Conference Program

HPC Everywhere, Everyday

SC13
Denver, CO | 2013

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

Conference Dates: November 17-22, 2013
Exhibition Dates: November 18-21, 2013

Denver Convention Center
Denver, CO

Sponsors: IEEE computer society, acm Association for Computing Machinery

SC13
Denver, CO | 2013

Table of Contents

Welcome to SC13

Bill Gropp, SC13 General Chair

Long-time HPC community members know that SC is the hub of our community. While there are many, many excellent technical conferences in HPC, SC, sponsored each year by the Association for Computing Machinery and the IEEE Computer Society, is the largest and most diverse conference in our community. Recent conferences have brought together more than 10,000 attendees during conference week from all over the world, truly making it the international conference for high performance computing, networking, storage, and analysis.

Each year SC pulls together the most complete cross-section of the technical work currently being pursued in supercomputing — this is reflected in the quality of the technical papers, the depth of the educational and engagement opportunities, and the breadth of technologies and research showcased on the exhibit floor. A quick look at the topics that will be featured this year shows that our tag line, "HPC Everywhere, Everyday" reflects the impact that HPC, and the people that attend SC, have on all aspects of society around the world.

Inside the convention center things will be bustling. We're expecting more than 350 exhibitors spread over nearly 140,000 square feet of exhibit space. With 26 meeting rooms dedicated to the Technical Program, Exhibitor Forum, Tutorials, Workshops, and HPC Interconnections, and two ballrooms for Technical Program events, there'll be plenty of space for all of our events, yet laid out in a way that makes it quick and easy to find your way around.

SC's continuing goal is to provide an informative, high-quality technical program that meets the highest academic standards. The Technical Program is highly competitive and one of the broadest of any HPC conference, with venues ranging from invited talks, panels, and research papers to tutorials, workshops, posters, and Birds-of-a-Feather (BOF) sessions. And this year we are honoring SC's legacy of success by extending peer review to cover every aspect of the technical program. I am also excited by the Emerging Technologies exhibits we have added to the program this year: these are peer-reviewed research projects that will be demonstrated on the show floor. Whether you are unveiling new research for the first time, or helping teach the HPC body of knowledge to the next generation, being part of the SC technical program is a rite of passage in any HPC career.

SC is fundamentally a technical conference, and anyone who has spent time on the show floor knows that the SC Exhibits program provides a unique opportunity to interact with the future of HPC. Far from being just a simple industry exhibition, our research and industry booths showcase recent developments in our field, with a rich combination of research labs, universities, and other organizations and vendors of all types of software, hardware, and services for HPC.

HPC Interconnections, or "HPCI" for short, is our new name for the SC Communities program. The new name reflects our renewed emphasis on providing an "on ramp" into the excitement that is SC and the field of high performance computing. HPCI is designed to help everyone get more out of the conference than ever, providing programs for everyone interested in building a stronger HPC community, including students, educators, researchers, international attendees and under-represented groups.

One of the defining characteristics of SC is SCinet, the conference network that, for the week of the conference, is one of the largest and most advanced networks in the world. While many attendees experience SCinet directly as wireless networking that is provided throughout the convention center, SCinet also provides a once-a-year opportunity for research and engineering groups to work with some of the most modern networking equipment and very high bandwidth wide area networking. This year, SCinet will exceed 1Terabit/second of networking bandwidth! During the conference please be sure to stop by the SCinet booth for more information on the networking infrastructure, and check out the SCinet research exposition presentations on Thursday to witness firsthand the future of high performance networks.

This year the conference will mark its 25th anniversary. in recognition of our silver anniversary, and of the lasting contribution that our participants have made to the HPC state of the art, we are hosting several new activities that look back on the progress in our field as seen through the lens of SC. There are exhibits of technology—many of which were first described or seen at SC—on the exhibit floor. A panel will talk about what we've learned over the years, and how that can help us as we confront the coming end of Moore's law and the changes in software, algorithms, and hardware that will be required in the coming decade. Also new this year is a new award for the most influential paper presented at SC that has stood the test of time in our "Test of Time" award.

Welcome to Denver!

For a quarter of a century, the Supercomputing Conference has served as the crossroads for the entire HPC community. Denver is a community at the crossroads of the country, located near the center of the continental United States, at the convergence of mountain and prairie, where high-rise buildings greet open range and the earth touches the sky. I can't think of a better place to celebrate SC's rich past, and our community's bright future. Welcome to Denver and to SC13!

Acknowledgements

No conference this size could be possible without the dedication, commitment and passion of the SC13 committee. The core committee includes more than 100 people who have largely volunteered their time planning for this event for more than three years. Added to that number are the more than 500 people who have helped review submissions and contributed to the planning and preparations. The full list of committee members is posted on the conference website at sc13.supercomputing.org.

SC13 Committee Management

Conference Chair
William Douglas Gropp, University of Illinois at Urbana-Champaign

Deputy General Chair
Trish Damkroger, Lawrence Livermore National Laboratory

Executive Committee Vice Chair
Scott Lathrop, Shodor/National Center for Supercomputing Applications

Executive Assistant
Beth McKown, National Center for Supercomputing Applications
Carolyn Peters, Argonne National Laboratory

Future Chair
Jackie Kern, University of Illinois at Urbana-Champaign

Society Reps - ACM
Donna Cappo, ACM
Ashley Cozzi, ACM

Society Reps - IEEE/CS
Brookes Little, IEEE Computer Society
Carmen Saliba, CMP, IEEE Computer Society

Communications Chair
John West, DOD HPC Modernization Program

Exhibits Chair
Jim Costa, Sandia National Laboratories

Finance Chair
Becky Verastegui, Oak Ridge National Laboratory

HPC Interconnections Chair
Elizabeth Jessup, University of Colorado Boulder

Infrastructure Chairs
Christine E. Cuicchi, HPC Modernization Program
Tim Yeager, US Air Force Research Laboratory

SCinet Chair
Trey Breckenridge, Mississippi State University

Silver Anniversary Chair
Mary Hall, University of Utah

Technical Program Chair
Satoshi Matsuoka, Tokyo Institute of Technology

STATE OF COLORADO

EXECUTIVE CHAMBERS

136 State Capitol
Denver, CO 80203 - 1792
Phone (303) 866-2471

John Hickenlooper
Governor

November 17, 2013

Greetings:

On behalf of the State of Colorado, it is my pleasure to welcome you to SC13, the International Conference for High Performance Computing (HPC), Networking, Storage, and Analysis. This year we are proud to be hosting this prestigious event at the Colorado Convention Center in downtown Denver, Colorado.

Attracting thousands of research and technology leaders from around the world, this conference represents a remarkable opportunity to meet new collaborators and experience the historic innovation that is changing the landscape of this industry.

We invite you to enjoy the beauty that is Colorado – from our majestic mountains and wide-open plains to our bustling cities and invigorating climate. In addition to our natural beauty, Colorado has also emerged as a hotspot for start-ups, entrepreneurship, business and innovation, and I hope you experience the energy and collaborative spirit of our welcoming state.

Best wishes for a successful conference, and we look forward to your return visit.

Sincerely,

John W. Hickenlooper
Governor

What's New at SC13

In addition to the programs you've come to expect at SC, you'll find these new opportunities to learn more about all things HPC:

- Emerging technologies, a peer-reviewed addition to the technical program that showcases new developments; you'll find this on the exhibit floor.

- HPC Impact Showcase highlights the use of HPC, from consumer goods to industrial applications. Also on the exhibit floor—come see how HPC is changing the world.

- Test of Time Award is for the paper, presented at SC, that has had the greatest sustained impact on our field. Presented on Thursday, along with the first ever ACM Athena Lecture at SC.

- HPC Interconnections is our new name for the SC Communities program, to emphasize the focus on providing an "on ramp" into the excitement of our field. New this year is a closer integration with the technical program. Have a look at the offerings of the HPC Educator and Broader Engagement programs; the sessions beginning Tuesday are open to all Technical Program attendees.

- In celebration of the 25th anniversary of SC we have special exhibits, including some at participating exhibitors, and a panel discussing the past and future of our field.

General Information

In this section, you'll find information on registration, exhibit hours, conference store hours, descriptions and locations of all conference social events, information booths and their locations, as well as convention center facilities and services.

General Information

General Information

Registration and Conference Store

The registration area and conference store are located on the 1st floor of the Denver Convention Center Lobby area.

Hours:

Saturday, November 16	**1pm – 6pm**
Sunday, November 17	**7am – 6pm**
Monday, November 18	**7am – 9pm**
Tuesday, November 19	**7:30am – 6pm**
Wednesday, November 20	**7:30am – 6pm**
Thursday, November 21	**7:30am – 5pm**
Friday, November 22	**8am – 11am**

Registration Pass Access

See page 11 for access grids.

Exhibit Hall Hours

Tuesday, November 19	**10am-6pm**
Wednesday, November 20	**10am-6pm**
Thursday, November 21	**10am-3pm**

Age Requirements Policy

- Technical Program attendees must be 16 years of age or older. Age verification is required.
- Exhibits-Only registration is available for children ages 12-16. Age verification is required.
- Children 12 and under are not permitted in the Exhibit Hall other than on Family Day (see page 13 details).
- Children under 16 are not allowed in the Exhibit Hall during installation, dismantling or before or after posted exhibit hours. Anyone under 16 must be accompanied by an adult at all times while visiting the exhibition.

SC13 Information Booth

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: one is in F Lobby, near the Blue Bear entrance, next to the conference store; the second booth is located in the open area near the Mile High ballroom and 600 series rooms.

Information Booth hours are:

Saturday, November 16	**1pm-6pm**
Sunday, November 17	**7am-6pm**
Monday, November 18	**7am-9pm**
Tuesday, November 19	**7:30am-6pm**
Wednesday, November 20	**7:30am-6pm**
Thursday, November 21	**7:30am-5pm**
Friday, November 22	**8am-11pm** (the booth near the Mile High ballroom is closed)

SC14 Preview Booth

Members of next year's SC committee will be available in the SC14 preview booth (located in A Lobby - across from Registration) to offer information and discuss next year's SC conference in New Orleans. 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, November 19	**10am-4pm**
Wednesday, November 20	**10am-4pm**
Thursday, November 21	**10am-4pm**

SC 25th Anniversary

SC13 plans several activities in celebration of the Silver Anniversary of the conference. The signature event will be a panel on Thursday afternoon titled "Retrospective on Supercomputing Technologies." In addition, there will be special exhibits, historical displays and recognition of attendees and exhibitors who have participated in all 25 SC conferences.

Social Events

Exhibitors' Party (Coors Field)

Sunday, November 17
6pm-9pm

Take me out to the ball park! SC13 is hosting an Exhibitors' Party for registered exhibitors. The party is SC13's way of thanking exhibitors for their participation and support of the conference. The party will be held at Coors Field, home of The Colorado Rockies professional baseball team. Everyone is encouraged to wear the jersey of their favorite sports team. The party will feature a live (and lively!) band, self-guided tours of the facility, including the visitor's locker room and dug out, games, and ballpark food and drinks throughout the facility.

Transportation is provided from B Lobby of the convention center to Coors Field starting at 5:45pm.

An Exhibitor badge, party ticket, and government-issued photo ID are required to attend this event.

Exhibits Gala Opening Reception (Exhibit Hall)

Monday, November 18
7pm-9pm

SC13 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 HPC Interconnections Program (Broader Engagement, HPC Educators and Student Volunteers) registrants.

Posters Reception (Mile High Pre-Function Area)

Tuesday, November 19
5:15pm-7pm

The reception is an opportunity for attendees to interact with poster presenters and includes research, electronic, and the ACM Student Research Competition (SRC) posters. The reception is open to all attendees with Technical Program registration. Complimentary refreshments and appetizers are available until 7pm.

Scientific Visualization Showcase Reception

(Mile High Pre-Function Area)
Tuesday, November 19
5:15pm-7pm

After you have viewed the posters at the Posters Reception, stop by the Scientific Visualization Showcase Reception. The reception is open to all attendees with Technical Program registration. The reception is located in the Mile High Ballroom Pre-function area.

Technical Program Conference Reception

(Denver Museum of Nature & Science)
Thursday, November 21
7pm-10pm
Buses Leave from B Lobby

SC13 is hosting a conference reception for all Technical Program attendees. Join us for delicious food, beverages, and entertainment at Denver Museum of Nature and Science (www.dmns.org). The Denver Museum of Nature & Science is the Rocky Mountain region's leading resource for informal science education. A variety of exhibitions and activities will be available for attendees to experience the natural wonders of Colorado, Earth, and the Universe.

A Technical 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 6:30pm-10pm from B Lobby of the convention center (look for buses with "DMNS" signs in the front window).

ACM SIGHPC Annual Members Meeting

See page 114.

Conference Services/Convention Center Facilities

SC13 Mobile App

Go Mobile with Boopsie—the SC13 App
Access the SC13 schedule, session and paper abstracts, information on speakers, exhibit maps, local information and more with your smart phone! Just go to your phone's app store and download the free SC13 app from Boopsie. You can access the SC13 app in one of two ways:

1. Direct from Boopsie (sc13.boopsie.com) from your smart phone.
2. Go to the App Store on your smart phone and search for 'SC13'.

The SC13 Boopsie app provides a searchable conference schedule, including session type, speaker(s), location, time and duration. It includes a complete roster of speakers, cross-correlated with the sessions in which they will appear. If you're looking for "What's Next" at SC13, there's even a handy feature to alert you to sessions starting in the next 60 minutes.Need to find a vendor's booth? There's a complete Exhibitor listing with booth numbers and a map of the exhibit floor. Included in the app are conference-related social media information including a direct link to the conference Facebook page and Twitter feeds, so you can see what is hot at the conference without needing to leave the app to check your phone's Twitter app.

ATMs

Three ATMs are located within the convention center:
F Lobby, B Lobby and D Lobby.

Attendee Lounge

Need to take a minute to relax before your next session? The attendee lounge is the place to do that. It is located in B Lobby, Sunday, November 17 through Friday, November 22 from 8:30am – 5pm.

Business Center

The Business Center is located off A Lobby in the red carpeted concourse of the convention center. The center is open 7:30am-5:30pm during convention days. For more information please contact: bborgerding@myofficeco.com or call 720.904.2300.

City and Dining Information

The information kiosk is located in the F Atrium, near the Blue Bear. It is staffed by Visit Denver.

Coat & Bag Check

The coat and bag check station is located in A Lobby. The hours are as follows:

Saturday, November 16	**7am-5:30pm**
Sunday, November 17	**7:00am-6pm**
Monday, November 18	**7:30am -9:30pm**
Tuesday, November 19	**7:30am-7:30pm**
Wednesday, November 20	**7:30am-7:30pm**
Thursday, November 21	**7:30pm-7:30pm**
Friday, November 22	**7:30pm-3:30pm**

Emergency Contact

For an onsite emergency, please call 303.228.8030 or dial 200 on a beige house phone.

Family Day

Family Day is Wednesday, November 20, 4pm-6pm. Adults and children 12 and over are permitted on the floor during these hours when accompanied by a registered conference attendee.

First-Aid Center

The first-aid center is located in E Lobby near room 507.

Lost Badge

There is a $40 processing fee to replace a lost badge.
Lost and found is located at the Coat Check station, A Lobby.

Parking

The Colorado Convention Center offers 1,000 on-site parking spaces with direct access to the center and the Denver Performing Arts Complex. It is open 24/7 to all visitors attending any event or business in the downtown area. For more information, see the website at: denverconvention.com/attend-an-event/parking.

Prayer & Meditation Room

The prayer and meditation room is located in Room 602 of the convention center. It is open 9am-6pm, Sunday, November 17 through Friday, November 22.

Restrooms

Restrooms are located conveniently throughout the convention center. See page 15 for locations.

Wheelchair/Scooter Rental

Wheelchairs and medical mobility scooters are available for rent at the Business Center. Wheelchairs are $4/hour or $15/day; scooters are $10/hour or $40/day.

Registration Pass Access

Registration Pass Access - Technical Program

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

Type of Event	Tutorials	Technical Program	Technical Program + Workshops	Workshop Only
Awards (Thursday)		*	*	*
Birds-of-a-Feather		*	*	
Broader Engagement & Educator Sessions (Sun/Mon)			*	*
Broader Engagement & Educator Sessions (Tue-Thu)		*	*	
Conference Reception (Thursday)		*	*	
Exhibit Floor		*	*	
Exhibitor Forum		*	*	
Exhibits Gala Opening (Monday)		*	*	
Invited Talks (Non-Plenary)		*	*	
Invited Talks (Plenary)		*	*	
Keynote (Tuesday)		*	*	
Panels (Tue-Thur)		*	*	
Panels (Friday Only)		*	*	
Papers		*	*	
Posters		*	*	
Poster Reception (Tuesday)		*	*	
Tutorial Lunch (Sun/Mon ONLY)	*			
Tutorial Sessions	*			
Student Cluster Competition		*	*	
Workshops			*	*

Registration Pass Access

Registration Pass Access - Exhibits

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

Type of Event	Exhibitor	Exhibit Hall Only
Awards (Thursday)	*	*
Birds-of-a-Feather	*	
Exhibit Floor	*	*
Exhibitor Forum	*	
Exhibits Gala Opening (Monday)	*	
Exhibitor's Reception	*	
Invited Talks (Plenary)	*	
Keynote (Tuesday)	*	
Panels (Friday Only)	*	
Posters	*	*
Student Cluster Competition	*	

Registration Pass Access - Press

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

Type of Event	Press
Awards (Thursday)	*
Birds-of-a-Feather	*
Broader Engagement & Educator Sessions (Sun/Mon)	*
Broader Engagement & Educator Sessions (Tue-Thu)	*
Conference Reception (Thursday)	*
Exhibit Floor (during show hours)	*
Exhibitor Forum	*
Exhibits Gala Opening (Monday)	*
Invited Talks (Non-Plenary)	*
Invited Talks (Plenary)	*
Keynote (Tuesday)	*

Maps/Daily Schedules

A schedule of each day's activities—by time/event/location—is provided in this section, along with a map of the Downtown area and meeting rooms.

Maps/Daily Schedules

VISIT DENVER
The Convention & Visitors Bureau

SC 2013 Hotel Map
Denver, CO
November 17-22, 2013

Hotel	Distance to CCC
1 The Brown Palace Hotel & Spa	5 Blocks
2 Comfort Inn Downtown	5 Blocks
3 Courtyard by Marriott Denver Downtown	3 Blocks
4 Crowne Plaza	2 Blocks
5 Curtis, a Doubletree Hotel	2 Blocks
6 Denver Marriott City Center	4 Blocks
7 Embassy Suites Denver Downtown @ CCC	50 Steps
8 Grand Hyatt Denver	3 Blocks
9 Hampton Inn & Suites Denver @ CCC	1 Block
10 Hampton Inn & Suites Downtown	0.9 Miles
11 Hilton Garden Inn Downtown Denver	1 Block
12 Homewood Suites Denver Downtown	1 Block
13 Hotel Monaco Denver	4 Blocks
14 Hotel Teatro	3 Blocks
15 Hyatt Regency Denver @ CCC	50 Steps
16 Magnolia Hotel	3 Blocks
17 Residence Inn Downtown Denver	1.5 Miles
18 Residence Inn by Marriott Denver City Center	5 Blocks
19 Ritz-Carlton Denver	7 Blocks
20 Sheraton Denver Downtown Hotel	3 Blocks
21 Westin Denver Downtown	5 Blocks

