPC Matters

Underdeveloped areas of the world have new access to clean water and sanitation...Coastal cities are evacuated before hurricanes claim lives...A new drug offers hope for patients with the early stages of Alzheimer's...A new helmet keeps a child safe while playing football...HPC matters because HPC is powering revolutionary advances in quality of life – for everyone.

In 2016 SC returns to beautiful Salt Lake City, Utah – a vibrant city with a breathtaking alpine backyard that combines metro and mountain. This metropolitan area of over one million people is adjacent to the beautiful Wasatch Mountains, and easy to get to from anywhere in the world. The Salt Lake City International Airport, rated #1 for on-time arrivals and departures year after year, is just an 8-minute ride to downtown Salt Lake and its convention district. The convention district's fare-free transit zone, with easy access to TRAX light-rail and buses, makes getting around Salt Lake City simple and affordable.

There is no better place to see why HPC matters to our lives, our future, our communities, and our world.



Sponsors:



Association for
Computing Machinery

**Salt Palace
Convention Center**
Salt Lake City, Utah