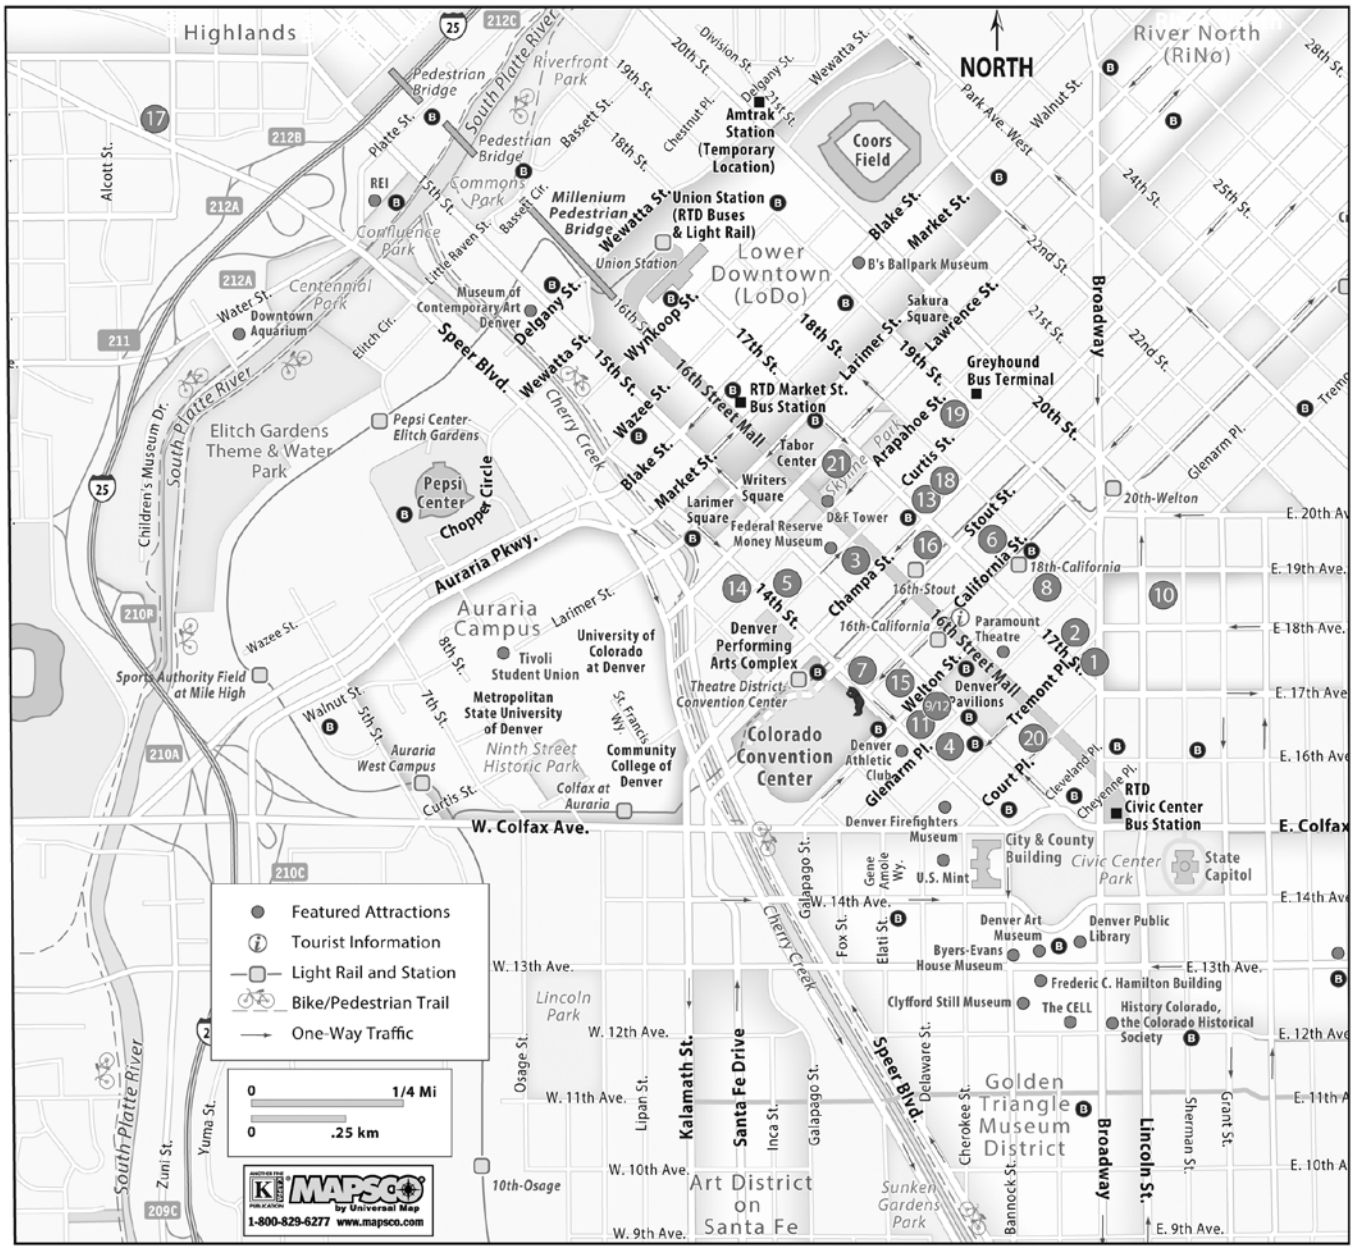

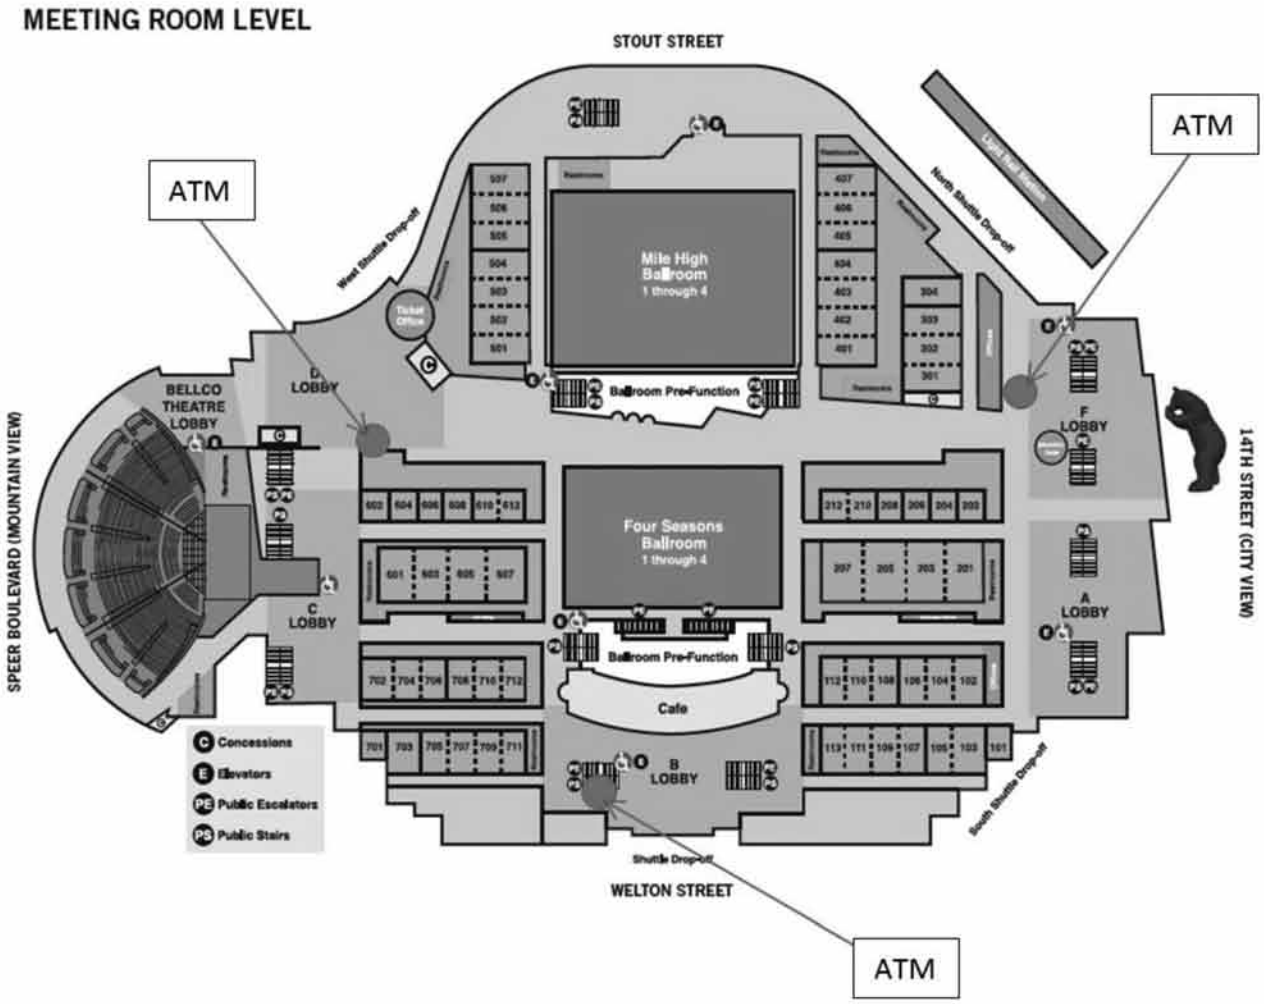

ATMs

Three ATMs are located within the convention center, and you can find these located in the F Lobby, B Lobby and the D Lobby as shown above.

Restrooms

Restrooms are located conveniently throughout the convention center, as follows:

Exhibit Halls:

Hall A south & east

Hall B south

Hall C south & west

Hall D west & north

Hall E north

Hall F north

Street Level (refer to map above, highlighted in blue)**:**

A, C, D, Lobby

B Lobby east & west

300, 400, 500 Corridor

F Lobby

Ballrooms:

Pre-function space

Bellco Theatre:

Lobby & Corridor

Saturday, November 16

Time	Event	Session	Location
Noon-4:30	HPC Educators	LittleFe Setup	702/704/706
1:00-5:30	HPC Educators	Registration	708/710/712
1:30-5:00	Broader Engagement	Registration	705/707/709/711
5:00-6:00	Broader Engagement	Orientation	705/707/709/711
6:00-7:00	HPC Educators	Orientation	705/707/709/711
7:00-9:00	BE/HPC Educators	Networking Event	B Lobby

Sunday, November 17

Time	Event	Session	Location
8:30-10:00	BE/HPC Educators	Plenary I: Making Parallelism Easy: A 25-Year Odyssey	705/707/709/711
8:30-Noon	Tutorial	Hybrid MPI and OpenMP Parallel Programming	407
8:30-Noon	Tutorial	InfiniBand/High-speed Ethernet for Dummies	201/203
8:30-Noon	Tutorial	A “Hands-on” Large Scale Visualization with ParaView	405
8:30-5:00	Tutorial	Introduction to OpenMP	402
8:30-5:00	Tutorial	Advanced MPI Programming	406
8:30-5:00	Tutorial	Debugging MPI & Hybrid/Heterogeneous Applications at Scale	404
8:30-5:00	Tutorial	Globus Online and the Science DMV…	403
8:30-5:00	Tutorial	Hands-on Practical Hybrid Parallel Application Performance Engineering	301
8:30-5:00	Tutorial	OpenACC: Productive, Portable Performance on Hybrid Systems…	401
8:30-5:00	Tutorial	Parallel Computing 101	303
8:30-5:00	Tutorial	Programming for the Intel Xeon Phi	205/207
8:30-5:00	Tutorial	Structured Parallel Programming with Patterns	302
9:00-5:30	Workshop	1st Workshop on Sustainable Software for Science…	503
9:00-5:30	Workshop	4th Workshop on Petascale (Big) Data Analytics	501
9:00-5:30	Workshop	6th Workshop on Many-Task Computing on Clouds, Grids…	502
9:00-5:30	Workshop	Building Energy Efficient HPC…	603
9:00-5:30	Workshop	IA^3 2013 - 3rd Workshop on Irregular Applications, Architectures…	504
9:00-5:30	Workshop	3rd Intl Workshop on Network-Aware Data Management	601
9:00-5:30	Workshop	4th Intl Workshop on Data-Intensive Computing in the Clouds	507
9:00-5:30	Workshop	8th Workshop on Ultrascale Visualization	505
9:00-5:30	Workshop	8th Workshop on Workflows in Support of…	506
10:30-Noon	Broader Engagement	Fortran 2008 Coarrays and Performance Analysis	705/707/709/711
10:30-Noon	HPC Educators	Supercomputing in Plain English	708/710/712
10:30-5:00	HPC Educators	LittleFe Buildout	702/704/706
1:30-3:00	Broader Engagement	Hands-on Exercise: Performance Analysis with TAU	705/707/709/711
1:30-5:00	HPC Educators	Curriculum Workshop: Mapping CS2013 & NSF/TCPP	708/710/712
1:30-5:00	Tutorial	Advanced Topics in InfiniBand/High-Speed Ethernet for Designing Hi-End…	201/203
1:30-5:00	Tutorial	Practical Fault Tolerance on Today’s HPC Systems	405
1:30-5:00	Tutorial	Scaling I/O Beyond 100,000 Cores using ADIOS	407

Sunday, November 17 (Continued)

Time	Event	Session	Location
3:30-5:00	Broader Engagement	Graphics and Visualization Technologies	705/707/709/711
6:00-9:00	Social Event	Exhibitors' Party	Buses depart from B Lobby

Monday, November 18

Time	Event	Session	Location
8:30-10:00	BE/HPC Educators	Plenary Talk II	705/707/709/711
8:30-Noon	Tutorial	Advanced PGAS Programming in UPC	303
8:30-Noon	Tutorial	Effective Procurement of Supercomputers	302
8:30-Noon	Tutorial	How to Analyze the Performance of Parallel Codes 101	301
8:30-5:00	Tutorial	Advanced OpenMP: Performance and 4.0 Features	406
8:30-5:00	Tutorial	Asynchronous Hybrid and Heterogeneous Parallel Programming with MPI/OmpSs...	401
8:30-5:00	Tutorial	Debugging and Optimizing MPI and OpenMP Applications Running on CUDA...	405
8:30-5:00	Tutorial	Ensuring Network Performance with perfSONAR	402
8:30-5:00	Tutorial	Linear Algebra Libraries for HPC: Scientific Computing with Multicore and Accelerators	407
8:30-5:00	Tutorial	OpenCL: A Hands-On Introduction	403
8:30-5:00	Tutorial	Parallel I/O in Practice	205/207
8:30-5:00	Tutorial	Python in HPC	201/203
8:30-5:00	Tutorial	The Practitioner's Cookbook for Good Parallel Performance on Multi- and...	404
9:00-Noon	Workshop	"Building" Energy Efficient High Performance Computing: 4th Annual EE HPC WG...	603
9:00-5:30	Workshop	4th Intl Workshop on Performance Modeling, Benchmarking and Simulation of...	502
9:00-5:30	Workshop	4th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (ScalA)	507
9:00-5:30	Workshop	Energy Efficient SuperComputing (E2SC)	210/120 (morning) 603 (afternoon)
9:00-5:30	Workshop	Extreme-Scale Programming Tools	501
9:00-5:30	Workshop	Python for High Performance and Scientific Computing (PyHPC 2013)	505
9:00-5:30	Workshop	The 6th Workshop on HPC Finance	601
9:00-5:30	Workshop	The 8th Parallel Data Storage Workshop	503
9:00-5:30	Workshop	WOLFHPC: Workshop on Domain-Specific Languages and High-Level Frameworks for HPC	506
9:00-6:00	Workshop	The 2nd International Workshop on Data Intensive Scalable Computing Systems (DISCS)	504
10:30-Noon	Broader Engagement	HPC Applications	705/707/709/711
10:30-5:00	HPC Educators	Workshop on Parallel, Distributed, and HPC in Undergraduate Curricula	708/710/712
10:30-5:00	HPC Educators	LittleFe Buildout	702/704/706
1:30-3:00	Broader Engagement	Session V: Mentor/Protege Session & Mixer	705/707/709/711

Monday, November 18 (Continued)

Time	Event	Session	Location
1:30-5:00	Tutorial	An Overview of Fault-Tolerant Techniques for HPC	301
1:30-5:00	Tutorial	Effective HPC Visualization and Data Analysis using VisIt	303
1:30-5:00	Tutorial	Introducing R: from Your Laptop to HPC and Big Data	302
3:00-5:00	HPC Undergraduate	Experiencing HPC For Undergraduates Orientation	703
3:30-5:00	BE/Student Volunteers	Professional Development I: Navigating a Computing Career	705/707/709/711
6:00-6:30	HPC Interconnections	SC First-Timers Introduction	705/707/709/711
6:00-9:00	Emerging Technologies	Emerging Technologies	Booth 3547
6:00-9:00	Scientific Viz Showcase	Exhibit	Mile High Pre-Function
7:00-8:00	HPC Impact Showcase	HPC Impact Showcase Monday	Booth 3947
7:00-9:00	Student Cluster Competition	Kickoff	Booth 3146

Tuesday, November 19

Time	Event	Session	Location
8:30-10:00	Keynote	The Secret Life of Data (Genevieve Bell)	Mile High
8:30-5:00	Posters	ACM Student Research Competition Posters	Mile High Pre-Function
8:30-5:00	Posters	Research Posters	Mile High Pre-Function
8:30-5:00	Scientific Viz Showcase	Exhibit	Mile High Pre-Function
10:00-5:30	Emerging Technologies	Emerging Technologies	Booth 3547
10:00-6:00	Student Cluster Competition	Student Cluster Competition	Booth 3146
10:20-2:30	HPC Impact Showcase	Showcase Tuesday	Booth 3947
10:30-Noon	Awards	ACM Gordon Bell Finalists I	201/203
10:30-Noon	Papers	Data Management in the Cloud	205/207
10:30-Noon	Papers	Graph Partitioning and Data Clustering	401/402/403
10:30-Noon	Papers	Inter-Node Communication	405/406/407
10:30-Noon	BE/Student Volunteers	Professional Development II: Building Your Technical Resume	601/603
10:30-Noon	Exhibitor Forum	Hardware & Architecture I	501/502
10:30-Noon	Exhibitor Forum	Moving, Managing & Storing Data I	503/504
10:30-Noon	HPC Educators	Special Session on Education and Workforce Development	708/710/712
10:30-5:00	HPC Educators	LittleFe Buildout	702/704/706
10:30-Noon	HPC Undergraduate	Introduction to HPC Research Topics	703
10:30-Noon	Invited Talks	Invited Talks (1)	Mile High
12:15-1:15	SIGHPC Meeting	ACM SIGHPC Annual Members Meeting	605
12:15-1:15	Birds of a Feather	Best Practices for Commissioning Liquid Cooling Infrastructure	404
12:15-1:15	Birds of a Feather	Building on the European Exascale Approach	210/212
12:15-1:15	Birds of a Feather	Collaborative Opportunities with the Open Science Data Cloud	205/207
12:15-1:15	Birds of a Feather	Getting Scientific Software Installed: Tools and Best Practices	708/710/712
12:15-1:15	Birds of a Feather	Hadoop 2 and HPC: Beyond MapReduce	507
12:15-1:15	Birds of a Feather	Integration of NVM Technologies in HPC Architectures: Challenges and Opportunities	601/603
12:15-1:15	Birds of a Feather	Maximize Data Center Power Efficiency through Energy Aware Computing	201/203
12:15-1:15	Birds of a Feather	Open MPI State of the Union	301/302/303

Tuesday, November 19 (Continued)

Time	Event	Session	Location
12:15-1:15	Birds of a Feather	Python for High Performance and Scientific Computing	401/402/403
12:15-1:15	Birds of a Feather	Scalable Adaptive Graphics Environment (SAGE) for Global Collaboration	501/502
12:15-1:15	Birds of a Feather	Seventh Graph500 List	705/707/709/711
12:15-1:15	Birds of a Feather	The Lustre Community: At the Intersection of HPC and Big Data	405/406/407
12:15-1:15	Birds of a Feather	The Role of Software-Defined Networking in the Big Data Equation	703
12:15-1:15	Birds of a Feather	TORQUE: Where Are We Now and Where Are We Going	503/504
1:30-3:00	Awards	Cray/Fernbach/Kennedy Award Recipients Talks	Mile High
1:30-3:00	Invited Talks	Invited Talks (2)	Mile High
1:30-3:00	Panel	Future of Memory Technology for Exascale and Beyond	301/302/303
1:30-3:00	Papers	Cloud Resource Management and Scheduling	205/207
1:30-3:00	Papers	Energy Management	401/402/403
1:30-3:00	Papers	Extreme-Scale Applications	201/203
1:30-3:00	Papers	Load Balancing	405/406/407
1:30-5:00	HPC Educators	Serious Computational Examples for Science Classes Featuring Python, Mathematica...	705/707/709/711
1:30-5:00	HPC Educators	Strategies for Introducing Parallelism with Python	708/710/712
3:30-5:00	BE/Student Volunteers	SV Professional Development III: HPC Careers and Networking Panel	601/603
3:30-5:00	Exhibitor Forum	Hardware & Architecture II	501/502
3:30-5:00	Exhibitor Forum	Moving, Managing & Storing Data II	503/504
3:30-5:00	Invited Talks	Invited Talks (3)	Mile High
3:30-5:00	Panel	Exascale Runtime Systems	301/302/303
3:30-5:00	Papers	Fault Tolerance and Migration in the Cloud	205/207
3:30-5:00	Papers	I/O Tuning	405/406/407
3:30-5:00	Papers	Physical Frontiers	201/203
3:30-5:00	Papers	System-wide Application Performance Assessments	401/402/403
5:15-7:00	Posters	ACM Student Research Competition Poster Reception	Mile High Pre-Function
5:15-7:00	Posters	Research Poster Reception	Mile High Pre-Function
5:15-7:00	Scientific Viz Showcase	Reception	Mile High Pre-Function
5:30-7:00	Birds of a Feather	Asynchronous and Adaptive Parallel Programming with Charm++	702/704/706
5:30-7:00	Birds of a Feather	Campus Bridging with XSEDE and Globus Online	601/603
5:30-7:00	Birds of a Feather	Defining BigData: Industry Views on Real-Time Data, Analytics, and HPC Technologies...	405/406/407
5:30-7:00	Birds of a Feather	Early Experiences Developing and Debugging on the Intel® Xeon Phi Coprocessor	401/402/403
5:30-7:00	Birds of a Feather	Eclipse Parallel Tools Platform (PTP)	210/212
5:30-7:00	Birds of a Feather	Energy Efficient High Performance Computing	205/207
5:30-7:00	Birds of a Feather	G8 Extreme Scale and Big Data Program: Progress on Current Exascale Projects...	501/502
5:30-7:00	Birds of a Feather	HPC and the Web	703
5:30-7:00	Birds of a Feather	Library of Mini-Applications for Exascale Component-Based Performance Modelling	708/710/712
5:30-7:00	Birds of a Feather	MPICH: A High-Performance Open-Source MPI Implementation	201/203
5:30-7:00	Birds of a Feather	OpenMP Goes Heterogeneous With OpenMP 4.0	301/302/303

Tuesday, November 19 (Continued)

Time	Event	Session	Location
5:30-7:00	Birds of a Feather	Science and Scientific Workflows: Putting Workflows to Work	404
5:30-7:00	Birds of a Feather	SUPReMM: Comprehensive Open Source Resource Management	507
5:30-7:00	Birds of a Feather	TOP500 Supercomputers	Mile High
5:30-7:00	Birds of a Feather	Towards Exascale Runtime Systems: Challenges and Opportunities	503/504
5:30-7:00	Birds of a Feather	Trends in Small HPC Center Management	705/707/709/711

Wednesday, November 20

Time	Event	Session	Location
8:30-10:00	Invited Talks	Invited Talks (4)	Mile High
8:30-5:00	Posters	ACM Student Research Competition Poster Exhibit	Mile High Pre-Function
8:30-5:00	Posters	Research Poster Exhibit	Mile High Pre-Function
8:30-5:00	Scientific Viz Showcase	Exhibit	Mile High Pre-Function
10:00-5:30	Emerging Technologies	Emerging Technologies	Booth 3547
10:00-3:00	HPC Interconnections	Student Job/Opportunity Fair	605/607
10:00-4:30	Student Cluster Competition	Student Cluster Competition	Booth 3146
10:20-2:30	HPC Impact Showcase	Showcase Wednesday	Booth 3947
10:30-Noon	Awards	ACM Gordon Bell Finalists II	201/203
10:30-Noon	Panel	Large Supercomputers and Networks (International Panel)	301/302/303
10:30-Noon	Papers	Optimizing Data Movement	205/207
10:30-Noon	Papers	Parallel Programming Models and Compilation	405/406/407
10:30-Noon	Papers	Performance Management of HPC Systems	401/402/403
10:30-Noon	Exhibitor Forum	Effective Application of HPC I	503/504
10:30-Noon	Exhibitor Forum	HPC Futures & Exascale I	501/502
10:30-Noon	HPC Educators	An Educator's Toolbox for CUDA (Part 1)	708/710/712
10:30-Noon	HPC Educators	Unveiling parallelization strategies at undergraduate level	702/704/706
10:30-Noon	HPC Undergraduate	Graduate Student Perspective	703
10:30-Noon	Invited Talks	Invited Talks (5)	Mile High
12:15-1:15	Birds of a Feather	Application Migration and Performance Expectation for Manycore Programming	205/207
12:15-1:15	Birds of a Feather	Applications of LLVM to HPC	201/203
12:15-1:15	Birds of a Feather	At the Intersection of Big Data and Extreme Computing	503/504
12:15-1:15	Birds of a Feather	Chapel Lightning Talks 2013	601/603
12:15-1:15	Birds of a Feather	Community MOOC's for Computing Technologies and Applications	705/707/709/711
12:15-1:15	Birds of a Feather	Heterogeneous Computing Platforms: Merging HPC and Embedded	708/710/712
12:15-1:15	Birds of a Feather	High Precision Arithmetic Operations: Libraries and Applications	301/302/303
12:15-1:15	Birds of a Feather	High-Performance Communications for High-Performance Computing	404
12:15-1:15	Birds of a Feather	HPC Job Scheduling Challenges and Successes from the PBS Community	702/704/706
12:15-1:15	Birds of a Feather	HPC Systems Engineering and Administration	501/502
12:15-1:15	Birds of a Feather	INCITE and Leadership-Class Systems	703
12:15-1:15	Birds of a Feather	Opportunities and Barriers for Computational Science Education	507
12:15-1:15	Birds of a Feather	PGAS: The Partitioned Global Address Space Programming Model	401/402/403

Wednesday, November 20 (Continued)

Time	Event	Session	Location
12:15-1:15	Birds of a Feather	The Open Community Runtime (OCR) Framework for Exascale Systems	405/406/407
12:15-1:15	Birds of a Feather	Total Power Usage Effectiveness: A New Take on PUE	210/212
1:30-3:00	Panel	RDMA: Scaling the I/O Architecture for Future Applications	301/302/303
1:30-3:00	Papers	Fault-Tolerant Computing	201/203
1:30-3:00	Papers	In-Situ Data Analytics and Reduction	205/207
1:30-3:00	Papers	Parallel Performance Tools	405/406/407
1:30-3:00	Papers	Preconditioners and Unstructured Meshes	401/402/403
1:30-3:00	Exhibitor Forum	Effective Application of HPC II	503/504
1:30-3:00	Exhibitor Forum	HPC Futures & Exascale II	501/502
1:30-3:00	HPC Educators	An Educator's Toolbox for CUDA (Part 2)	708/710/712
1:30-3:00	Invited Talks	Invited Talks (6)	Mile High
1:30-3:00	Student Cluster Competition	Get Involved in the Student Cluster Competition	703
1:30-5:00	HPC Educators	CSinParallel: Using Map-Reduce to Teach Data-Intensive Scalable Computing...	702/704/706
3:30-5:00	Posters	ACM Student Research Competition Poster Presentations	404
3:30-5:00	Broader Engagement	Session VI: Journal Publishing 101	703
3:30-5:00	Exhibitor Forum	HPC Futures & Exascale III	501/502
3:30-5:00	HPC Educators	An Educator's Toolbox for CUDA (Part 3)	708/710/712
3:30-5:00	Invited Talks	Invited Talks (7)	Mile High
3:30-5:00	Panel	Fault Tolerance/Resilience at Petascale/Exascale: Is it Really Critical?...	301/302/303
3:30-5:00	Papers	Engineering Scalable Applications	401/402/403
3:30-5:00	Papers	Improving Large-Scale Computation and Data Resources	205/207
3:30-5:00	Papers	Tools for Scalable Analysis	405/406/407
4:30-5:30	Student Cluster Competition	Grand Finale	Booth 3146
5:30-6:00	Student Cluster Competition	"Celebrity Pro-Am Cluster Challenge" Kickoff	Booth 3146
5:30-7:00	Birds of a Feather	Big Data, Big Compute: Data-Intensive and Extreme-Scale Computing	301/302/303
5:30-7:00	Birds of a Feather	Codesign for the Department of Energy's Computational Science Community	702/704/706
5:30-7:00	Birds of a Feather	Cost-Benefit Quantification for HPC: An Inevitable Challenge	601/603
5:30-7:00	Birds of a Feather	Creating a Training Roadmap for New Computational Consultants and Their Users	703
5:30-7:00	Birds of a Feather	Critically Missing Pieces in Heterogeneous Accelerator Computing	401/402/403
5:30-7:00	Birds of a Feather	Drilling Down: Understanding User-Level Activity on Today's Supercomputers	205/207
5:30-7:00	Birds of a Feather	Evolution of and Experiences with OpenACC	705/707/709/711
5:30-7:00	Birds of a Feather	Exascale IO Initiative: Progress Status	708/710/712
5:30-7:00	Birds of a Feather	OpenCL: Version 2.0 and Beyond	405/406/407
5:30-7:00	Birds of a Feather	OpenSHMEM: Further Developing a Standard for the PGAS & SHMEM Community	201/203
5:30-7:00	Birds of a Feather	Reconfigurable Supercomputing	503/504
5:30-7:00	Birds of a Feather	Research Data Alliance (RDA) for HPC	501/502
5:30-7:00	Birds of a Feather	Super-R: Supercomputing and R for Data-Intensive Analysis	404
5:30-7:00	Birds of a Feather	Techniques and Strategies for Extreme Scale Debugging	210/212
5:30-7:00	Birds of a Feather	The Green500 List and Its Evolution	Mile High

Wednesday, November 20 (Continued)

Time	Event	Session	Location
5:30-7:00	Birds of a Feather	Women in HPC Around the World	507
6:00-10:00	Student Cluster Competition	Dinner	Booth 3146

Thursday, November 21

Time	Event	Session	Location
7:30-8:30	HPC Undergraduate	Experiencing HPC For Undergraduates Breakfast	703
8:30-10:00	Invited Talks	Invited Talks (8)	Mile High
8:30-Noon	Scientific Viz Showcase	Exhibit	Mile High Pre-Function
8:30-5:00	Posters	ACM Student Research Competition Poster Exhibit	Mile High Pre-Function
8:30-5:00	Posters	Research Poster Exhibit	Mile High Pre-Function
10:00-14:00	Emerging Technologies	Emerging Technologies	Booth 3547
10:00-11:30	Student Cluster Competition	“Celebrity Pro-Am Cluster Challenge”	Booth 3146
10:20-1:00	HPC Impact Showcase	Showcase Thursday	Booth 3947
10:30-Noon	Panel	Emerging Technologies and Big Data (Euro-Centric)	301/302/303
10:30-Noon	Papers	GPU Programming	205/207
10:30-Noon	Papers	Matrix Computations	401/402/403
10:30-Noon	Papers	Memory Resilience	405/406/407
10:30-Noon	Papers	Performance Analysis of Applications at Large Scale	201/203
10:30-Noon	Doctoral Showcase	Dissertation Research	601/603
10:30-Noon	Exhibitor Forum	Moving, Managing & Storing Data III	503/504
10:30-Noon	Exhibitor Forum	Software for HPC I	501/502
10:30-Noon	HPC Undergraduate	Careers in HPC	703
10:30-Noon	Invited Talks	Invited Talks (9)	Mile High
10:30-Noon	Network Research Exhibition	Network Research Exhibition	404
10:30-12/1:30-3:00	HPC Educators	Going Parallel with C++11 (Parts 1 &2)	702/704/706
12:15-1:15	Birds of a Feather	Application Grand Challenges in the Heterogeneous Accelerator Era	401/402/403
12:15-1:15	Birds of a Feather	Computer Architecture Repositories for Open Simulation and Modeling Tools	404
12:15-1:15	Birds of a Feather	Coprocessors, GPUs, MICs: Enticing Users to Jump onto Accelerators	405/406/407
12:15-1:15	Birds of a Feather	Discussing an I/O Framework to Accelerate Improvements in Application I/O Performance	601/603
12:15-1:15	Birds of a Feather	Ethernet’s Rate Roadmap	702/704/706
12:15-1:15	Birds of a Feather	Harnessing Accelerator Technology for Next-Gen Sequencing Bioinformatics	503/504
12:15-1:15	Birds of a Feather	HPC Training Perspectives and Collaborations from PRACE, XSEDE and RIKEN AICS	501/502
12:15-1:15	Birds of a Feather	New Developments in the APGAS Programming Model and X10	507
12:15-1:15	Birds of a Feather	Petascale Systems with Intel Xeon Phi Co-Processors	301/302/303
12:15-1:15	Birds of a Feather	Slurm User Group Meeting	205/207
12:15-1:15	Birds of a Feather	The 2013 HPC Challenge Awards	201/203
12:15-1:15	Birds of a Feather	The Message Passing Interface: Version 3.0 and What Comes Next?	705/707/709/711
12:15-1:15	Birds of a Feather	The UDT Forum: A Community for UDT Developers and Users	703
12:30-1:30	Awards	SC13 Conference Awards Presentations	Mile High
1:30-3:00	Awards	ACM Athena Award / SC13 Test of Time Award Special Lectures	Mile High

Thursday, November 21 (Continued)

Time	Event	Session	Location
1:30-3:00	Doctoral Showcase	Dissertation Research	601/603
1:30-3:00	Exhibitor Forum	Moving, Managing & Storing Data IV	503/504
1:30-3:00	Exhibitor Forum	Software for HPC II	501/502
1:30-3:00	Network Research Exhibition	Network Research Exhibition	404
1:30-3:00	Panel	Extreme Computing and Industrial Competitiveness	301/302/303
1:30-3:00	Papers	Memory Hierarchy	405/406/407
1:30-3:00	Papers	MPI Performance and Debugging	201/203
1:30-3:00	Papers	Sorting and Graph Algorithms	205/207
3:00-5:00	Broader Engagement	Wrap Up/Scavenger Hunt Awards	705/707/709/711
3:30-5:00	Doctoral Showcase	Early Research Showcase	601/603
3:30-5:00	Exhibitor Forum	Moving, Managing & Storing Data V	503/504
3:30-5:00	Exhibitor Forum	Software for HPC III	501/502
3:30-5:00	HPC Educators	High-level parallel programming using Chapel (Part 2)	708/710/712
3:30-5:00	Panel	SC13 Silver Anniversary: Retrospective on Supercomputing Technologies	Mile High
3:30-5:00	Papers	Application Performance Characterization	405/406/407
3:30-5:00	Papers	Optimizing Numerical Code	401/402/403
7:00-10:00	Special & Invited Event	Conference Reception	B Lobby for Buses

Friday, November 22

Time	Event	Session	Location
8:30-10:00	Panel	Programming and Managing Systems At Scale	401/402/403
8:30-10:00	Panel	The Importance of Metadata in a World Where HPC is Everywhere	405/406/407
8:30-Noon	Workshop	1st Intl Workshop on Software Engr for HPC in Computational Science and Engineering	210/212
8:30-Noon	Workshop	3rd International Workshop on HPC, Networking and Analytics for the Power Grid	601/603
8:30-Noon	Workshop	8th Workshop on Virtualization in High-Performance Cloud Computing (VHPC '13)	503/504
8:30-Noon	Workshop	Exascale MPI	201/203
8:30-Noon	Workshop	Obtaining Bitwise Reproducible Results: Perspectives and Latest Advances	301/302/303
8:30-Noon	Workshop	VISTech Workshop: Visualization Infrastructure and Systems Technology	205/207
8:30-Noon	Workshop	Workshop on Large-Scale Data Analytics (LSDA)	501/502
10:30-Noon	Panel	Big Computing: From the Exa-Scale to the Sensor-Scale	405/406/407
10:30-Noon	Panel	Investing in Higher Education and Software for the Future	401/402/403

Keynote/Invited Talks/Panels

In the past, SC conferences have featured a variety of invited talks under various names such as Masterworks, plenary talks, and state of the field. To improve the conference experience, SC13 combines them under a single banner: Invited Talks.

Overall, SC serves as the premier venue for high performance computing, networking and storage. Thus, the Invited Talks feature leaders who detail innovative work in these areas and their application to the world's most challenging problems. These talks often provide a longer-term perspective than individual research papers and puts multiple research insights into broader context. At SC13 you can hear about the hardest challenges, the latest innovations in supercomputing and data analytics, and how new approaches are addressing the toughest and most complex questions of our time.

Panels at SC13 will be, as in past years, among the most important and heavily attended events of the conference. Panels will bring together the key thinkers and producers in the field to consider in a lively and rapid-fire context some of the key questions challenging HPC this year.

Panels bring a rare opportunity for mutual engagement of community leaders and broad mainstream contributors in a face-to-face exchange through audience participation and questioning. Surprises are the norm at panels, which make for exciting and lively hour-and-a-half sessions. Panels can explore topics in depth by capturing the opinions of a wide range of people active in the relevant fields. Panels represent state of the art opinions, and can be augmented with social media, twitter and video feeds, and even real-time audience polling. Please plan on attending one or more of the panel offerings at SC13. We look forward to your help, through your participation, in making panels at SC13 a major success and lots of fun.

Keynote

Tuesday, November 19

8:30am-10am
Room: Mile High

Keynote Address: The Secret Life of Data

Genevieve Bell
Technology Anthropologist, Intel Fellow

Today Big Data is one of the hottest buzzwords in technology, but from an anthropological perspective Big Data has been with us for millennia, in forms such as census information collected by ancient civilizations. The next 10 years will be shaped primarily by new algorithms that make sense of massive and diverse datasets and discover hidden value. Could we ignite our creativity by looking at data from a fresh perspective? What if we designed for data like we design for people? This talk will explore the secret life of data from an anthropological point of view to allow us to better understand its needs—and its astonishing potential—as we embark to create a new data society.

Biography: Born and raised in Australia, Dr. Genevieve Bell is an Intel Fellow and the Director of Interaction and Experience Research at Intel Corporation in Portland, Oregon. Genevieve is the driving force behind Intel's consumer-centered focus. Gathering a team of anthropologists, interaction designers and human factors engineers to transform consumer-centric product innovation, she has fundamentally changed how Intel envisions, plans and develops its platforms. Genevieve's team is responsible for setting research directions, conducting global comparative qualitative and quantitative research, leading new product strategy and definition and championing consumer-centric innovation and thinking across all of Intel's platforms. Dr. Bell has a PhD in anthropology from Stanford University and is co-author of a book on "Divining a Digital Future" from MIT Press. She was recognized by Fast Company magazine as one of the 100 most innovative people in business.

Invited Talks

Tuesday, November 19

Invited Talks 1

Chair: William Kramer (National Center for Supercomputing Applications)
10:30am-12pm
Room: Mile High

Extreme Scale Computational Biology

Klaus Schulten (University of Illinois Urbana-Champaign)

The molecular dynamics and visualization programs NAMD and VMD serve over 300,000 registered users in many fields of biology and medicine. These programs are on every kind of computer from personal laptops to massive supercomputers, and will soon be able to generate atomic-level views of entire living cells, opening a treasure chest of data for biotechnology, pharmacology and medicine. This lecture will summarize the impressive NAMD and VMD development over the last several years that led to the programs' excellent performance on Blue Waters, Titan and Stampede. It also describes the spectacular research discoveries enabled by petascale computing, which become more exciting every day and are leading to a veritable revolution in biology as researchers become able to describe very large structures which were previously inaccessible. I will further discuss ongoing effort in power-conscientious NAMD and VMD computing on ARM+GPU processors needed for utilizing the next generation of computers.

Biography: Klaus Schulten holds a Swanlund Chair in Physics at the University of Illinois, where he is also affiliated with the Department of Chemistry and the Center for Biophysics and Computational Biology. He is a full-time faculty member at the Beckman Institute of Applied Science and Technology, the founder and director of the Theoretical and Computational Biophysics Group, an NIH Resource for Macromolecular Modeling and Bioinformatics, and the co-director of the NSF Physics Frontiers Center, Center for the Physics of Living Cells, at Illinois. His research focuses on the structure and function of supramolecular systems in the living cell and on the development of non-equilibrium statistical mechanical descriptions and powerful, efficient computing tools for physical biology. In addition to his profound contributions to biophysics, Schulten has pioneered the development and application of high-performance computers to modeling large macromolecular systems in realistic environments. He has made ground-breaking contributions to our understanding of biomolecular processes coupled with mechanical force, bioelectronic processes in metabolism and vision, and the function and mechanisms of membrane proteins.

Visualization of Strong Ground Motion and Tsunami for the Great Earthquake of 2011 Off Tohoku, Japan

Takashi Furumura (Earthquake Research Institute, University of Tokyo)

The strong impact caused by the 2011 earthquake off Tohoku, Japan, (Mw=9.0) was reproduced by a large-scale FDM simulation using the K-Computer and the Earth Simulator with a detail earthquake source rupture model and a heterogeneous subsurface structure model. The visualized derived seismic wavefield illustrates the significance of this earthquake and the process in which large and long-duration ground motion is developed in populated cities and destructive tsunami is amplified along the Tohoku coast. The results of the simulation are compared the synthetic and observed waveforms from high-density nation-wide seismic network and offshore ocean bottom tsunami network both deployed over Japan. We believe the reliable disaster information should be effective in earthquake disaster reduction studies. Towards this goal, the development of integrated earthquake disaster prevention simulations by linking sequences of simulations for earthquake nucleation, seismic wave and tsunami propagation, and building oscillation in interaction with soft soil is presented.

Biography: Takashi Furumura is a Professor of the Research Center for Large-scale Earthquake, Tsunami and Disaster, Earthquake Research Institute at the University of Tokyo. He is also a Professor of the Center for Integrated Disaster Information Research, Interfaculty Initiative in Information Studies, Graduate School of Interdisciplinary Information Studies. He received his Ph.D. in Geophysics in 1992 from Hokkaido University. Since then, he had been working in Hokkaido University of Education till 2000 and then moved to the University of Tokyo. He specializes in the field of large-scale computer simulation, data assimilation, and visualization of earthquake ground motion and tsunami for understanding the strong ground motion and tsunami development process and for mitigating of disasters. He has been directing national projects aiming development of integrated disaster mitigation simulation based on HPC technology.

Invited Talks 2

Chair: Irene Qualters (NSF)
1:30pm-3pm
Room: 601/603

The Role of HPC in Education

Thom Dunning (University of Illinois at Urbana-Champaign)

Science and engineering research has been revolutionized by computation and opportunities abound for revolutionizing teaching and learning with computing technologies. But, to date computing has largely been used to organize, prepare for and disseminate courses. The potential of using these technologies to teach students the fundamental principles of a subject through authentic computational simulation is largely unexplored.

We have been using computational tools and simulations to help teach high school and college chemistry. The simulations enable students to gain a deep, rich appreciation for the basic principles of chemistry. Further, the use of computational tools is enthusiastically embraced by teachers and students, results in improved performance of both, and leads to increased student interest in chemistry.

This presentation will cover the experiences that we have had using computational tools and simulations in the high school chemistry classroom as well as in undergraduate organic chemistry courses.

Biography: Thom Dunning is the director of the Institute for Advanced Computing Applications and Technologies and the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign. He also holds an endowed position as Distinguished Chair for Research Excellence in Chemistry and professor in the Department of Chemistry. As leader of IACAT/NCSA, Dunning leads a staff of approximately 300 technologists and scientists who: • provide and support high-performance computing, data-intensive computing, and networking resources; • explore innovative computing architectures and techniques to achieve petascale (and beyond) science; • develop cyberenvironments tailored to the needs of research communities and software and tools to improve cybersecurity; • create artful visualizations of scientific phenomena; • and help prepare the next generation of scientists and engineers. Dunning previously held leadership positions at the Joint Institute for Computational Sciences at the University of Tennessee and Oak Ridge National Laboratory, the University of North Carolina System, the Office of Science at the U.S. Department of Energy, the Environmental Molecular Sciences Laboratory at Pacific Northwest National Laboratory, Argonne National Laboratory, and Los Alamos National Laboratory. He was instrumental in creating DOE's Scientific Discovery through Advanced Computing (SciDAC) program, the federal government's first comprehensive program aimed at developing the software infrastructure needed for leadership-class scientific computing. He is a fellow of the American Chemical Society, the American Physical Society, and the American Association for the Advancement of Science. Dunning received the American Chemical Society's Computers in Chemical and Pharmaceutical Research Award in 2011.

Gathering Intelligence from Massive Graphs

David Bader (Georgia Institute of Technology)

Abstract: TBA

Biography: David A. Bader is a Full Professor in the School of Computational Science and Engineering, College of Computing, at Georgia Institute of Technology, and Executive Director for High Performance Computing. He received his Ph.D. in 1996 from The University of Maryland, and his research is supported through highly-competitive research awards, primarily from NSF, NIH, DARPA, and DOE. Dr. Bader serves as a Board member of the Computing Research Association (CRA), and on the Steering Committees of the IPDPS and HiPC conferences. He is an associate editor-in-chief of the Journal of Parallel and Distributed Computing (JPDC), and serves as an associate editor for several high impact publications including IEEE Transactions on Computers (TC), ACM Transactions on Parallel Computing (TOPC), ACM Journal of Experimental Algorithmics (JEA), IEEE DSOnline, Parallel Computing, and Journal of Computational Science, and has been an associate editor for the IEEE Transactions on Parallel and Distributed Systems (TPDS). He was elected as chair of the IEEE Computer Society Technical Committee on Parallel Processing (TCPP) and as chair of the SIAM Activity Group in Supercomputing (SIAG/SC). He is also a leading expert on multicore, manycore, and multithreaded computing for data-intensive applications such as those in massive-scale graph analytics. He has co-authored over 130 articles in peer-reviewed journals and conferences, and his main areas of research are in parallel algorithms, combinatorial optimization, massive-scale social networks, and computational biology and genomics. Prof. Bader is a Fellow of the IEEE and AAAS, a National Science Foundation CAREER Award recipient, and has received numerous industrial awards from IBM, NVIDIA, Intel, Cray, Oracle/Sun Microsystems, and Microsoft Research. Bader is a co-founder of the Graph500 List for benchmarking "Big Data" computing platforms. Bader is recognized as a "RockStar" of High Performance Computing by InsideHPC and as HPCwire's People to Watch in 2012.

Invited Talks 3

Chair: Robert F. Lucas (Information Sciences Institute)
3:30pm-5pm
Room: Mile High

Performance Evaluation for Large Scale Systems: Closed Loop Control with Appropriate Metrics

Lizy Kurian John (University of Texas at Austin)

The number of cores/systems utilized in high end systems is increasing. On one hand, the availability of large number of nodes gives opportunities for isolating faulty nodes and finding enough working nodes; on the other hand, the presence of too many systems makes it difficult to isolate problems. Performance/power problems at one node, if undetected in a timely fashion, can result in massive problems. This talk will discuss challenges in evaluating large systems, and metrics to detect performance/power bugs, and potential approaches to creating reliable high-performance and energy-efficient systems with appropriate feedback.

Biography: Lizy Kurian John is B. N. Gafford Professor of Electrical Engineering at the University of Texas at Austin. Her research interests include computer architecture, performance evaluation, workload characterization, energy efficient computing, reconfigurable architectures, etc. Her research has been supported by the National Science Foundation, the State of Texas Advanced Technology program, Lockheed Martin, Semiconductor research Corporation (SRC), IBM, Intel, Motorola, DELL, AMD, Samsung, Oracle and Microsoft Corporations. She is recipient of several awards, and most recently The Pennsylvania State University Outstanding Engineering Alumnus (2011). She coauthored a book on Digital Systems Design using VHDL (Cengage Publishers), and edited a book on Computer Performance Evaluation and Benchmarking. She is an IEEE Fellow.

Uncertainty Quantification in Computational Models

Habib N. Najm (Sandia National Laboratories)

Models of physical systems involve inputs/parameters determined from empirical measurements, and therefore exhibiting a degree of uncertainty. Estimating the propagation of this uncertainty into computational model output predictions is crucial for purposes of validation, design optimization, and decision support. Recent years have seen significant developments in probabilistic methods for efficient uncertainty quantification (UQ) in computational models. These methods are grounded in the use of functional representations for random variables. In this context, the utility of Polynomial Chaos (PC) UQ methods has been demonstrated in a range of physical models, including structural mechanics, porous media, fluid dynamics, heat transfer, and chemistry. While high-dimensionality remains a challenge, great strides have been made in dealing with moderate dimensionality along with non-linearity and oscillatory dynamics. This talk will outline PC UQ methods, and their utility in the estimation of uncertain input parameters from empirical data, and for propagation of parametric uncertainty to model outputs.

Biography: Habib N. Najm is a Distinguished Member of the Technical Staff at Sandia National Laboratories in Livermore, CA. He received his PhD degree in Mechanical Engineering from MIT in 1989. He worked with Texas Instruments from 1989 through 1993, before joining Sandia National Laboratories. His work at Sandia involves a range of computational science research, including development of numerical methods and computational tools for modeling and analysis of reacting flow, and for uncertainty quantification in computational

models. He is the director of QUEST, a multi-institutional UQ research institute, focusing on uncertainty quantification in extreme scale computations. He has co-authored over 80 journal articles and 11 US patents.

Wednesday, November 20

Invited Talks 4

8:30am-10am
Room: Mile High

Climate Earth System Modeling for the IPCC Sixth Assessment Report (AR6): Higher Resolution and Complexity

Warren Washington (National Center for Atmospheric Research)

It is expected that higher resolution coupled model components will be used for another round of climate and Earth system simulations using the Community Earth System Model (CESM). In addition, the treatment of physical processes will be more complex and tightly coupled to other components. A challenge is that high performance computer systems will be much faster and more difficult to efficiently use. Some preliminary simulations will be presented with computational details.

Biography: Dr. Warren Washington is a Senior Scientist at the National Center for Atmospheric Research (NCAR) in Boulder, Colorado. This group uses state-of-the-art computer climate models to study present and future climate change. He has engaged in research for more than fifty years. He has had Presidential Appointments under Carter, Reagan, Clinton, and Bush, Jr. Administrations. More recently, he served on the National Science Board from 1994 to 2006 and he was Chair from 2002 to 2006. He has over 150 publications and co-authored with Claire Parkinson a book considered a standard reference on climate modeling -- Dr. Washington has many awards including being a member of the National Academy of Engineering, Presidency of the American Meteorological Society (1994), a member of American Philosophical Society and the American Academy of Arts and Sciences. He has honorary degrees from OSU and Bates College. This fall he will receive another honorary degree from the University of Massachusetts. In November 2010, he was awarded the National Medal of Science by President Obama, the nation's highest science award.

Data, Computation, and the Fate of the Universe

Saul Perlmutter (Lawrence Berkeley National Laboratory)

This talk will reach into the past to explain how integrating big data -- and careful analysis -- led to the discovery of the acceleration of the universe's expansion. It will go on to discuss the increasing importance and impact of coupling scientists with data, analysis, and simulation to gain future insights.

Biography: Saul Perlmutter is a 2011 Nobel Laureate, sharing the prize in Physics for the discovery of the accelerating expansion of the Universe. He is a professor of physics at the University of California, Berkeley, and Lawrence Berkeley National Laboratory, and the leader of the international Supernova Cosmology Project. In addition to other awards and honors, he is a member of the National Academy of Sciences and the American Academy of Arts and Sciences and a fellow of the American Physical Society and the American Association for the Advancement of Science. Perlmutter has also written popular articles, and has appeared in numerous PBS, Discovery Channel, and BBC documentaries. His interest in teaching scientific approaches to problem-solving for non-scientists led to Berkeley courses on Sense and Sensibility and Science and Physics & Music.

Invited Talks 5

Chair: Irene Qualters (NSF)
10:30am-12pm
Room: Mile High

The Transformative Impact of Computation in a Data-Driven World

Farnam Jahanian (National Science Foundation)

We are witnessing unprecedented growth of scientific and social data, deep integration of the cyber and physical worlds, wireless connectivity at broadband speeds, and seamless access to resources in the cloud. These advances are transforming the way we work, play, communicate, learn, and discover. Investments in ambitious, long-term research and infrastructure, as well as in the development of a workforce empowered by computation and data, enable these advances and are a national imperative. This talk will focus on the technological advances and emerging trends that are shaping our future and accelerating the pace of discovery and innovation across all science and engineering disciplines. It will also describe how these trends inform priorities and programs at National Science Foundation.

Biography: Farnam Jahanian leads the National Science Foundation Directorate for Computer and Information Science and Engineering (CISE). He guides CISE, with a budget of over $850 million, in its mission to uphold the Nation's leadership in scientific discovery and engineering innovation through its support of fundamental research in computer and information science and engineering and of transformative advances in cyberinfrastructure. Dr. Jahanian is on leave from the University of Michigan, where he holds the Edward S. Davidson Collegiate Professorship and served as Chair for Computer Science and Engineering from 2007 – 2011 and as Director of the Software Systems Laboratory from 1997 – 2000. His research on Internet infrastructure security formed the basis for the Internet security company Arbor Networks, which he co-founded in

2001 and served as Chairman until its acquisition in 2010. Dr. Jahanian holds a master's degree and a Ph.D. in Computer Science from the University of Texas at Austin. He is a Fellow of ACM, IEEE and AAAS.

Europe's Place in a Global Race

Richard Kenway (University of Edinburgh)

This talk describes Europe's response to the opportunities and challenges of high-performance computing over the past 25 years. Today, those opportunities continue to grow, but so do the challenges: how to exploit exascale architectures and big data, how to organize e-infrastructure on a continental scale, and how to deliver the promised impact on the economy, science and society? Europe perceives itself to be in a global race. Will it win?

Biography: Vice-Principal Professor Richard Kenway OBE FRSE FInstP CPhys DPhil BSc University of Edinburgh Professor Kenway was appointed to the Tait Chair of Mathematical Physics at the University of Edinburgh in 1994. His research explores non-perturbative aspects of theories of elementary particles using computer simulation of lattice gauge theories, particularly the strong interactions of quarks and gluons described by Quantum Chromodynamics (QCD). He led UK participation in the QCDOC project to build three 10 teraflop/s computers to simulate QCD, jointly with the USA and Japan, and these machines operated successfully from 2004 to 2011. He is the principal investigator on grants for a follow-on project with IBM and Columbia University to build and exploit a 1 petaflop/s prototype BlueGene/Q computer, which was installed at the University of Edinburgh in 2011. In 2002, he initiated the International Lattice Data Grid project, which provides a global infrastructure for sharing simulation data. As Vice-Principal, Professor Kenway is responsible for the UK High-Performance Computing Service and for promoting advanced computing technology to benefit academia and industry. For ten years, until it closed in 2011, his responsibilities included the UK National e-Science Centre. From 2008 to 2011, he was Head of the School of Physics and Astronomy. In the Queen's 2008 Birthday Honours, Professor Kenway was awarded an Officer of the Most Excellent Order of the British Empire (OBE) for services to science. He led the establishment of the Scientific Steering Committee of the Partnership for Advanced Computing in Europe (PRACE) and chaired it from 2010 to 2012. Currently, he chairs the Board of the Scottish Universities Physics Alliance (SUPA). He is a founder member of the UK e-Infrastructure Leadership Council.

Invited Talks 6

Chair: William Tang (Princeton University)
1:30pm-3pm
Room: Mile High

Scalable Computing on Tianhe-2 System

Yutong Lu (National University of Defense Technology)

Scalability is one of the big challenges for the post-petascale and exascale systems, we need integrated technologies from architecture design to system software stack design. This talk expresses the efforts from NUDT to design and implement the Tianhe-2 system to get better scalability. From the heterogenous to nero-heterogenous architecture, we got some comparison among cpu, gpu and mic. We design the high bandwidth, low latency communication system with offload collective operations, to improve the scalability of interconnection. We improve the multiple protocols in custom MPI system to support different application models. We design the hybrid hierarchy I/O architecture and file system to tackle the I/O bottleneck. The co-design policy through whole system for scalable computing could support the domain applications extend to even large scale efficiently.

Biography: Yutong Lu, Professor, Director of system software laboratory, school of computer science, National University of Defense Technology (NUDT), Changsha, China. She held the position of professor in the State key Laboratory of High Performance Computing in China. She got her B.S, M.S, and Dr. degree from NUDT. She has been involved in research and development work for five generations of Chinese domestic supercomputers. She was director designer for Tianhe-1A and Tianhe-2, which got the number 1 system in Nov. 2010 and Jun. 2013 Top500 lists respectively. Her research interests include parallel OS, high-speed communication, global file system, programming environment and MPI.

Invited Talks 7

Chair: Robert F. Lucas (Information Sciences Institute)
3:30pm-5pm
Room: Mile High

Integration of Quantum Computing into High Performance Computing

Colin P. Williams (D-Wave Systems, Inc.)

The media likes to portray quantum computing as being in a head-to-head race with high performance computing. In reality, we foresee ways for quantum computing to complement and enhance high performance computing and vice versa. In this talk I will describe D-Wave's approach to quantum computing including its operating principles, system architecture, evidence of quantumness, and report on our latest

performance data. In particular, I will describe D-Wave's new "Vesuvius" 512-qubit quantum processor and describe several examples of computational problems we have mapped to our architecture. I will then describe strategies for integrating our quantum processor into mainstream high performance computing systems. As our quantum processor is naturally well suited to solving discrete combinatorial optimization, machine learning and artificial intelligence problems the talk should be of broad interest to computer scientists, physicists, and engineers with interests in a wide range of application areas.

Biography: Colin is Director of Business Development & Strategic Partnerships at D-Wave Systems Inc. – the World's first quantum computer company – where he works with corporate and government clients to infuse D-Wave quantum computing technology into their products and services. In his doctoral work Colin developed an artificial intelligence system for reasoning about the physical world via qualitative and quantitative differential equations. At Xerox PARC he became interested in the links between statistical physics and computer science, invented the theory of computational phase transitions, and applied it to understanding the deep structure of NP-Hard problems. Later, Colin became interested in the connections between quantum physics and computer science. He wrote the first book on quantum computing, "Explorations in Quantum Computing", and followed it up with two others, started the Quantum Computing Group at the NASA Jet Propulsion Laboratory, Caltech, and quickly broadened its scope to include research on quantum communications, quantum key distribution, quantum sensors, and quantum metrology. In addition, he personally invented quantum algorithms for solving NP-Hard problems, computing quantum wavelet transforms, accelerating average case quantum search, and gravitational tomography. He also showed how to perform arbitrary non-unitary quantum computation probabilistically, wrote a CAD tool for automatically designing quantum circuits that implemented arbitrary desired quantum computations, and specialized it to use gate sets optimized for linear optical quantum computing, spintronic quantum computing and superconducting quantum computing. In 2012 he published a greatly expanded and updated edition of "Explorations in Quantum Computing". Colin holds a Ph.D. in artificial intelligence from the University of Edinburgh, a M.Sc. and D.I.C. in atmospheric physics and dynamics from Imperial College, University of London, and a B.Sc. (with Honors) in mathematical physics from the University of Nottingham.

Thursday, November 21

Invited Talks 8

Chair: William Kramer (National Center for Supercomputing Applications)
8:30am-10am
Room: Mile High

BIG DATA + BIG COMPUTE = An Extreme Scale Marriage for SMARTER SCIENCE?

Alok Choudhary (Northwestern University)

Knowledge discovery has been driven by theory, experiments and by large-scale simulations on high-performance computers. Modern experiments and simulations involving satellites, telescopes, high-throughput instruments, sensors, and supercomputers yield massive amounts of data. What has changed recently is that the world is creating massive amounts of data at an astonishing pace and diversity.

Processing, mining and analyzing this data effectively and efficiently will be a critical component as we can no longer rely upon traditional ways of dealing with the data due to its scale and speed. But there is a social aspect of acceleration, which is sharing of "big data" and unleashing thousands to ask questions and participate in discovery. This talk addresses the fundamental question "what are the challenges and opportunities for extreme scale systems to be an effective platform" for not only traditional simulations, but their suitability for data-intensive and data driven computing to accelerate time to insights.

Biography: Alok Choudhary is the Henry & Isabelle Dever Professor of Electrical Engineering and Computer Science and a professor at Kellogg School of Management. He is also the founder, chairman and chief scientist (served as its CEO during 2011-2013) of Voxsup Inc., a big data analytics and social media marketing company. He received the National Science Foundation's Young Investigator Award in 1993. He is a fellow of IEEE, ACM and AAAS. His research interests are in high-performance computing, data intensive computing, scalable data mining, computer architecture, high-performance I/O systems, software and their applications in science, medicine and business. Alok Choudhary has published more than 400 papers in various journals and conferences and has graduated 33 PhD students. Techniques developed by his group can be found on every modern processor and scalable software developed by his group can be found on many supercomputers. Alok Choudhary's work and interviews have appeared in many traditional media including New York Times, Chicago Tribune, The Telegraph, ABC, PBS, NPR, AdExchange, Business Daily and many international media outlets all over the world.

The Interplay Between Internet Security and Scale

Vern Paxson (University of California, Berkeley)

Internet Security poses fundamentally hard problems due to its adversarial nature, but the challenges become especially complex at large scales. This talk will examine a range of such scaling issues: network speed, traffic volume, user diversity, forensic analysis, and difficulties that attackers themselves face. I will discuss these both in the context of operational security and in terms of how such issues arise in cybersecurity research.

Biography: Vern Paxson is a professor of Electrical Engineering and Computer Sciences at UC Berkeley and the director of the Networking and Security group at the International Computer Science Institute in Berkeley. His research focuses heavily on measurement-based analysis of network activity and Internet attacks. He has worked extensively on high performance network monitoring and on cybercrime, and has had a staff affiliation with the Lawrence Berkeley National Lab for decades.

Invited Talks 9

Chair: William Kramer (National Center for Supercomputing Applications)
10:30am-12pm
Room: Mile High

Big Data for Big Cities

Steven Koonin (Center for Urban Science and Progress, New York University)

For the first time in history, more than half of the world's population lives in urban areas; by mid-century, 70 percent of the world's more than 9 billion people will live in cities. Enabling those cities to deliver services effectively, efficiently, and sustainably while keeping citizens safe, healthy, prosperous, and well-informed, is one of the most important undertakings in this century. Success will require that we understand the physical, environmental, and human dimensions of cities in a synthetic and synoptic manner. New technologies enable the acquisition, integration, and analysis of urban "big data" to inform that understanding. I will discuss some of the rationale, applications, and challenges in the nascent field of urban informatics, as well as the ways in which advances in extreme computing and analytics technologies can accelerate progress.

Biography: Steven E. Koonin was appointed as the founding Director of NYU's Center for Urban Science and Progress in April 2012. That consortium of academic, corporate, and government partners will pursue research and education activities to develop and demonstrate informatics technologies for urban problems in the "living laboratory" of New York City. He previously served as the U.S. Department of Energy's second Senate-confirmed Under Secretary for Science from May 19, 2009 through November 18, 2011. As Under Secretary for Science, Dr. Koonin functioned as the Department's chief scientific officer, coordinating and overseeing research across the DOE. Dr. Koonin particularly championed research programs in High Performance Simulation, Exascale Computing, Inertial Fusion Energy, and Greenhouse Gas Monitoring, Reporting, and Verification. He also provided technical counsel on diverse nuclear security matters. He joined the Caltech faculty in 1975, was a research fellow at the Neils Bohr Institute during 1976-1977, and was an Alfred P. Sloan Foundation Fellow during 1977-1979. He became a professor of theoretical physics at Caltech in 1981 and served as Chairman of the Faculty from 1989-1991. Dr. Koonin was the seventh provost of Caltech from 1995-2004. He has served on numerous advisory committees for the Department of Energy, the National Science Foundation, and the Department of Defense, including the Defense Science Board and the CNO's Executive Panel. He is a member of the National Academy of Sciences, the Council on Foreign Relations, and the Trilateral Commission, and a fellow of the American Physical Society, the American Association for the Advancement of Science, and the American Academy of Arts and Sciences. In 1985, Dr. Koonin received the Humboldt Senior U.S. Scientist Award and, in 1998 the Department of Energy's E.O. Lawrence Award for his "broad impact on nuclear many-body physics, on astrophysics, and on a variety of related fields where sophisticated numerical methods are essential; and in particular, for his breakthrough in nuclear shell model calculations centered on an ingenious method for dealing with the huge matrices of heavy nuclei by using path integral methods combined with the Monte Carlo technique."

Panels

Tuesday, November 19

Future of Memory Technology for Exascale and Beyond

Chair: Subhash Saini (NASA Ames Research Center)
1:30pm-3pm
Room: 301/302/303

Moderator: Richard Murphy (Micron Technology, Inc.)
Panelists: Shekhar Borkar (Intel Corporation), Bill Dally (NVIDIA), Andreas Hansson (ARM), Doug Joseph (IBM), Peter Kogge (University of Notre Dame), Troy Manning (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. This panel will focus on three critical questions: (1) What is the right memory technology for exascale and beyond? (2) Given that memory capacity per core has declined on supercomputers since 2003, what can be done to continue to facilitate scaling? (3) What about tighter processor/memory operations and other potentially disruptive architectures?

Exascale Runtime Systems

Chair: Alice Koniges (Lawrence Berkeley National Laboratory)
3:30pm-5pm
Room: 301/302/303

Moderator: Pavan Balaji (Argonne National Laboratory)
Panelists: Laxmikant Kale (University of Illinois at Urbana-Champaign), Kathy Yelick (University of California, Berkeley), Bronis de Supinski (Lawrence Livermore National Laboratory), Vivek Sarkar (Rice University), Thomas Sterling (Indiana University)

As we head to exascale, the computational hardware will undergo a significant revamp toward low-energy usage. General-purpose computational cores might be augmented with heterogeneous accelerator cores. Memory subsystems would move toward deeper multi-level caches, smaller cache-coherence domains and non-DRAM memories. Networks will trend toward more cost-scalable and performance-localized setups. How runtime systems can bridge the gap between such architectures and modern complex and dynamic HPC applications is an open question. Would the answer be an evolutionary hybrid approach using MPI between nodes and an OpenMP or OpenACC-like model within the node? Would the answer be a completely revolutionary dynamic runtime system that would rethink the computation and data movement model from the ground up? This panel will be a half-and-half setup---the first half will be short "appetizer" presentations by experts discussing both evolutionary and revolutionary solutions; the second half will be an open debate among the participants.

Wednesday, November 20

Large Supercomputers and Networks (International Panel)

Chair: Anne C. Elster (Norwegian University of Science & Technology/University of Texas at Austin)
10:30am-12pm
Room: 301/302/303

Moderator: Pete Beckman (Exascale Technology and Computing Institute, ANL)
*Panelists: Yutong Lu (National University of Defense Technology, China), Yutaka Ishikawa (University of Tokyo), Catherine Riviere (PRACE/GENCI), Marek T. Michalewicz (A*STAR Computational Resource Center, Singapore)*

This panel will discuss experiences and thoughts on large supercomputer installations as well as networking. What are the main challenges? Funding? Power? People? Appropriate applications? How can we best utilize all our hardware and what kind of network infrastructure do we need? What kind of large supercomputers are we expected and/or wanting to see next year? In five years?

The panel will include representatives from the world's largest supercomputing centers, including PRACE (Partnership for Advanced Computing in Europe) whose 25 membership countries provide a pan-European supercomputing infrastructure.

RDMA: Scaling the I/O Architecture for Future Applications

Chair: Anne C. Elster (Norwegian University of Science & Technology/University of Texas at Austin)
1:30pm-3pm
Room: 301/302/303

Moderator: Paul Grun (Cray Inc.)
Panelists: Sudip Dosanjh (Lawrence Berkeley National Laboratory), Rick Stevens (Argonne National Laboratory), Gary Grider (Los Alamos National Laboratory)

The goal for this panel is to bridge the gap between users of large-scale machines and the providers of I/O architectures on which they are built. RDMA has been a dominant technology in the evolution of HPC, but since its introduction in 2000 there have been significant changes in the underlying compute model, how those machines are deployed and the ways in which they are used is producing a growing mismatch between user programs and the I/O model. To reach exascale, we have to re-think how I/O is architected. For example, are there opportunities to: •Improve the scalability and reduce the power consumption of data movement to meet exascale requirements? •Improve performance allowing researchers to solve ever-larger problems? •Take advantage of co-processing, multi-core, emerging trends toward large scale solid state devices and other technologies? •Better align I/O systems with important models like distributed shared memory?

Fault Tolerance/Resilience at Petascale/Exascale: Is it Really Critical? Are Solutions Necessarily Disruptive?

Chair: Alice Koniges (Lawrence Berkeley National Laboratory)
3:30pm-5pm
Room: 301/302/303

Moderator: Franck Cappello (Argonne National Laboratory)
Panelists: Marc Snir (Argonne National Laboratory), Bronis De Supinski (Lawrence Livermore National Laboratory), Al Geist (Oak Ridge National Laboratory), John Daly (Department of Defense), Ana Gainaru (University of Illinois at Urbana Champaign), Satoshi Matsuoka (Tokyo Institute of Technology)

This panel will explore important questions regarding fault tolerance and resilience in present petascale supercomputers and future exascale systems. For domain scientists, application developers, system administrators, students, and researchers, it is of utmost importance to understand what fault tolerance/resilience experts are observing today on the latest generation of petascale systems and are projecting for future exascale systems.

Unfortunately, there is no common view of this problem. The degree of criticality of the problem, the way to approach it (engineering problem versus a research problem), the feasibility and applicability of classic and disruptive techniques are raising controversies.

In order to expose the diversity of the community perception on this problem, the panel will gather six experts with differing opinions on key questions. These experts will present their background, experience, and actions in this domain and respond to questions from the panel chair and the audience.

Thursday, November 21

Emerging Technologies and Big Data (Euro-Centric)

Chair: Anne C. Elster (Norwegian University of Science & Technology/University of Texas at Austin)
10:30am-12pm
Room: 301/302/303

Moderator: Marie-Christine Sawley (Intel Paris Exascale Lab, France)
Panelists: Sverre Jarp (CERN, Switzerland), Alison Kennedy (EPCC,UK), Alex Ramirez (Barcelona Supercomputing Center, Spain), Catherine Riviere (GENCI)

This panel will discuss emerging technologies such as low-powered systems and other European initiatives. Big Data and the European-based Mont-Blanc project which has selected Samsung's flagship Exynos platform as the building block to power its first integrated low-power HPC prototype featuring a dual-core 1.7GHz mobile CPU based on ARM Cortex-A15 architecture paired with an integrated Mali-T604 GPU will also be included in the discussions.

Extreme Computing and Industrial Competitiveness

Chair: Alice Koniges (Lawrence Berkeley National Laboratory)
1:30pm-3pm
Room: 301/302/303

Moderator: Walter Kirchner (Council on Competitiveness)
Panelists: Steven E. Koonin (New York University), J. Michael McQuade (United Technologies Corporation), Dona Crawford (Lawrence Livermore National Laboratory)

Industry, from pharmaceutical to oil and gas to automotive to aeronautics, has computing challenges that cannot be fully addressed using current high-end computing capabilities. Examples of these are atomistic to continuum simulations for materials and chemical macro parameters; turbulent combustion in closed and open volumetric geometries; and enhanced

oil & gas reservoir modeling. These challenges limit, despite marked advances in computing capabilities, the rate of innovation and of technology development to solve enterprise and societal problems. The next level of extreme computing, exascale, will provide a transformative opportunity for these industries and others to increase their innovation capacity and decrease time to market.

Nations that encourage and support the development and use of next -generation extreme computing for national and economic security will lead 21st century competitive economies. The undergirding technologies to accomplish the next highest level of compute, exascale, will become pervasive in other consumer electronic products.

SC13 Silver Anniversary: Retrospective on Supercomputing Technologies

Chair: Mary Hall (University of Utah)
3:30pm-5pm
Room: Mile High

Moderator: Mary Hall (University of Utah)
Panelists: Fran Berman (Rensselaer Polytechnic Institute), David Keyes (King Abdullah University of Science & Technology), Ken-ichi Miura (Fujitsu Laboratories, Ltd.), Warren Washington (National Center for Atmospheric Research), Hans Zima (University of Vienna)

What was the state-of-the-art in supercomputing when the conference began 25 years ago? What promising technologies have failed and what unlikely technologies have prospered? How has context from external factors influenced changes in technology? This panel will engage a set of experts with broad interests (from Architecture, Programming Models, Libraries, Applications, Analysis and Tools) in a retrospective of the supercomputing field over the 25 years of the conference's history.

Friday, November 22

Programming and Managing Systems At Scale

Chair: Alice Koniges (Lawrence Berkeley National Laboratory)
8:30am-10am
Room: 401/402/403

Moderator: Chad Harrington (Adaptive Computing)
Panelists: Don Maxwell (Oak Ridge National Laboratory), William Kramer (University of Illinois at Urbana Champaign), Michael Jackson (Adaptive Computing), Larry Kaplan (Cray Inc.), Morris Jette (SchedMD), John Hengeveld (Intel Corporation)

In "Programming and Managing Systems at Scale," an all-star roster of software and hardware providers as well as industry experts from the world's largest supercomputing centers will discuss how system architecture must be re-engineered in order to reach the scales the industry is capable of today and going forward. The age of exascale computing is fast approaching, but for that future to become fully realized, current performance obstacles—such as the inefficiencies of diagnosing failures around network, storage and single-threaded programs—must be overcome.

The panel, which includes leading experts from Titan, Blue Waters, Adaptive Computing, SchedMD, Intel and Cray, will discuss how refining architectural techniques will provide the necessary foundation to achieve exascale computing. Panelists will share how to conquer the new frontier in supercomputing technology with lessons learned from operating some of the world's most powerful and fastest supercomputers - Titan and Blue Waters.

The Importance of Metadata in a World Where HPC is Everywhere

Chair: Anne C. Elster (Norwegian University of Science & Technology/University of Texas at Austin)
8:30am-10am
Room: 405/406/407

Moderator: Tracey D. Wilson (CSC)
Panelists: Daniel Q. Duffy (NASA), Bernie Siebers (NOAA), Henry Newman (Instrumental, Inc), Keith R. Michaels (Boeing), Reid Bingham (General Atomics)

HPC is everywhere. The fact of high-performance or high end computing expanding to other markets with small to medium sized clusters of systems has long been foreseen. But the expansion brings about another issue: Data is everywhere. The data could be stored locally with the systems, stored with the user, distributed amongst several sites in small or large federations, or stored in a cloud. To the user, data is the real golden egg and the system the goose they must use to acquire or manipulate it. As we generate more and more data, data management becomes increasingly difficult.

This panel will discuss key issues with the expansion of HPC as it relates to metadata for data provenance, data mining and introspection for users, and what metadata and metadata rules organizations should use when dealing with cloud environments.

Big Computing: From the Exa-Scale to the Sensor-Scale

Chair: Alice Koniges (Lawrence Berkeley National Laboratory)
10:30am-12pm
Room: 405/406/407

Moderator: Frederica Darema (AFSOR)
Panelists: James Kahle (IBM Research), Sangtae Kim (University of Wisconsin at Madison), Robert Lucas (University of Southern California), Burton Smith (Microsoft Corporation), Kathy Yelick (University of California, Berkeley)

The panel will discuss the potential for new opportunities through a vision of additional dimensions to exascale and Big Data, namely those driven by the emerging ubiquitous instrumentation of systems through multitudes of distributed and heterogeneous collections of sensor- and controller-networks - the next wave of Big Data and Big Computing. Thus far, conquering the exascale has been considered as having unique challenges for power efficiency at the multicore level, dynamic management of multitudes of such resources, optimized performance, fault tolerance and resilience, and new application algorithms. Ubiquitous instrumentation environments have corresponding requirements for power efficiencies, fault tolerance, etc. Potentially, similar kinds of multicores will be the building blocks of Exascale platforms and of sensors and controllers. Leveraging common technologies is also driven by trends in large application systems. Such considerations entice questions on what new opportunities are created through synergistic advances at the two extreme computing and data scales.

Investing in Higher Education and Software for the Future

Chair: Anne C. Elster (Norwegian University of Science & Technology/University of Texas at Austin)
10:30am-12pm
Room: 401/402/403

Moderator: David Kahaner (ATIP (Asian Technology Information Program))

Panelists: Depei Qian (Beihang University,China), N. Balakrishnan (Indian Institute of Science), Kum Won Cho (National Institute of Supercomputing & Networking, South Korea), Rob Cook (Queensland Cyber Infrastructure Foundation, Australia) Taisuke Boku (University of Tsukuba, Japan)

As computer systems are becoming more complex, so is the software that will be needed for these systems. This panel will discuss how several countries have invested in higher education and training for our next generation. China is projected to produce more than 50% of the world's IT students, and the country has founded several recent tech-oriented universities and invested heavily in HPC. India is known for its IITs, and Saudi Arabia has invested heavily in KAUST. How should the rest of us keep up.

Awards

Each year, SC showcases not only the best and brightest stars of high performance computing, but also its rising stars and those who have made a lasting impression. These awards are one way these people are recognized at SC. New this year are the “Test of Time” award—recognizing a paper from a past conference that has greatly influenced the HPC discipline—and the “ACM Athena Lecturer” award—celebrating women researchers who have made fundamental contributions to computer science.

Awards

Tuesday, November 19

ACM Gordon Bell Prize Finalists I

Chair: Michael L. Norman (San Diego Supercomputer Center)
10:30am-12pm
Room: 201/203

Taking a Quantum Leap in Time to Solution for Simulations of High-Tc Superconductors

Peter Staar (ETH Zurich), Thomas A. Maier (Oak Ridge National Laboratory), Raffaele Solca (ETH Zurich), Gilles Fourestey (ETH Zurich), Michael Summers (Oak Ridge National Laboratory), Thomas C. Schulthess (ETH Zurich)

We present a new quantum cluster algorithm to simulate models of high-Tc superconductors. This algorithm extends current methods with continuous lattice self-energies, thereby removing artificial long-range correlations. This cures the fermionic sign problem in the underlying quantum Monte Carlo solver for large clusters and realistic values of the Coulomb interaction in the entire temperature range of interest. We find that the new algorithm improves time-to-solution by nine orders of magnitude compared to current, state of the art quantum cluster simulations. An efficient implementation is given, which ports to multi-core as well as hybrid CPU-GPU systems. Running on 18,600 nodes on ORNL's Titan supercomputer enables us to compute a converged value of T_c/t=0.053+-0.0014 for a 28 site cluster in the 2D Hubbard model with U/t=7 at 10% hole doping. Typical simulations on Titan sustain between 9.2 and 15.4 petaflops (double precision measured over full run), depending on configuration and parameters used.

20 Petaflops Simulation of Protein Suspensions in Crowding Conditions

Massimo Bernaschi , Mauro Bisson (Italian Council of Research (C.N.R.)), Massimiliano Fatica (NVIDIA Corporation), Simone Melchionna (Italian Council of Research (C.N.R.))

We present performance results for the simulation of protein suspensions in crowding conditions obtained with MUPHY, a computational platform for multi-scale simulations of real-life biofluidic problems. Previous versions of MUPHY have been used in the past for the simulation of blood flow through the human coronary arteries and DNA translocation across nanopores. The simulation exhibits excellent scalability up to 18,000 K20X Nvidia GPUs and achieves almost 20 Petaflops of aggregate sustained performance with a peak performance of 27.5 Petaflops for the most intensive computing component. Those figures demonstrate once again the flexibility of MUPHY in simulating biofluidic phenomena exploiting at their best the features of the architecture in use (preliminary results were obtained in the present case on a completely different platform, the IBM Blue Gene/Q). The combination of novel mathematical models, computational algorithms, hardware technology, code tuning and parallelization techniques required to achieve these results are presented.

11 PFLOP/s Simulations of Cloud Cavitation Collapse

Diego Rossinelli, Babak Hejazialhosseini , Panagiotis Hadjidoukas (ETH Zurich), Costas Bekas, Alessandro Curioni (IBM Zurich Research Laboratory), Adam Bertsch, Scott Futral (Lawrence Livermore National Laboratory), Steffen Schmidt, Nikolaus Adams (Technical University Munich), Petros Koumoutsakos (ETH Zurich)

We present unprecedented, high throughput simulations of cloud cavitation collapse on 1.6 million cores of Sequoia reaching 55% of its nominal peak performance, corresponding to 11 PFLOP/s. The destructive power of cavitation reduces the lifetime of energy critical systems such as internal combustion engines and hydraulic turbines, yet it has been harnessed for water purification and kidney lithotripsy. The present two-phase flow simulations enable the quantitative prediction of cavitation using 13 trillion grid points to resolve the collapse of 15'000 bubbles. We advance by one order of magnitude the current state-of-the-art in terms of time to solution, and by two orders the geometrical complexity of the flow. The software successfully addresses the challenges that hinder the effective solution of complex flows on contemporary supercomputers, such as limited memory bandwidth, I/O bandwidth and storage capacity. The present work redefines the frontier of high performance computing for fluid dynamics simulations.

Cray/Fernbach/Kennedy Award Recipients Presentations

Chair: Daniel A. Reed (University of Iowa)
1:30pm-3pm
Room: Mile High

The Seymour Cray Computer Science and Engineering Award recognizes innovative contributions to HPC systems that best exemplify the creative spirit of the late Seymour Cray. It includes a $10,000 honorarium and is sponsored by the IEEE Computer Society. The awards session will be chaired by Horst Simon and Larry Smarr, chairs of the selection committees.

The Sidney Fernbach Memorial Award honors innovative uses of HPC in problem solving, recognizing the pioneering contributions of the late Sidney Fernbach. It includes a $2,000 honorarium and is sponsored by the IEEE Computer Society. The awards session will be chaired by Horst Simon and Larry Smarr, chairs of the selection committees.

The Ken Kennedy Award recognizes substantial contributions to programmability and productivity in computing and substantial community service or mentoring contributions. The award honors the remarkable research, service, and mentoring contributions of the late Ken Kennedy. It includes a $5,000 honorarium, and is co-sponsored by the ACM and IEEE Computer Society. The awards session will be chaired by Horst Simon and Larry Smarr, chairs of the selection committees.

Wednesday, November 20

ACM Gordon Bell Prize Finalists II

Chair: Taisuke Boku (University of Tsukuba)
10:30am-12pm
Room: 201/203

The Origin of Mass

Peter Boyle (University of Edinburgh), Michael I. Buchoff (Lawrence Livermore National Laboratory), Norman Christ (Columbia University), Taku Izubuchi, Chulwoo Jung (Brookhaven National Laboratory), Thomas C. Luu (Lawrence Livermore National Laboratory), Robert Mawhinney (Columbia University), Chris Schroeder, Ron Soltz , Pavlos Vranas, Joseph Wasem (Lawrence Livermore National Laboratory), Zhongjie Lin, Hantao Yin (Columbia University)

The origin of mass is one of the deepest mysteries in science. Neutrons and protons, which account for almost all visible mass in the Universe, emerged from a primordial plasma through a cataclysmic phase transition microseconds after the Big Bang. However, most mass in the Universe is invisible. The existence of dark matter, which interacts with our world so weakly that it is essentially undetectable, has been established from its galactic-scale gravitational effects. Here we describe results from the first truly physical calculations of the cosmic phase transition and a groundbreaking first-principles investigation into composite dark matter, studies impossible with previous state-of-the-art methods and resources. By inventing a powerful new algorithm, ``DSDR,'' and implementing it effectively for contemporary supercomputers, we attain excellent strong scaling, perfect weak scaling to the LLNL BlueGene/Q two million cores, sustained speed of 7.2 petaflops, and time-to-solution speedup of more than 200 over the previous state-of-the-art.

Radiative Signatures of the Relativistic Kelvin-Helmholtz Instability

Michael Bussmann, Heiko Burau, Thomas E. Cowan, Alexander Debus, Axel Huebl (Helmholtz-Zentrum Dresden-Rossendorf), Guido Juckeland (Technical University Dresden), Thomas Kluge (Helmholtz-Zentrum Dresden-Rossendorf), Wolfgang E. Nagel (Technical University Dresden), Richard Pausch (Helmholtz-Zentrum Dresden-Rossendorf), Felix Schmitt (Technical University Dresden), Ulrich Schramm (Helmholtz-Zentrum Dresden-Rossendorf), Joseph Schuchart (Oak Ridge National Laboratory), Rene Widera (Helmholtz-Zentrum Dresden-Rossendorf)

We present a particle-in-cell simulation of the relativistic Kelvin-Helmholtz Instability (KHI) that for the first time delivers angularly resolved radiation spectra of the particle dynamics during the formation of the KHI. This enables studying the formation of the KHI with unprecedented spatial, angular and spectral resolution. Our results are of great importance for understanding astrophysical jet formation and comparable plasma phenomena by relating the particle motion observed in the KHI to its radiation signature. The innovative methods presented here on the implementation of the particle-in-cell algorithm on graphic processing units can be directly adapted to any many-core parallelization of the particle-mesh method. With these methods we see a peak performance of 7.176 PFLOP/s (double-precision) plus 1.449 PFLOP/s (single-precision), an efficiency of 96% when weakly scaling from 1 to 18432 nodes, an efficiency of 68.92% and a speed up of 794 (ideal: 1152) when strongly scaling from 16 to 18432 nodes.

HACC: Extreme Scaling and Performance Across Diverse Architectures

Salman Habib , Vitali A. Morozov, Nicholas Frontiere, Hal Finkel, Adrian Pope, Katrin Heitmann, Kalyan Kumaran, Venkat Vishwanath, Tom Peterka, Joseph A. Insley (Argonne National Laboratory), David Daniel, Patricia Fasel (Los Alamos National Laboratory), Zarija Lukic (Lawrence Berkeley National Laboratory)

Supercomputing is evolving towards hybrid and accelerator-based architectures with millions of cores. The HACC (Hardware/Hybrid Accelerated Cosmology Code) framework exploits this diverse landscape at the largest scales of problem size, obtaining high scalability and sustained performance. Developed to satisfy the science requirements of cosmological surveys, HACC melds particle and grid methods using a novel algorithmic structure that flexibly maps across architectures, including CPU/GPU, multi/many-core, and Blue Gene systems. We demonstrate the success of HACC on two very different machines, the CPU/GPU system Titan and the BG/Q systems Sequoia and Mira, attaining unprecedented levels of scalable performance. We demonstrate strong and weak scaling on Titan, obtaining up to 99.2% parallel efficiency, evolving 1.1

trillion particles. On Sequoia, we reach 13.94 PFlops (69.2% of peak) and 90% parallel efficiency on 1,572,864 cores, with 3.6 trillion particles, the largest cosmological benchmark yet performed. HACC design concepts are applicable to several other supercomputer applications.

Thursday, November 21

SC13 Conference Awards Presentations

Chair: Daniel A. Reed (University of Iowa)
12:30pm-1:30pm
Room: Mile High

Daniel Reed (University of Iowa), Dennis Gannon (Microsoft Corporation)

The awards managed by the SC13 conference, as well as selected ACM and IEEE awards, will be presented. These include the ACM Gordon Bell Prize, ACM-W Athena Lecturer Award, and the IEEE Technical Committee on Scalable Computing (TCSC) Young Investigators in Scalable Computing Awards; the SC13 Best Paper, Best Student Paper, and Best Poster Awards; George Michael Memorial HPC Ph.D. Fellowship; ACM Student Research Competition; and Student Cluster Competition. The SC13 Test of Time paper awardee will also be recognized.

ACM Athena Lecturer Award
SC13 Test of Time Award Special Lectures

Chair: Franck Cappello (Argonne National Laboratory)
1:30pm-3pm
Room: Mile High

The ACM Athena Lecturer Award celebrates women researchers who have made fundamental contributions to Computer Science. Sponsored by the ACM, the award includes a $10000 honorarium. This year's ACM Athena Lecturer Award winner is Katherine Yelick, Professor of Electrical Engineering and Computer Sciences, University of California, Berkeley and Associate Lab Director for Computing Sciences, Lawrence Berkeley National Laboratory.

The SC Test of Time Award recognizes a seminal technical paper from past SC conferences that has transformed high performance computing, storage, or networking. The inaugural winner of the SC Test of Time Award is William Pugh, emeritus Professor of Computer Science at the University of Maryland at College Park.

Papers

The SC13 Technical Papers program received 457 submissions covering a wide variety of research topics in high performance computing. We followed a rigorous peer review process with an author rebuttal period, careful management of conflicts, and four reviews per submission (in most cases). At a two-day face-to-face committee meeting on June 24-25 in Denver, over 230 technical paper committee members discussed every paper and finalized the selections. At the conclusion of the meeting, the committee had accepted 90 papers, reflecting an acceptance rate of 20 percent. Additionally, 13 of the 90 accepted papers have been selected as finalists for the Best Paper and Best Student Paper awards. Winners will be announced during the SC13 Awards Session on Thursday, November 21, 2013.

Papers

Papers

Tuesday, November 19

Data Management in the Cloud

Chair: Erwin Laure (KTH Royal Institute of Technology)
10:30am-12pm
Room: 205/207

Efficient Data Partitioning Model for Heterogeneous Graphs in the Cloud

Kisung Lee, Ling Liu (Georgia Institute of Technology)

As the size and variety of information networks continue to grow in many scientific domains, we witness a growing demand for efficient processing of large heterogeneous graphs using a cluster of compute nodes in the Cloud. One of the main open issues is how to effectively partition a large graph to process complex graph operations efficiently. In this paper, we present a distributed data partitioning framework for efficient processing of large-scale graphs in the Cloud. First, we introduce extended vertex blocks as graph partitioning building blocks. Second, we propose vertex block grouping algorithms which group those vertex blocks that have high correlation in graph structure to the same partition. Third, we propose a partition-guided query partitioning model which transforms graph queries into vertex block-based graph query patterns for parallel processing of graph queries. We conduct extensive experiments on several real-world graphs to show the effectiveness and scalability of our framework.

SDQuery DSI: Integrating Data Management Support with a Wide Area Data Transfer Protocol

Yu Su (Ohio State University), Yi Wang (Ohio State University), Gagan Agrawal (Ohio State University), Rajkumar Kettimuthu (Argonne National Laboratory)

In many science areas where datasets need to be transferred or shared, rapid growth in dataset size, coupled with lack of increase in wide area data transfer bandwidth, is making it extremely hard for scientists to analyze the data. This paper addresses the current limitations by developing SDQuery DSI, a GridFTP plug-in which supports flexible server-side data subsetting over HDF5 and NetCDF data formats. The GridFTP server is able to dynamically load this tool to download any data subset. Different queries types (query over dimensions, coordinates and values) are supported by our tool. A number of optimizations for improving indexing (parallel indexing), data subsetting (performance model) and data transfer (parallel streaming) are also applied. We have extensively evaluated our implementation. We compared our GridFTP SDQuery DSI with GridFTP default File DSI and showed that in different network environments, our method can achieve better efficiency in almost all cases.

Design and Performance Evaluation of NUMA-Aware RDMA-Based End-to-End Data Transfer Systems

Yufei Ren (Stony Brook University), Tan Li (Stony Brook University), Dantong Yu (Brookhaven National Laboratory), Shudong Jin (Stony Brook University), Thomas Robertazzi (Stony Brook University)

Data-intensive applications place stringent requirements on the performance of both back-end storage systems and front-end network interfaces. However, for transferring data at ultra high-speed, for example, at 100 Gbps and higher, the effects of multiple bottlenecks along a full end-to-end path have not been resolved efficiently. In this paper, we describe our implementation of end-to-end data transfer software at such high-speeds. At the back-end, we construct a storage-area network with the iSCSI protocols and utilize efficient RDMA technology. At the front-end, we design network communication software to transfer data in parallel, and utilize NUMA techniques to maximize the performance of multiple network interfaces. We demonstrate that our system can deliver the full 100 Gbps end-to-end data transfer throughput. The software product is tested rigorously and is demonstrated to be applicable to supporting various data-intensive applications that constantly move bulk data within and across data centers.

Graph Partitioning and Data Clustering

Chair: Alex Pothen (Purdue University)
10:30am-12pm
Room: 401/402/403

Scalable Parallel OPTICS Data Clustering Using Graph Algorithmic Techniques

Md. Mostofa Ali Patwary, Diana Palsetia, Ankit Agrawal (Northwestern University), Wei-keng Liao (Northwestern University), Fredrik Manne (University of Bergen), Alok Choudhary (Northwestern University)

OPTICS is a hierarchical density-based data clustering algorithm that discovers arbitrarily-shaped clusters and eliminates noise using adjustable reachability distance thresholds. Parallelizing OPTICS is challenging as the algorithm exhibits a strongly sequential data access order. We present a scalable parallel OPTICS algorithm (POPTICS) designed using graph algorithmic concepts. To break the data access sequentiality, POPTICS exploits the similarities between the OPTICS algorithm and Prim's Minimum Spanning Tree algorithm. Additionally, we use the disjoint-set data structure to achieve a high parallelism for distributed cluster extraction. Using high dimensional datasets containing up to a billion floating point numbers, we show scalable speedups of up to 27.5 for our OpenMP implementation on a 40-core shared-memory machine, and up to 3,008 for our MPI implementation on a 4,096-core distributed-memory machine. We also show that the quality of the results given by POPTICS are comparable to those given by the classical OPTICS algorithm.

Scalable Matrix Computations on Large Scale-Free Graphs Using 2D Graph Partitioning

Erik G. Boman, Karen D. Devine, Sivasankaran Rajamanickam (Sandia National Laboratories)

Scalable parallel computing is essential for processing large scale-free (power-law) graphs. The data distribution becomes important on distributed-memory computers with thousands of cores. Recently, it has been shown that 2D layouts (edge partitions) have significant advantages over traditional 1D layouts. However, the simple 2D block distribution does not use the structure of the graph, and more advanced 2D partitioning methods are too expensive for large graphs. We propose a new partitioning algorithm that combines graph partitioning with the 2D block distribution. The cost is essentially the same as 1D graph partitioning. We study the performance of sparse matrix-vector multiplication for large scale-free graphs from, e.g., social networks using several partitioners and data layouts, both 1D and 2D. We demonstrate that our new 2D method consistently outperforms the other methods considered, both for SpMV and an eigensolver, on matrices up to 1.6 billion non-zeros and up to 16,384 cores.

Scalable Parallel Graph Partitioning

Shad Kirmani, Padma Raghavan (Pennsylvania State University)

We consider partitioning a graph in parallel using a large number of processors. Parallel multilevel partitioners, such as Pt-Scotch and ParMetis, produce good quality partitions but their performance scales poorly. Coordinate bisection schemes such as those in Zoltan, which can be applied only to graphs with coordinates, scale well but partition quality is often compromised. We seek to address this gap by developing a scalable parallel scheme which imparts coordinates to a graph through a lattice-based multilevel embedding. Partitions are computed with a parallel formulation of a geometric scheme that has been shown to provide provably good cuts on certain classes of graphs. We analyze the parallel complexity of our scheme and we observe speed-ups and cut-sizes on large graphs. Our results indicate that our method is substantially faster than ParMetis and Pt-Scotch for hundreds to thousands of processors, while producing high quality cuts.
Award: Best Student Paper Finalists

Inter-Node Communication

Chair: Rich Graham (Mellanox Technologies)
10:30am-12pm
Room: 405/406/407

Channel Reservation Protocol for Over-Subscribed Channels and Destinations

George Michelogiannakis (Stanford University and Lawrence Berkeley National Laboratory), Nan Jiang (Stanford University), Daniel U. Becker (Stanford University), William J. Dally (Stanford University and NVIDIA Corporation)

Channels in system-wide networks tend to be over-subscribed due to the cost of bandwidth and increasing traffic demands. To make matters worse, workloads can overstress specific destinations, creating hotspots. Lossless networks offer attractive advantages compared to lossy networks but suffer from tree saturation. This led to the development of explicit congestion notification (ECN). However, ECN is very sensitive to its configuration parameters and acts only after congestion forms. We propose channel reservation protocol (CRP) to enable sources to reserve bandwidth in multiple resources in advance of packet transmission and with a single request, but without idling resources like circuit switching. CRP prevents congestion from ever occurring and thus reacts instantly to traffic changes, whereas ECN requires 300,000 cycles to stabilize in our experiments. Furthermore, ECN may not prevent congestion formed by short-lived flows generated by a large combination of source–destination pairs.

Enabling Highly-Scalable Remote Memory Access Programming with MPI-3 One Sided

Robert Gerstenberger, Maciej Besta, Torsten Hoefler (ETH Zurich)

Modern interconnects offer remote direct memory access (RDMA) features. Yet most applications rely on explicit message passing for communications albeit their unwanted overheads. The MPI-3.0 standard defines a programming interface for exploiting RDMA networks directly; however, its scalability and practicability has to be demonstrated in practice. In this work, we develop scalable bufferless protocols that implement the MPI-3.0 specification. Our protocols support scaling to millions of cores with negligible memory consumption while providing highest performance and minimal overheads. To arm programmers, we provide a spectrum of performance models for all critical functions and demonstrate the usability of our library and models with several application studies with up to half a million processes. We show that our design is comparable to, or better than UPC and Fortran Coarrays in terms of latency, bandwidth, and message rate. We also demonstrate application performance improvements with comparable programming complexity.
Awards: Best Paper Finalists, Best Student Paper Finalists

MVAPICH-PRISM: A Proxy-Based Communication Framework Using InfiniBand and SCIF for Intel MIC Clusters

Sreeram Potluri , Devendar Bureddy, Khaled Hamidouche, Akshay Venkatesh, Krishna Kandalla, Hari Subramoni, Dhabaleswar K. Panda (Ohio State University)

Xeon Phi packs up to 1TFlops of performance on a single chip while providing x86_64 compatibility. InfiniBand is currently one of the most popular interconnects for supercomputing systems. The software stack on Xeon Phi allows processes to directly access the InfiniBand HCA, providing a low latency path for internode communication. However, drawbacks in state-of-the-art chipsets like Sandy Bridge limit the bandwidth available for these transfers. We propose MVAPICH-PRISM, a novel proxy-based framework to optimize communication performance on such systems. We present several designs and evaluate them using micro-benchmarks and application kernels. Our designs improve internode MPI latency between Xeon Phi processes by up to 65% and internode bandwidth by up to five times. Our designs improve the performance of MPI_Alltoall operation by up to 65%, with 256 processes. They improve the performance of 3D Stencil communication kernel and P3DFFT library by 56% and 22% with 1,024 and 512 processes, respectively.

Load Balancing

Chair: Sriram Krishnamoorthy (Pacific Northwest National Laboratory)
1:30pm-3pm
Room: 405/406/407

A Framework for Load Balancing of Tensor Contraction Expressions via Dynamic Task Partitioning

Pai-Wei Lai, Kevin Stock, Samyam Rajbhandari (Ohio State University), Sriram Krishnamoorthy (Pacific Northwest National Laboratory), P. Sadayappan (Ohio State University)

In this paper, we introduce the Dynamic Load-balancing Tensor Contractions (DLTC), a domain-specific library for efficient task parallel execution of tensor contraction expressions, a class of computation encountered in quantum chemistry and physics. Our framework decomposes each contraction into smaller unit of tasks, represented by a novel abstraction referred to as iterators. We exploit an extra level of parallelism by having tasks across independent contractions to be executed concurrently through a dynamic load balancing runtime. We demonstrated the improved performance, scalability, as well as flexibility for the computation of tensor contraction expressions on parallel computers using examples from Coupled Cluster (CC) methods.

Load-Balanced Pipeline Parallelism

Md Kamruzzaman, Steven Swanson, Dean Tullsen (University of California, San Diego)

Accelerating a single thread in current parallel systems remains a challenging problem, because sequential threads do not naturally take advantage of the additional cores. Recent work shows that automatic extraction of pipeline parallelism is an effective way to speed up single thread execution. However, two problems remain challenging: load balancing and inter-thread communication. This work shows a new mechanism to exploit pipeline parallelism that naturally solves the load balancing and communication problems. This compiler-based technique automatically extracts the pipeline stages and executes them in a data parallel fashion, using token-based chunked synchronization to handle sequential stages. This technique provides linear speedup for several applications, and outperforms prior techniques to exploit pipeline parallelism by as much as 50%.

A Distributed Dynamic Load Balancer for Iterative Applications

Harshitha Menon, Laxmikant Kale (University of Illinois at Urbana-Champaign)

For many applications, computation load varies over time. Such applications require dynamic load balancing to improve performance. Centralized load balancing schemes, which perform the load balancing decisions at a central location, are not scalable. In contrast, fully distributed strategies are scalable but typically do not produce a balanced work distribution as they tend to consider only local information. This paper describes a fully distributed algorithm for load balancing that uses partial information about the global state of the system to perform load balancing. This algorithm, referred to as GrapevineLB, consists of two stages: global information propagation using a lightweight algorithm inspired by epidemic algorithms, and work unit transfer using a randomized algorithm. We provide analysis of the algorithm along with detailed simulation and performance comparison with other load balancing strategies. We demonstrate the effectiveness of GrapevineLB for adaptive mesh refinement and molecular dynamics on up to 131,072 cores of BlueGene/Q.
Award: Best Student Paper Finalist

Cloud Resource Management and Scheduling

Chair: Kate Keahey (Argonne National Laboratory)
1:30pm-3pm
Room: 205/207

Exploring Portfolio Scheduling for Long-Term Execution of Scientific Workloads in IaaS Clouds

Kefeng Deng, Junqiang Song, Kaijun Ren (National University of Defense Technology), Alexandru Iosup (Delft University of Technology)

Long-term execution of scientific applications often leads to dynamic workloads and varying application requirements. When the execution uses resources provisioned from IaaS clouds, and thus consumption-related payment, efficient and online scheduling algorithms must be found. Portfolio scheduling, which selects dynamically a suitable policy from a broad portfolio, may provide a solution to this problem. However, selecting online the right policy from possibly tens of alternatives remains challenging. In this work, we introduce an abstract model to explore this selection problem. Based on the model, we present a comprehensive portfolio scheduler that includes tens of provisioning and allocation policies. We propose an algorithm that can enlarge the chance of selecting the best policy in limited time, possibly online. Through trace-based simulation, we evaluate various aspects of our portfolio scheduler, and find performance improvements from 7% to 100% in comparison with the best constituent policies and high improvement for bursty workloads.

Cost-Effective Cloud HPC Resource Provisioning by Building Semi-Elastic Virtual Clusters

Shuangcheng Niu, Jidong Zhai (Tsinghua University), Xiaosong Ma (North Carolina State University), Xiongchao Tang, Wenguang Chen (Tsinghua University)

Recent studies have found cloud environments increasingly appealing for executing HPC applications, including tightly coupled parallel simulations. While public clouds offer elastic, on-demand resource provisioning and pay-as-you-go pricing, individual users setting up their on-demand virtual clusters may not be able to take full advantage of common cost-saving opportunities, such as reserved instances.

In this paper, we propose a Semi-Elastic Cluster (SEC) computing model for organizations to reserve and dynamically resize a virtual cloud-based cluster. We present a set of integrated batch scheduling plus resource scaling strategies uniquely enabled by SEC, as well as an online reserved instance provisioning algorithm based on job history. Our trace-driven simulation results show that such a model has a 61.0% cost saving than individual users acquiring and managing cloud resources without causing longer average job wait time. Meanwhile, the overhead of acquiring/maintaining shared cloud instances is shown to take only a few seconds.

Exploiting Application Dynamism and Cloud Elasticity for Continuous Dataflows

Alok Gautam Kumbhare, Yogesh Simmhan, Viktor K. Prasanna (University of Southern California)

Contemporary continuous dataflow systems use elastic scaling on clouds to handle variable data rates to meet applications' needs while attempting to maximize resource utilization. However, virtualized clouds present an added challenge due to the variability in resource performance thereby impacting the QoS. Elastic use of cloud resources and their allocation to dataflows thus need to adapt to such infrastructure dynamism. We develop the concept of ``dynamic dataflows'' that utilizes alternate tasks and allows additional control over the dataflow's cost and QoS. We formalize an optimization problem to allow trade-off between application's value against resource cost. We present two novel heuristics based on variable sized bin-packing heuristics and evaluate them through simulations with different data profiles using VM performance traces from a private cloud. The results show that the heuristics are effective in intelligently utilizing cloud elasticity to mitigate the effect of both input data rate and performance variabilities on QoS.

Energy Management

Chair: Taisuke Boku (University of Tsukuba)
1:30pm-3pm
Room: 401/402/403

A `Cool' Way of Improving the Reliability of HPC Machines

Osman Sarood, Esteban Meneses, Laxmikant Kale (University of Illinois at Urbana-Champaign)

Soaring energy consumption, accompanied by declining reliability, together loom as the biggest hurdles for the next generation of supercomputers. Reliability at exascale level could degrade to the point where failures become the norm. HPC researchers are focusing on improving existing fault tolerance protocols. Research on improving hardware reliability has also been making progress independently. In this paper, we try to bridge this gap and combine both software and hardware aspects towards improving reliability of HPC machines. Fault rates are known to double for every 10 rise in core temperature. We leverage this notion to experimentally demonstrate the potential of restraining processor temperatures and load balancing to achieve two-fold benefits: improving reliability of parallel machines and reducing total execution time required by applications. For a 350K socket machine, regular checkpoint/restart fails to make progress, whereas our validated model predicts an efficiency of 20% by improving the machine reliability by 2.29X.

Coordinated Energy Management in Heterogeneous Processors

Indrani Paul (Advanced Micro Devices, Inc. and Georgia Institute of Technology), Vignesh Ravi , Srilatha Manne (Advanced Micro Devices, Inc.), Manish Arora (Advanced Micro Devices, Inc. and University of California, San Diego), Sudhakar Yalamanchili (Georgia Institute of Technology)

This paper examines energy management in a heterogeneous processor consisting of an integrated CPU-GPU for high-performance computing (HPC) applications. Energy management for HPC applications is challenged by their uncompromising performance requirements and complicated by the need for coordinating energy management across distinct core types – a new and less understood problem.

We examine the intra-node CPU-GPU frequency sensitivity of HPC applications on tightly coupled CPU-GPU architectures as the first step in understanding power and performance optimization for a heterogeneous multi-node HPC system. The insights from this analysis form the basis of a coordinated energy management scheme, called DynaCo, for integrated CPU-GPU architectures. We implement DynaCo on a modern heterogeneous processor and compare its performance to a state-of-the-art power- and performance-management algorithm. DynaCo improves measured average energy-delay squared (ED^2) product by up to 30% with less than 2% average performance loss across several exascale and other HPC workloads.

Award: Best Paper Finalist

Integrating Dynamic Pricing of Electricity into Energy Aware Scheduling for HPC Systems

Xu Yang, Zhou Zhou, Sean Wallace, Zhiling Lan (Illinois Institute of Technology), Wei Tang, Susan Coghlan, Michael E. Papka (Argonne National Laboratory)

The research literature to date mainly aimed at reducing energy consumption in HPC environments. In this paper we propose a job power aware scheduling mechanism to reduce HPC's electricity bill without degrading the system utilization. The novelty of our job scheduling mechanism is its ability to take the variation of electricity price into consideration as a means to make better decisions of the timing of scheduling jobs with diverse power profiles. We verified the effectiveness of our design by conducting trace-based experiments on an IBM Blue Gene/P and a cluster system as well as a case study on Argonne's 48-rack IBM Blue Gene/Q system. Our preliminary results show that our power aware algorithm can reduce electricity bill of HPC systems as much as 23%.

Extreme-Scale Applications

Chair: Michael A. Heroux (Sandia National Laboratories)
1:30pm-3pm
Room: 201/203

Petascale Direct Numerical Simulation of Turbulent Channel Flow on up to 786K Cores

Myoungkyu Lee, Nicholas Malaya, Robert D. Moser (University of Texas at Austin)

We present results of performance optimization for direct numerical simulation (DNS) of wall bounded turbulent flow (channel flow). DNS is a technique in which the fluid flow equations are solved without subgrid modeling. Of particular interest are high Reynolds number(Re) turbulent flows over walls, because of their importance in technological applications. Simulating high Re turbulence is a challenging computational problem, due to the high spatial and temporal resolution requirements.

An optimized code was developed using spectral methods, the method of choice for turbulent flows. Optimization was performed to address three major issues: efficiency of banded matrix linear algebra, cache reuse and memory access, and communication for the global data transposes.

Results show that performance is highly dependent on characteristics of the communication network, rather than single-core performance. In our tests, it exhibits approximately 80% strong scaling parallel efficiency at 786K cores relative to performance on 65K cores.

Award: Best Student Paper Finalists

Solving the Compressible Navier-Stokes Equations on up to 1.97 Million Cores and 4.1 Trillion Grid Points

Ivan Bermejo-Moreno, Julien Bodart (Stanford University), Johan Larsson (University of Maryland), Blaise Barney (Lawrence Livermore National Laboratory), Joseph Nichols, Steve Jones (Stanford University)

We present weak and strong scaling studies as well as performance analyses of the Hybrid code, a finite-difference solver of the compressible Navier-Stokes equations on structured grids used for the direct numerical simulation of isotropic turbulence and its interaction with shock waves. Parallelization is achieved through MPI, emphasizing the use of non-blocking communication with concurrent computation. The simulations, scaling and performance studies were done on the Sequoia, Vulcan and Vesta Blue Gene/Q systems, the first two accounting for a total of 1,966,080 cores when combined. The maximum number of grid points simulated was 4.12 trillion, with a memory usage of approximately 1.6 PB. We discuss the use of hyperthreading, which significantly improves the parallel performance of the code on this architecture.

Petascale WRF Simulation of Hurricane Sandy: Deployment of NCSA's Cray XE6 Blue Waters

Peter Johnsen (Cray Inc.), Mark Straka (University of Illinois at Urbana-Champaign), Melvyn Shapiro, Alan Norton, Thomas Galarneau (National Center for Atmospheric Research)
The National Center for Atmospheric Research (NCAR)

Weather Research and Forecasting (WRF) model has been employed on the largest yet storm prediction model using real data of over 4 billion points to simulate the landfall of Hurricane Sandy. Using an unprecedented 13,680 nodes (437,760 cores) of the Cray XE6 "Blue Waters" at NCSA at the University of Illinois, researchers achieved a sustained rate of 285 Tflops while simulating an 18-hour forecast. A grid of size 9120x9216x48 (1.4Tbytes of input) was used, with horizontal resolution of 500 meters and a 2-second time step. 86 Gbytes of forecast data was written every 6 forecast hours at a rate of up to 2 Gbytes/second and collaboratively post-processed and displayed using the Vapor suite at NCAR. Opportunities to enhance scalability in the source code, run-time, and operating system realms were exploited. The output of this numerical model is now under study for model validation.

System-wide Application Performance Assessments

Chair: Dominik Ulmer (Cray Switzerland)
3:30pm-5pm
Room: 401/402/403

Supercomputing with Commodity CPUs: Are Mobile SoCs Ready for HPC?

Nikola Rajovic, Paul M. Carpenter, Isaac Gelado, Nikola Puzovic, Alex Ramirez, Mateo Valero (Barcelona Supercomputing Center)

In the late 1990s, powerful economic forces led to the adoption of commodity desktop processors in high-performance computing. This transformation has been so effective that the June 2013 TOP500 list is still dominated by x86.

In 2013, the largest commodity market in computing is not PCs or servers, but mobile computing, comprising smartphones and tablets, most of which are built with ARM-based SoCs. This leads to the suggestion that once mobile SoCs deliver sufficient performance, mobile SoCs can help reduce the cost of HPC.

This paper addresses this question in detail. We analyze the trend in mobile SoC performance, comparing it with the similar trend in the 1990s. We also present our experience evaluating performance and efficiency of mobile SoCs, deploying a cluster and evaluating the network and scalability of production applications. In summary, we give a first answer as to whether mobile SoCs are ready for HPC.
Award: Best Student Paper Finalist

There Goes the Neighborhood: Performance Degradation due to Nearby Jobs

Abhinav Bhatele, Kathryn Mohror , Steven H. Langer (Lawrence Livermore National Laboratory), Katherine E. Isaacs (University of California Davis)

Predictable performance is important for understanding and alleviating application performance issues; quantifying the effects of source code, compiler, or system software changes; estimating the time required for batch jobs; and determining the allocation requests for proposals. Our experiments show that on a Cray XE system, the execution time of a communication-heavy parallel application ranges from 28% faster to 41% slower than the average observed performance. Blue Gene systems, on the other hand, demonstrate no noticeable run-to-run variability. In this paper, we focus on Cray machines and investigate potential causes for performance variability such as OS jitter, shape of the allocated partition, and interference from other jobs sharing the same network links. Reducing such variability could improve overall throughput at a computer center and save energy costs.

CooMR: Cross-Task Coordination for Efficient Data Management in MapReduce Programs

Xiaobing Li, Yandong Wang, YIzheng Jiao, Cong Xu, Weikuan Yu (Auburn University)

Hadoop is a popular open-source implementation of a MapReduce programming model for big data processing. It represents system resources as map and reduce slots and schedules them to various tasks. This execution model gives little regard to the need of cross-task coordination on the use of shared resources on each node, which results in task interference. In addition, the existing merge algorithm causes excessive I/O. In this study, we undertake an effort to address both issues. Accordingly, we introduce a cross-task coordination framework called CooMR for efficient data management in MapReduce programs. CooMR consists of three component schemes including cross-task opportunistic memory sharing and log-structured I/O consolidation, which aim to facilitate task coordination, and a key-based in-situ merge algorithm designed to enable the sorting/merging of intermediate data without actually moving the pairs. Our evaluation demonstrates that CooMR can increase task coordination, improve resource utilization, and effectively accelerate MapReduce programs.

Fault Tolerance and Migration in the Cloud

Chair: Henry Tufo (University of Colorado)
3:30pm-5pm
Room: 205/207

Optimization of Cloud Task Processing with Checkpoint-Restart Mechanism

Sheng Di (INRIA and Argonne National Laboratory (USA)), Yves Robert (ENS Lyon and University of Tennessee, Knoxville), Frederic Vivien (INRIA), Derrick Kondo (INRIA), Cho-Li Wang (University of Hong Kong), Franck Cappello (INRIA, Argonne National Laboratory and University of Illinois at Urbana-Champaign)

In this paper, we aim at optimizing fault-tolerance techniques based on a checkpointing/restart mechanism, in the context of cloud computing. Our contribution is three-fold. (1) We derive a fresh formula to compute the optimal number of checkpoints for cloud jobs with varied distributions of failure events. Our analysis is not only generic with no assumption on failure probability distribution, but attractively simple to apply in practice. (2) We design an adaptive algorithm to optimize the impact of checkpointing regarding various costs like checkpointing/restart overhead. (3) We evaluate our optimized solution in a real cluster environment with hundreds of virtual machines and Berkeley Lab Checkpoint/Restart tool. Task failure events are emulated via a production trace produced on a large-scale Google data center. Experiments confirm that our solution is fairly suitable for Google systems. Our optimized formula outperforms Young's formula by 3-10 percent, reducing wall-clock-lengths by 50-100 seconds per job on average.

Scalable Virtual Machine Deployment Using VM Image Caches
Kaveh Razavi, Thilo Kielmann (VU University Amsterdam)
In IaaS clouds, VM startup times are frequently perceived as slow, negatively impacting both dynamic scaling of web applications and the startup of high-performance computing applications consisting of many VM nodes. A significant part of the startup time is due to the large transfers of VM image content from a storage node to the actual compute nodes, even when copy-on-write schemes are used. We have observed that only a tiny part of the VM image is needed for the VM to be able to start up. Based on this observation, we propose using small caches for VM images to overcome the VM startup bottlenecks. We have implemented such caches as an extension to KVM/QEMU. Our evaluation with up to 64 VMs shows that using our caches reduces the time needed for simultaneous VM startups to the one of a single VM.

Guide-Copy: Fast and Silent Migration of Virtual Machine for Datacenters

Jihun Kim, Dongju Chae, Jangwoo Kim, Jong Kim (Pohang University of Science and Technology)

Cloud infrastructure providers deploy Dynamic Resource Management (DRM) to minimize the cost of datacenter operation, while maintaining the Service Level Agreement (SLA). Such DRM schemes depend on the capability to migrate virtual machine (VM) images. However, existing migration techniques are not suitable for highly utilized clouds due to their latency and bandwidth critical memory transfer mechanisms. In this paper, we propose guide-copy migration, a novel VM migration scheme to provide a fast and silent migration, which works nicely under highly utilized clouds. The guide-copy migration transfers only the memory pages accessed at the destination node in the near future by running a guide version of the VM at the source node and a migrated VM at the destination node simultaneously during the migration. The guide-copy migration's highly accurate and low-bandwidth memory transfer mechanism enables a fast and silent VM migration to maintain the SLA of all VMs in the cloud.

I/O Tuning

Chair: Andre Brinkmann (Johannes Gutenberg-University Mainz)
3:30pm-5pm
Room: 405/406/407

Characterization and Modeling of PIDX Parallel I/O for Performance Optimization

Sidharth Kumar, Avishek Saha (University of Utah), Venkatram Vishwanath, Philip Carns (Argonne National Laboratory), John A. Schmidt (University of Utah), Robert Latham (Argonne National Laboratory), Giorgio Scorzelli (University of Utah), Hemanth Kolla (Sandia National Laboratories), Robert Ross (Argonne National Laboratory), Jackie Chen (Sandia National Laboratories), Michael E. Papka (Argonne National Laboratory), Ray Grout (National Renewable Energy Laboratory), Valerio Pascucci (University of Utah)

Parallel I/O library performance can vary greatly in response to user-tunable parameter values such as aggregator count, file count, and aggregation strategy. Unfortunately, manual selection of these values is time consuming and dependent on characteristics of the target machine, the underlying file system, and the dataset itself. Some characteristics, such as the amount of memory per core, can also impose hard constraints on the range of viable parameter values. In this work we address these problems by using machine learning techniques to model the performance of the PIDX parallel I/O library and select appropriate tunable parameter values. We characterize both the network and I/O phases of PIDX on a Cray XE6 as well as an IBM Blue Gene/P system. We use the results of this study to develop a machine learning model for parameter space exploration and performance prediction.

Taming Parallel I/O Complexity with Auto-Tuning

Babak Behzad, Huong Vu Thanh Luu (University of Illinois at Urbana-Champaign), Joseph Huchette (Rice University), Surendra Byna , Prabhat Mr. (Lawrence Berkeley National Laboratory), Ruth Aydt (HDF Group), Quincey Koziol (HDF Group), Marc Snir (University of Illinois at Urbana-Champaign)

We present an auto-tuning system for optimizing I/O performance of HDF5 applications and demonstrate its value across platforms, applications, and at scale. The system uses a genetic algorithm to search a large space of tunable parameters and to identify effective settings at all layers of the parallel I/O stack. The parameter settings are applied transparently by the auto-tuning system via dynamically intercepted HDF5 calls.

To validate our auto-tuning system, we applied it to three I/O benchmarks that replicate the I/O activity of their respective applications. We tested the system with different weak-scaling configurations that generate 30 GB to 1 TB of data, and executed these configurations on diverse HPC platforms (Cray XE6, IBM BG/P, and Dell Cluster). In all cases, the auto-tuning framework identified tunable parameters that substantially improved write performance over default system settings. We consistently demonstrate I/O write speedups between 2x and 50x for test configurations.

Toward Millions of File System IOPS on Low-Cost Commodity Hardware

Da Zheng (Johns Hopkins University), Randal Burns (Johns Hopkins University), Alexander S. Szalay (Johns Hopkins University)

We describe a storage system that removes I/O bottlenecks to achieve more than one million IOPS based on a user space file abstraction for arrays of commodity SSDs. The file abstraction refactors I/O scheduling and placement for extreme parallelism and non-uniform memory and I/O. The system includes a set-associative, parallel page cache in the user space. We redesign page caching to eliminate CPU overhead and lock-contention in non-uniform memory architecture machines. We evaluate our design on a 32 core NUMA machine with four, eight-core processors. Experiments show that our design delivers 1.23 million 512-byte read IOPS. The page cache realizes the scalable IOPS of Linux asynchronous I/O (AIO) and increases user-perceived I/O performance linearly with cache hit rates. The parallel, set-associative cache matches the cache hit rates of the global Linux page cache under real workloads.

Physical Frontiers

Chair: Richard Vuduc (Georgia Institute of Technology)
3:30pm-5pm
Room: 201/203

Physics-Based Seismic Hazard Analysis on Petascale Heterogeneous Supercomputers

Yifeng Cui (San Diego Supercomputer Center), Efecan Poyraz (University of California, San Diego), Kim B. Olsen (San Diego State University), Jun Zhou (University of California, San Diego), Kyle Withers (San Diego State University), Scott Callaghan (University of Southern California), Jeff Larkin (NVIDIA Corporation), Clark C. Guest (University of California, San Diego), Dong Ju Choi (San Diego Supercomputer Center), Amit Chourasia (San Diego Supercomputer Center), Zheqiang Shi, Steven M. Day (San Diego State University), Philip J. Maechling, Thomas H. Jordan (University of Southern California)

We have developed a highly scalable and efficient GPU-based finite-difference code (AWP) for earthquake simulation that implements high throughput, memory locality, communication reduction and communication/computation overlap and achieves linear scalability on Cray XK7 Titan at ORNL and NCSA's Blue Waters system. AWP delivered excellent performance while simulating realistic 0-10 Hz earthquake ground motions relevant to building engineering design. Moreover, we show that AWP provides a speedup by a factor of 110 in key strain tensor calculations critical to probabilistic seismic hazard analysis (PSHA). These performance improvements to critical scientific application software, coupled with improved co-scheduling capabilities of our workflow-managed systems, make a statewide hazard model a goal reachable with existing supercomputers. The performance improvements of GPU-based AWP are expected to save millions of core-hours over the next few years as physics-based seismic hazard analysis is developed using heterogeneous petascale supercomputers.

A Scalable Parallel Algorithm for Dynamic Range-Limited N-Tuple Computation in Many-Body Molecular Dynamics Simulation

Manaschai Kunaseth, Rajiv K. Kalia, Aiichiro Nakano, Ken-ichi Nomura, Priya Vashishta (University of Southern California)

Recent advancements in reactive molecular dynamics (MD) simulations based on many-body interatomic potentials necessitate efficient dynamic n-tuple computation, where a set of atomic n-tuples within a given spatial range is constructed at every time step. Here, we develop a computation-pattern algebraic framework to mathematically formulate general n-tuple computation. Based on translation/reflection-invariant properties of computation patterns within this framework, we design a shift-collapse (SC) algorithm for cell-based parallel MD. Theoretical analysis quantifies the compact n-tuple search space and small communication cost of SC-MD for arbitrary n, which are reduced to those in best pair-computation ap-

proaches (e.g. eighth-shell method) for n = 2. Benchmark tests show that SC-MD outperforms our production MD code at the finest grain, with 9.7- and 5.1-fold speedups on Intel-Xeon and BlueGene/Q clusters. SC-MD also exhibits excellent strong scalability.

2HOT: An Improved Parallel Hashed Oct-Tree N-body Algorithm for Cosmological Simulation

Michael S. Warren (Los Alamos National Laboratory)

We report on improvements made over the past two decades to our adaptive treecode N-body method (HOT). A mathematical and computational approach to the cosmological N-body problem is described, with performance and scalability measured up to 256k (2^18) processors. We present error analysis and scientific application results from a series of more than ten 69 billion (4096^3) particle cosmological simulations, accounting for 4 x 10^20 floating point operations. These results include the first simulations using the new constraints on the standard model of cosmology from the Planck satellite. Our simulations set a new standard for accuracy and scientific throughput, while meeting or exceeding the computational efficiency of the latest generation of hybrid TreePM N-body methods.

Award: Best Paper Finalist

Wednesday, November 20

Parallel Programming Models and Compilation

Chair: Ganesh Gopalakrishnan (University of Utah)
10:30am-12pm
Room: 405/406/407

Parallelizing the Execution of Sequential Scripts

Zhao Zhang, Daniel S. Katz, Timothy G. Armstrong (University of Chicago), Justin M. Wozniak, Ian Foster (Argonne National Laboratory)

Scripting is often used in science to create applications via the composition of existing programs. Parallel scripting systems allow the creation of such applications, but each system introduces the need to adopt a somewhat specialized programming model. We present an alternative scripting approach, AMFS Shell, that lets programmers express parallel scripting applications via minor extensions to existing sequential scripting languages, such as Bash, and then execute them in-memory on large-scale computers. We define a small set of commands between the scripts and a parallel scripting runtime system, so that programmers can compose their scripts in a familiar scripting language. The underlying AMFS implements both collective (fast file movement) and functional (transformation based on content) file management. Tasks are handled by AMFS's built-in execution engine. AMFS Shell is expressive enough for a wide range of applications, and the framework can run such applications efficiently on large-scale computers.

Deterministic Scale-Free Pipeline Parallelism with Hyperqueues

Hans Vandierendonck (Queen's University Belfast), Kallia Chronaki (Barcelona Supercomputing Center), Dimitrios S. Nikolopoulos (Queen's University Belfast)

Deterministic and scale- free programming models built on a task abstraction offer a viable path to ubiquitous parallelism. However, it remains hard to reconcile these attributes with pipeline parallelism, where the number of pipeline stages is typically hard-coded in the program and defines the degree of parallelism.

This paper introduces hyperqueues, a programming abstraction that enables deterministic and scale-free pipeline parallel programs. Hyperqueues extend the concept of Cilk++'s hyperobjects to provide thread-local views on a shared data structure. While hyperobjects are organized around private local views, hyperqueues require shared concurrent views on the underlying data structure. We define the semantics of hyperqueues and describe their implementation in a work-stealing scheduler. We demonstrate scalable performance on pipeline parallel PARSEC benchmarks and find that hyperqueues provide comparable or better performance than implementations that are highly tuned to the number of available processing cores using POSIX threads and Intel's Threading Building Blocks.

Compiling Affine Loop Nests for Distributed-Memory Parallel Architectures

Uday Bondhugula (Indian Institute of Science)

We present new techniques for compilation of arbitrarily nested loops with affine dependencies for distributed-memory parallel architectures. Our framework is implemented as a source-level transformer that uses the polyhedral model, and generates parallel code with communication expressed with the Message Passing Interface (MPI) library. Compared to all previous approaches, ours is a significant advance either (1) with respect to the generality of input code handled, or (2) efficiency of communication code, or both. We provide experimental results on a cluster of multicores demonstrating its effectiveness. In some cases, code we generate outperforms manually parallelized codes, and in another case is within 25% of it. To the best of our knowledge, this is the first work reporting end-to-end fully automatic distributed-memory parallelization and code generation for input programs and transformation techniques as general as those we allow.

Performance Management of HPC Systems

Chair: Sadaf R. Alam (Swiss National Supercomputing Centre)
10:30am-12pm
Room: 401/402/403

Enabling Fair Pricing on HPC Systems with Node Sharing

Alex D. Breslow (University of California, San Diego), Ananta Tiwari (San Diego Supercomputer Center), Martin Schulz (Lawrence Livermore National Laboratory), Laura Carrington (San Diego Supercomputer Center), Lingjia Tang (University of Michigan), Jason Mars (University of Michigan)

Co-location, where multiple jobs share compute nodes in large-scale HPC systems, has been shown to increase aggregate throughput and energy efficiency by 10 to 20%. However, system operators disallow co-location due to fair-pricing concerns, i.e., a pricing mechanism that considers performance interference from co-running jobs. In the current pricing model, application execution time determines the price, which results in unfair prices paid by the minority of users whose jobs suffer from co-location.

This paper presents POPPA, a runtime system that enables fair pricing by delivering precise online interference detection and facilitates the adoption of supercomputers with co-locations. POPPA leverages a novel shutter mechanism, a cyclic, fine-grained interference sampling mechanism to accurately deduce the interference between co-runners, to provide unbiased pricing of jobs that share nodes. POPPA is able to quantify inter-application interference within 4% mean absolute error on a variety of co-located benchmarks and real scientific workloads.
Award: Best Paper Finalists, Best Student Paper Finalists

ACIC: Automatic Cloud I/O Configurator for HPC Applications

Mingliang Liu (Tsinghua University), Ye Jin (North Carolina State University), Jidong Zhai (Tsinghua University), Yan Zhai (University of Wisconsin-Madison), Qianqian Shi (Tsinghua University), Xiaosong Ma (North Carolina State University), Wenguang Chen (Tsinghua University)

The cloud has become a promising alternative to traditional HPC centers or in-house clusters. This new environment highlights the I/O bottleneck problem, typically with top-of-the-line compute instances but sub-par communication and I/O facilities. It has been observed that changing cloud I/O system configurations leads to significant variation in the performance and cost efficiency of I/O intensive HPC applications. However, storage system configuration is tedious and error-prone to do manually, even for experts.

This paper proposes ACIC, which takes a given application running on a given cloud platform, and automatically searches for optimized I/O system configurations. ACIC utilizes machine learning models to perform black-box performance/cost predictions. To tackle the high-dimensional parameter exploration space unique to cloud platforms, we enable affordable, reusable, and incremental training guided by Plackett and Burman Matrices. Results with four representative applications indicate that ACIC consistently identifies near-optimal configurations among a large group of candidate settings.

COCA: Online Distributed Resource Management for Cost Minimization and Carbon Neutrality in Data Centers

Shaolei Ren (Florida International University), Yuxiong He (Microsoft Research)

Due to the enormous energy consumption and associated environmental concerns, data centers have been increasingly pressured to reduce long-term net carbon footprint to zero, i.e., carbon neutrality. In this paper, we propose an online algorithm, called COCA (optimizing for COst minimization and CArbon neutrality), for minimizing data center operational cost while satisfying carbon neutrality without long-term future information. Unlike the existing research, COCA enables distributed server-level resource management: each server autonomously adjusts its processing speed and optimally decides the amount of workloads to process. We prove that COCA achieves a close-to-minimum operational cost (incorporating both electricity and delay costs) compared to the optimal algorithm with future information, while bounding the potential violation of carbon neutrality. We also perform trace-based simulation studies to complement the analysis, and the results show that COCA reduces cost by more than 25% (compared to state of the art) while resulting in a smaller carbon footprint.

Optimizing Data Movement

Chair: Dhabaleswar K. (DK) Panda (Ohio State University)
10:30am-12pm
Room: 205/207

SIDR: Structure-Aware Intelligent Data Routing in Hadoop

Joe Buck, Noah Watkins, Greg Levin, Adam Crume, Kleoni Ioannidou, Scott Brandt, Carlos Maltzahn, Neoklis Polyzotis, Aaron Torres (Los Alamos National Laboratory)

The MapReduce framework is being extended for domains quite different from the web applications for which it was designed, including the processing of big structured data, e.g., scientific and financial data. Previous work using MapReduce to process scientific data ignores existing structure when assigning intermediate data and scheduling tasks. In this paper, we present a method for incorporating knowledge of the structure of scientific data and executing query into the MapReduce communication model. Built in SciHadoop, a version of the Hadoop MapReduce framework for scientific data, SIDR intelligently partitions and routes intermediate data, allowing it to: remove Hadoop's global barrier and execute Reduce tasks prior to all Map tasks completing; minimize intermediate key skew; and produce early, correct results. SIDR executes

queries up to 2.5 times faster than Hadoop and 37% faster than SciHadoop; produces initial results with only 6% of the query completed; and produces dense, contiguous output.

Using Cross-Layer Adaptations for Dynamic Data Management in Large Scale Coupled Scientific Workflows

Tong Jin, Fan Zhang, Qian Sun, Hoang Bui, Manish Parashar (Rutgers University), Hongfeng Yu (University of Nebraska-Lincoln), Scott Klasky, Norbert Podhorszki, Hasan Abbasi (Oak Ridge National Laboratory)

As system scales and application complexity grow, managing and processing simulation data has become a significant challenge. While recent approaches based on data staging and in-situ/in-transit data processing are promising, dynamic data volumes and distributions,such as those occurring in AMR-based simulations, make the efficient use of these techniques challenging. In this paper we propose cross-layer adaptations that address these challenges and respond at runtime to dynamic data management requirements. Specifically we explore (1) adaptations of the spatial resolution at which the data is processed, (2) dynamic placement and scheduling of data processing kernels, and (3) dynamic allocation of in-transit resources. We also exploit coordinated approaches that dynamically combine these adaptations at the different layers. We evaluate the performance of our adaptive cross-layer management approach on the Intrepid IBM-BlueGene/P and Titan Cray-XK7 systems using Chombo-based AMR applications, and demonstrate its effectiveness in improving overall time-to-solution and increasing resource efficiency.

Exploring the Future of Out-Of-Core Computing with Compute-Local Non-Volatile Memory

Myoungsoo Jung (University of Texas at Dallas), Ellis H. Wilson III, Wonil Choi (Pennsylvania State University), John Shalf, Hasan Metin Aktulga, Chao Yang (Lawrence Berkeley National Laboratory), Erik Saule, Umit V. Catalyurek (Ohio State University), Mahmut Kandemir (Pennsylvania State University)

Drawing parallels to the rise of general purpose graphical processing units (GPGPUs) as accelerators for specific high-performance computing (HPC) workloads, the HPC community is also witness to a recent rise in the use of non-volatile memory (NVM), such as flash-based solid-state drives (SSDs), as accelerators for certain I/O-intensive workloads. These NVM acceleration architectures have been shown to work particularly well for read-intensive scientific applications such as out-of-core numerical linear algebra, where matrices too large for individual nodes' main memory tend to be calculated once and read many times. However, existing works have explored use of NVM within dedicated I/O nodes, which are distant from the compute nodes that actually need such acceleration. As NVM bandwidth begins to out-pace point-to-point network capacity, we argue for the need to break from the archetype of completely separated storage and consider compute-local solutions involving NVM.

Award: Best Paper Finalists, Best Student Paper Finalists

Fault-Tolerant Computing

Chair: Pavan Balaji (Argonne National Laboratory)
1:30pm-3pm
Room: 201/203

ACR: Automatic Checkpoint/Restart for Soft and Hard Error Protection

Xiang Ni, Esteban Meneses , Nikhil Jain, Laxmikant Kale (University of Illinois at Urbana-Champaign)

As machines increase in scale, many researchers have predicted that failure rates will correspondingly increase. Soft errors do not inhibit execution, but may silently generate incorrect results. Recent trends have shown that soft error rates are increasing, and hence they must be detected and handled to maintain correctness. We present a holistic methodology for automatically detecting and recovering from soft or hard faults with minimal application intervention. This is demonstrated by ACR: an automatic checkpoint/restart framework that performs application replication and automatically adapts the checkpoint period using online information about the current failure rate. ACR performs an application- and user-oblivious recovery. We empirically test ACR by injecting failures that follow different distributions for five applications and show low overhead when scaled to 131,072 cores. We also analyze the interaction between soft and hard errors and propose three recovery schemes that explore the trade-off between performance and reliability requirements.

SPBC: Leveraging the Characteristics of MPI HPC Applications for Scalable Checkpointing

Thomas Ropars (Swiss Federal Institute of Technology Lausanne), Tatiana Martsinkevich (INRIA and University of Paris Sud), Amina Guermouche (Université de Versailles Saint-Quentin en Yveline), André Schiper (Swiss Federal Institute of Technology Lausanne), Franck Cappello (Argonne National Laboratory)

The high failure rate expected for future supercomputers requires the design of new fault tolerant solutions. Most checkpointing protocols are designed to work with any message-passing application but suffer from scalability issues at extreme scale. We take a different approach: We identify a property common to many HPC applications, namely channel-determinism, and introduce a new partial order relation, called always-happens-before relation, between events of such applications. Leveraging these two concepts, we design a protocol that combines an unprecedented set of features. Our protocol called SPBC combines in a hierarchical way coordinated checkpointing and message logging. It is the first protocol that provides failure containment without logging any information reliably apart from process checkpoints, and this, without penalizing recovery performance. Experiments run with a representative set of HPC workloads demonstrate a good performance of our protocol during both, failure-free execution and recovery.

Using Simulation to Explore Distributed Key-Value Stores for Extreme-Scale System Services

Ke Wang (Illinois Institute of Technology), Abhishek Kulkarni (Indiana University), Michael Lang (Los Alamos National Laboratory), Dorian Arnold (University of New Mexico), Ioan Raicu (Illinois Institute of Technology)

Owing to the significant high rate of component failures at extreme scales, system services will need to be failure-resistant, adaptive and self-healing. A majority of HPC services are still designed around a centralized paradigm and hence are susceptible to scaling issues. Peer-to-peer services have proved themselves at scale for wide-area internet workloads. Distributed key-value stores (KVS) are widely used as a building block for these services, but are not prevalent in HPC services. In this paper, we simulate KVS for various service architectures and examine the design trade-offs as applied to HPC service workloads to support extreme-scale systems. The simulator is validated against existing distributed KVS-based services. Via simulation, we demonstrate how failure, replication, and consistency models affect performance at scale. Finally, we emphasize the general use of KVS to HPC services by feeding real HPC service workloads into the simulator and presenting a KVS-based distributed job launch prototype.

Parallel Performance Tools

Chair: Shirley Moore (University of Texas El Paso)
1:30pm-3pm
Room: 405/406/407

A Data-Centric Profiler for Parallel Programs

Xu Liu, John Mellor-Crummey (Rice University)

It is difficult to manually identify opportunities for enhancing data locality. To address this problem, we extended the HPCToolkit performance tools to support data-centric profiling of scalable parallel programs. Our tool uses hardware counters to directly measure memory access latency and attributes latency metrics to both variables and instructions. Different hardware counters provide insight into different aspects of data locality (or lack thereof). Unlike prior tools for data-centric analysis, our tool employs scalable measurement, analysis, and presentation methods that enable it to analyze the memory access behavior of scalable parallel programs with low runtime and space overhead. We demonstrate the utility of HPCToolkit's new data-centric analysis capabilities with case studies of five well-known benchmarks. In each benchmark, we identify performance bottlenecks caused by poor data locality and demonstrate non-trivial performance optimizations enabled by this guidance.

On the Usefulness of Object Tracking Techniques in Performance Analysis

German Llort, Harald Servat, Juan Gonzalez, Judit Gimenez, Jesus Labarta (Barcelona Supercomputing Center)

Understanding the behavior of a parallel application is crucial if we are to tune it to achieve its maximum performance. Yet the behavior the application exhibits may change over time and depend on the actual execution scenario: particular inputs and program settings, the number of processes used, or hardware-specific problems. So beyond the details of a single experiment a far more interesting question arises: how does the application behavior respond to changes in the execution conditions?

In this paper, we demonstrate that object tracking concepts from computer vision have huge potential to be applied in the context of performance analysis. We leverage tracking techniques to analyze how the behavior of a parallel application evolves through multiple scenarios where the execution conditions change. This method provides comprehensible insights on the influence of different parameters on the application behavior, enabling us to identify the most relevant code regions and their performance trends.

Detection of False Sharing Using Machine Learning

Sanath Jayasena (University of Moratuwa), Saman Amarasinghe (Massachusetts Institute of Technology), Asanka Abeyweera, Gayashan Amarasinghe, Himeshi De Silva, Sunimal Rathnayake (University of Moratuwa), Xiaoqiao Meng, Yanbin Liu (IBM Corporation)

False sharing is a major class of performance bugs in parallel applications. Detecting false sharing is difficult as it does not change the program semantics. We introduce an efficient and effective approach for detecting false sharing based on machine learning.

We develop a set of mini-programs in which false sharing can be turned on and off. We then run the mini-programs both with and without false sharing, collect a set of hardware performance event counts and use the collected data to train a classifier. We can use the trained classifier to analyze data from arbitrary programs for detection of false sharing.

Experiments with the PARSEC and Phoenix benchmarks show that our approach is indeed effective. We detect published false sharing regions in the benchmarks with zero false positives. Our performance penalty is less than 2%. Thus, we believe that this is an effective and practical method for detecting false sharing.

In-Situ Data Analytics and Reduction

Chair: Carlos Maltzahn (University of California, Santa Cruz)
1:30pm-3pm
Room: 205/207

Assessing the Effects of Data Compression in Simulations Using Physically Motivated Metrics

Daniel E. Laney, Steven H. Langer, Christopher R. Weber, Peter G. Lindstrom (Lawrence Livermore National Laboratory), Al Wegener (Samplify Systems)

This paper examines whether lossy compression can be used effectively in physics simulations as a possible strategy to combat the expected data-movement bottleneck in future high performance computing architectures. We show that, for the codes and simulations we tested, compression levels of 3--5X can be applied without causing significant changes to important physical quantities.

Rather than applying signal processing error metrics, we utilize physics-based metrics appropriate for each code to assess the impact of compression. We evaluate three different simulation codes: a Lagrangian shock-hydrodynamics code, an Eulerian higher-order hydrodynamics turbulence modeling code, and an Eulerian coupled laser-plasma interaction code. We compress relevant quantities after each time-step to approximate the effects of tightly coupled compression and study the compression rates to estimate memory and disk-bandwidth reduction. We find that the error characteristics of compression algorithms must be carefully considered in the context of the underlying physics being modeled.
Award: Best Paper Finalist

Exploring Power Behaviors and Trade-offs of In-Situ Data Analytics

Marc Gamell, Ivan Rodero, Manish Parashar (Rutgers University), Janine C. Bennett (Sandia National Laboratories), Attila Gyulassy (University of Utah), Valerio Pascucci (University of Utah and Pacific Northwest National Laboratory), Patrick McCormick (Los Alamos National Laboratory), Scott Klasky (Oak Ridge National Laboratory), Scott Pakin (Los Alamos National Laboratory), Hemanth Kolla (Sandia National Laboratories), Peer-Timo Bremer (Lawrence Livermore National Laboratory), Jacqueline Chen (Sandia National Laboratories), Aaditya G. Landge (University of Utah)

As scientific applications target exascale, challenges related to data and energy are becoming dominating concerns. For example, coupled simulation workflows are increasingly adopting in-situ data processing and analysis techniques to address costs and overheads due to data movement and I/O. However it is also critical to understand these overheads and associated trade-offs from an energy perspective. The goal of this paper is exploring data-related energy/performance trade-offs for end-to-end simulation workflows running at scale on current high-end computing systems. Specifically, this paper presents: (1) an analysis of the data-related behaviors of a combustion simulation workflow with an in-situ data analytics pipeline, running on the Titan system at ORNL; (2) a power model based on system power and data exchange patterns, which is empirically validated; and (3) the use of the model to characterize the energy behavior of the workflow and to explore energy/performance trade-offs on current as well as emerging systems.

GoldRush: Resource Efficient In Situ Scientific Data Analytics Using Fine-Grained Interference Aware Execution

Fang Zheng (Georgia Institute of Technology), Hongfeng Yu (University of Nebraska-Lincoln), Can Hantas, Matthew Wolf, Greg Eisenhauer, Karsten Schwan (Georgia Institute of Technology), Hasan Abbasi, Scott Klasky (Oak Ridge National Laboratory)

Severe I/O bottlenecks on high end computing platforms call for running data analytics in situ. Demonstrating that there exist considerable resources in compute nodes unused by typical high end scientific simulations, we leverage this fact by creating an agile runtime, termed GoldRush, that can harvest those otherwise wasted, idle resources to efficiently run in situ data analytics. GoldRush uses fine-grained scheduling to “steal” idle resources, in ways that minimize interference between the simulation and in situ analytics. This involves recognizing the potential causes of on-node resource contention and then using scheduling methods that prevent them. Experiments with representative science applications at large scales show that resources harvested on compute nodes can be leveraged to perform useful analytics, significantly improving resource efficiency, reducing data movement costs incurred by alternate solutions, and posing negligible impact on scientific simulations.

Preconditioners and Unstructured Meshes

Chair: George Biros (University of Texas at Austin)
1:30pm-3pm
Room: 401/402/403

A Scalable, Efficient Scheme for Evaluation of Stencil Computations over Unstructured Meshes

James King, Mike Kirby (University of Utah)

Stencil computations are a common class of operations that appear in many computational scientific and engineering applications. Stencil computations often benefit from compile-time analysis, exploiting data-locality, and parallelism. Post-processing of discontinuous Galerkin (dG) simulation solutions with B-spline kernels is an example of a numerical method which requires evaluating computationally intensive stencil operations over a mesh. Previous work on stencil computations has focused on structured meshes, while giving little attention to unstructured meshes. Performing stencil operations over an unstructured mesh requires sampling of heterogeneous elements which often leads to inefficient memory access patterns and limits data locality/reuse. In this paper, we present

an efficient method for performing stencil computations over unstructured meshes which increases data-locality and cache efficiency, and a scalable approach for stencil tiling and concurrent execution. We provide experimental results in the context of post-processing of dG solutions that demonstrate the effectiveness of our approach.

Scalable Domain Decomposition Preconditioners For Heterogeneous Elliptic Problems

Pierre Jolivet, Frédéric Hecht (Laboratoire Jacques-Louis Lions), Frédéric Nataf, Christophe Prud'homme (Institut de Recherche Mathématique Avancée)

Domain decomposition methods are, alongside multigrid methods, one of the dominant paradigms in contemporary large-scale partial differential equation simulation. In this paper, a lightweight implementation of a theoretically and numerically scalable preconditioner is presented in the context of overlapping methods. The performance of this work is assessed by numerical simulations executed on thousands of cores, for solving various highly heterogeneous elliptic problems in both 2D and 3D with billions of degrees of freedom. Such problems arise in computational science and engineering, in solid and fluid mechanics. While focusing on overlapping domain decomposition methods might seem too restrictive, it will be shown how this work can be applied to a variety of other methods, such as non-overlapping methods and abstract deflation based preconditioners. It is also presented how multilevel preconditioners can be used to avoid communication during an iterative process such as a Krylov method.

Award: Best Paper Finalists

Parallel Design and Performance of Nested Filtering Factorization Preconditioner

Long Qu (INRIA and Laboratoire de Recherche en Informatique), Laura Grigori (INRIA), Frédéric Nataf (Laboratoire Jacques-Louis Lions)

We present the parallel design and performance of the nested filtering factorization preconditioner (NFF), which can be used for solving linear systems arising from the discretization of a system of PDEs on unstructured grids. NFF has limited memory requirements, and it is based on a two level recursive decomposition that exploits a nested block arrow structure of the input matrix, obtained priorly by using graph partitioning techniques. It also allows to preserve several directions of interest of the input matrix to alleviate the effect of low frequency modes on the convergence of iterative methods. For a boundary value problem with highly heterogeneous coefficients, discretized on three-dimensional grids with 64 millions unknowns and 447 millions nonzero entries, we show experimentally that NFF scales up to 2048 cores of Genci's Bull system (Curie), and it is up to 2.6 times faster than the domain decomposition preconditioner Restricted Additive Schwarz implemented in PETSc.

Award: Best Student Paper Finalist

Tools for Scalable Analysis

Chair: Dorian C. Arnold (University of New Mexico)
3:30pm-5pm
Room: 405/406/407

Effective Sampling-Driven Performance Tools for GPU-Accelerated Supercomputers

Milind Chabbi, Karthik Murthy, Michael Fagan, John Mellor-Crummey (Rice University)

Performance analysis of GPU-accelerated systems requires a system-wide view that considers both CPU and GPU components. In this paper, we describe how to extend system-wide, sampling-based performance analysis methods to GPU-accelerated systems. Since current GPUs do not support sampling, our implementation required careful coordination of instrumentation-based performance data collection on GPUs with sampling-based methods employed on CPUs. In addition, we also introduce a novel technique for analyzing systemic idleness in CPU/GPU systems. We demonstrate the effectiveness of our techniques with application case studies on Titan and Keeneland. Some of the highlights of our case studies are: 1) we improved performance for LULESH 1.0 by 30%, 2) we identified a hardware performance problem on Keeneland, 3) we identified a scaling problem in LAMMPS derived from CUDA initialization, and 4) we identified a performance problem that is caused by GPU synchronization operations that suffer delays due to blocking system calls.

Rethinking Algorithm-Based Fault Tolerance with a Cooperative Software-Hardware Approach

Dong Li (Oak Ridge National Laboratory), Zizhong Chen, Panruo Wu (University of California, Riverside), Jeffrey S. Vetter (Oak Ridge National Laboratory)

Algorithm-based fault tolerance (ABFT) is a highly efficient resilience solution for many widely-used scientific computing kernels. However, in the context of the resilience ecosystem, ABFT is completely opaque to any underlying hardware resilience mechanisms. As a result, some data structures are over-protected by ABFT and hardware, which leads to unnecessary costs in terms of performance and energy. In this paper, we rethink ABFT using an integrated view including both software and hardware with the goal of improving performance and energy efficiency of ABFT-enabled applications. In particular, we study how to coordinate ABFT and error-correcting code (ECC) for main memory, and investigate the impact of this coordination on performance, energy, and resilience for ABFT-enabled applications. Scaling tests and analysis indicate that our approach saves up to 25% for system energy (and up to 40% for dynamic memory energy) with up to 18% performance improvement over traditional approaches of ABFT with ECC.

Using Automated Performance Modeling to Find Scalability Bugs in Complex Codes

Alexandru Calotoiu, Torsten Hoefler (ETH Zurich), Marius Poke, Felix Wolf (German Research School for Simulation Sciences)

Many parallel applications suffer from latent performance limitations that may prevent them from scaling to larger machine sizes. Often, such scalability bugs manifest themselves only when an attempt to scale the code is actually being made - a point where remediation can be difficult. However, creating analytical performance models that would allow such issues to be pinpointed earlier is so laborious that application developers attempt it at most for a few selected kernels, running the risk of missing harmful bottlenecks. In this paper, we show how both coverage and speed of this scalability analysis can be substantially improved. Generating an empirical performance model automatically for each part of a parallel program, we can easily identify those parts that will reduce performance at larger core counts. Using a climate simulation as an example, we demonstrate that scalability bugs are not confined to those routines usually chosen as kernels.

Engineering Scalable Applications

Chair: Andrew M. Canning (Lawrence Berkeley National Laboratory)
3:30pm-5pm
Room: 401/402/403

Kinetic Turbulence Simulations at Extreme Scale on Leadership-Class Systems

Bei Wang (Princeton University), Stephane Ethier, William Tang (Princeton Plasma Physics Laboratory), Timothy Williams (Argonne National Laboratory), Khaled Ibrahim (Lawrence Berkeley National Laboratory), Kamesh Madduri (Pennsylvania State University), Samuel Williams, Leonid Oliker (Lawrence Berkeley National Laboratory)

Reliable predictive simulation capability addressing confinement properties in magnetically confined fusion plasmas is critically-important for ITER, a 20 billion dollar international burning plasma device under construction in France. The complex study of kinetic turbulence, which can severely limit the energy confinement and impact the economic viability of fusion systems, requires simulations at extreme scale for such an unprecedented device size. Our newly optimized, global, “ab initio” particle-in-cell code solving the nonlinear equations underlying gyrokinetic theory achieves excellent performance with respect to “time to solution” at the full capacity of the IBM Blue Gene/Q on 786,432 cores of “Mira” at ALCF and recently of the 1,572,864 cores of “Sequoia” at LLNL. Recent multithreading and domain decomposition optimizations in the new GTC-P code represent critically important software advances for modern, low memory per core systems by enabling routine simulations at unprecedented size (130 million grid points ITER-scale) and resolution (65 billion particles).

Swendsen-Wang Multi-Cluster Algorithm for the 2D/3D Ising Model on Xeon Phi and GPU

Florian Wende, Thomas Steinke (Zuse Institute Berlin)

Simulations of the critical Ising model by means of local update algorithms suffer from critical slowing down. One way to partially compensate for the influence of this phenomenon on the runtime of simulations is using increasingly faster and parallel computer hardware. Another approach is using algorithms that do not suffer from critical slowing down, such as cluster algorithms. This paper reports on the Swendsen-Wang multi-cluster algorithm on Intel Xeon Phi coprocessor 5110P, Nvidia Tesla M2090 GPU, and x86 multi-core CPU. We present shared memory versions of the said algorithm for the simulation of the two- and three-dimensional Ising model. We use a combination of local cluster search and global label reduction by means of atomic hardware primitives. Further, we describe an MPI version of the algorithm on Xeon Phi and CPU, respectively. Significant performance improvements over known implementations of the Swendsen-Wang algorithm are demonstrated.

Mr. Scan: Extreme Scale Density-Based Clustering Using a Tree-Based Network of GPGPU Nodes

Benjamin R. Welton, Evan H. Samanas, Barton P. Miller (University of Wisconsin - Madison)

Density-based clustering algorithms are a widely-used class of data mining techniques that can find irregularly shaped clusters and cluster data without prior knowledge of the number of clusters it contains. DBSCAN is the most well-known density-based clustering algorithm. We introduce our version of DBSCAN, called Mr. Scan, which uses a hybrid parallel implementation that combines the MRNet tree-based distribution network with GPGPU-equipped nodes. Mr. Scan avoids the problems of existing implementations by effectively partitioning the point space and by optimizing DBSCAN’s computation over dense data regions. We tested Mr. Scan on both a geolocated Twitter dataset and image data obtained from the Sloan Digital Sky Survey. At its largest scale, Mr. Scan clustered 6.5 billion points from the Twitter dataset on 8,192 GPU nodes on Cray Titan in 17.3 minutes. All other parallel DBSCAN implementations have only demonstrated the ability to cluster up to 100 million points.

Improving Large-Scale Computation and Data Resources

Chair: Robert A. Ballance (Sandia National Laboratories)
3:30pm-5pm
Room: 205/207

The Science DMZ: A Network Design Pattern for Data-Intensive Science

Eli Dart, Lauren Rotman, Brian Tierney, Mary Hester (Lawrence Berkeley National Laboratory), Jason Zurawski (Internet2)

The ever-increasing scale of scientific data has become a significant challenge for researchers that rely on networks to interact with remote computing systems and transfer results to collaborators worldwide. Despite the availability of high-capacity connections, scientists struggle with inadequate cyberinfrastructure that cripples data transfer performance, and impedes scientific progress.

The Science DMZ paradigm comprises a proven set of network design patterns that collectively address these problems for scientists. We explain the Science DMZ model, including network architecture, system configuration, cybersecurity, and performance tools, that creates an optimized network environment for science. We describe use cases from universities, supercomputing centers and research laboratories, highlighting the effectiveness of the Science DMZ model in diverse operational settings.

In all, the Science DMZ model is a solid platform that supports any science workflow, and flexibly accommodates emerging network technologies. As a result, the Science DMZ vastly improves collaboration, accelerating scientific discovery.
Award: Best Paper Finalist

Enabling Comprehensive Data-Driven System Management for Large Computational Facilities

James C. Browne (University of Texas at Austin), Robert L. DeLeon, Charng-Da Lu, Thomas R. Furlani, Matthew D. Jones, Steven M. Gallo, Abani K. Patra (SUNY at Buffalo), William L. Barth (University of Texas at Austin), John Hammond (Intel Corporation)

This paper presents a comprehensive system meeting the information requirements of all stakeholders (users, application developers, consultants, systems administrators and system management) of large cluster computers. Accounting, scheduler and event logs are integrated with data from a new measurement tool (called TACC_Stats). TACC_Stats periodically records resource use including many hardware counters for each node executing for each job and by aggregation, system level metrics. Analysis of this data by the XDMOD reporting system generates many analyses and reports which have not previously been available for open-source, Linux-based software systems. This paper systematically identifies all of the information that is needed for effective execution of the role of each stakeholder class and generates example reports and discusses their value propositions. We believe this system to be the first fully comprehensive system for supporting the information needs of all of the stakeholders for open-source software based cluster systems.

Insights for Exascale IO APIs from Building a Petascale IO API

Jay Lofstead (Sandia National Laboratories), Robert Ross (Argonne National Laboratory)

Near the dawn of the petascale era, IO libraries had reached a stability in their function and data layout with only incremental changes being incorporated. The shift in technology, particularly the scale of parallel file systems and the number of compute processes, prompted revisiting best practices for optimal IO performance.

Along with other efforts like PLFS, the project that led to ADIOS (ADaptable IO System) was motivated by both the shift in technology and the historical requirement for optimal IO performance to change how simulations performed IO depending on the platform. To solve both issues, the ADIOS team, in consultation with other leading IO experts, sought to build a new IO platform based on the assumptions inherent in the petascale hardware platforms.

This paper helps inform the design of future IO platforms with a discussion of lessons learned as part of the process of designing and building ADIOS.

Thursday, November 21

GPU Programming

Chair: Michael Garland (NVIDIA)
10:30am-12pm
Room: 205/207

General Transformations for GPU Execution of Tree Traversals

Michael Goldfarb, Youngjoon Jo, Milind Kulkarni (Purdue University)

With the advent of programmer-friendly GPU computing, there has been much interest in offloading workloads that can exploit the massive parallelism available on modern GPUs. Exploiting this parallelism and optimizing for the GPU memory hierarchy is well-understood for regular applications that operate on dense data structures such as arrays and matrices. However, there has been significantly less work in the area of irregular algorithms and even less so when dynamic data structures are involved. Recently, irregular algorithms such

as Barnes-Hut and kd-tree traversals have been implemented on GPUs, yielding significant performance gains over CPU implementations. However, the implementations often rely on exploiting application-specific semantics to get acceptable performance. We argue that there are general-purpose techniques for implementing irregular algorithms on GPUs that exploit similarities in algorithmic structure rather than application-specific knowledge. We demonstrate these techniques on several tree traversal algorithms, achieving speedups of up to 38x over 32-thread CPU versions.

A Large-Scale Cross-Architecture Evaluation of Thread-Coarsening

Alberto Magni, Christophe Dubach, Michael F.P. O'Boyle (University of Edinburgh)

OpenCL has become the de facto data parallel programming model for parallel devices in today's high-performance supercomputers. OpenCL was designed with the goal of guaranteeing program portability across hardware from different vendors. However, achieving good performance is hard, requiring manual tuning of the program and expert knowledge of each target device.

In this paper we consider a data parallel compiler transformation — thread-coarsening — and evaluate its effects across a range of devices by developing a source-to-source OpenCL compiler based on LLVM. We thoroughly evaluate this transformation on 17 benchmarks and five platforms with different coarsening parameters giving over 43,000 different experiments. We achieve speedups over 9x on individual applications and average speedups ranging from 1.15x on the NVIDIA Kepler GPU to 1.50x on the AMD Cypress GPU. Finally, we use statistical regression to analyze and explain program performance in terms of hardware-based performance counters.

Semi-Automatic Restructuring of Offloadable Tasks for Many-Core Accelerators

Nishkam Ravi, Yi Yang (NEC Laboratories America), Tao Bao (Purdue University), Srimat Chakradhar (NEC Laboratories America)

We propose a new directive to add relaxed semantics to directive-based languages for accelerators. The compiler identifies and generates one or more off-loadable tasks in the neighborhood of the code region marked by the directive. Central to our contribution is the idea of sub-offload and super-offload. In sub-offload, only part of the code region marked by the developer is offloaded to the accelerator, while the other part executes on the CPU in parallel. In super-offload, a code region larger than the one specified by the developer is declared as the offloadable task (e.g., a parent loop).

We develop Elastic Offload Compiler (EOC) for use alongside existing directive-based languages. The current implementation supports LEO for the new Intel Xeon Phi (MIC) architecture. We evaluate EOC with respect to SpecOMP and NAS Parallel Benchmarks. Speedups range between 1.3x-4.4x with the CPU version as baseline and 1.2x-24x with the CPU-MIC version as baseline.

Performance Analysis of Applications at Large Scale

Chair: Darren J. Kerbyson (Pacific Northwest National Laboratory)
10:30am-12pm
Room: 201/203

Tera-Scale 1D FFT with Low-Communication Algorithm on Intel Xeon Phi Coprocessors

Jongsoo Park, Ganesh Bikshandi, Karthikeyan Vaidyanathan, Ping Tak Peter Tang, Pradeep Dubey, Daehyun Kim (Intel Corporation)

This paper demonstrates the first tera-scale performance of Intel Xeon Phi coprocessors on 1D FFT computations. Applying a disciplined performance programming methodology of sound algorithm choice, valid performance model, and well-executed optimizations, we break the tera-flop mark on a mere 64 nodes of Xeon Phi and reach 6.7 TFLOPS with 512 nodes, which is 1.5x than achievable on a same number of Intel Xeon nodes. It is a challenge to fully utilize the compute capability presented by many-core wide-vector processors for bandwidth-bound FFT computation. We leverage a new algorithm, Segment-of-Interest FFT, with low inter-node communication cost, and aggressively optimize data movements in node-local computations, exploiting caches. Our coordination of low communication algorithm and massively parallel architecture for scalable performance is not limited to running FFT on Xeon Phi; it can serve as a reference for other bandwidth-bound computations and for emerging HPC systems that are increasingly communication limited.

A Framework for Hybrid Parallel Flow Simulations with a Trillion Cells in Complex Geometries

Christian Godenschwager, Florian Schornbaum, Martin Bauer, Harald Köstler, Ulrich Rüde (Friedrich-Alexander-Universität Erlangen-Nürnberg)

WaLBerla is a massively parallel software framework for simulating complex flows with the lattice Boltzmann method (LBM). Performance and scalability results are presented for SuperMUC, the world's fastest x86-based supercomputer ranked number 6 on the Top500 list, and JUQUEEN, a Blue~Gene/Q system ranked as number 5.

We reach resolutions with more than one trillion cells and perform up to 1.93 trillion cell updates per second using 1.8 million threads. The design and implementation of WaLBerla is driven by a careful analysis of the performance on current petascale supercomputers. Our fully distributed data structures and algorithms allow for efficient, massively parallel simulations on these machines. Elaborate node level optimizations and vectorization using SIMD instructions result in highly optimized compute kernels for the single- and two-relaxation-time LBM. Excellent weak and strong scaling is achieved for a complex vascular geometry of the human coronary tree.

A New Routing Scheme for Jellyfish and its Performance with HPC Workloads

Xin Yuan, Santosh Mahapatra, Wickus Nienaber (Florida State University), Scott Pakin (Los Alamos National Laboratory), Michael Lang (Los Alamos National Laboratory)

The jellyfish topology where switches are connected using a random graph has recently been proposed for large scale data-center networks. It has been shown to offer higher bisection bandwidth and better permutation throughput than the corresponding fat-tree topology with a similar cost. In this work, we propose a new routing scheme for jellyfish that out-performs existing schemes by more effectively exploiting the path diversity, and comprehensively compare the performance of jellyfish and fat-tree topologies with HPC workloads. The results indicate that both jellyfish and fat-tree topologies offer comparable high performance for HPC workloads on systems that can be realized by 3-level fat-trees using the current technology and the corresponding jellyfish topologies with similar costs. Fat-trees are more effective for smaller systems while jellyfish is more scalable.

Memory Resilience

Chair: Scott Pakin (Los Alamos National Laboratory)
10:30am-12pm
Room: 405/406/407

Feng Shui of Supercomputer Memory: Positional Effects in DRAM and SRAM Faults

Vilas Sridharan (Advanced Micro Devices, Inc.), Jon Stearley (Sandia National Laboratories), Nathan DeBardeleben, Sean Blanchard (Los Alamos National Laboratory), Sudhanva Gurumurthi (Advanced Micro Devices, Inc.)

Several recent publications confirm that faults are common in high-performance computing systems. Therefore, further attention to the faults experienced by such computing systems is warranted. In this paper, we present a study of DRAM and SRAM faults in large high-performance computing systems. Our goal is to understand the factors that influence faults in production settings. We examine the impact of aging on DRAM, finding a marked shift from permanent to transient faults in the first two years of DRAM lifetime. We examine the impact of DRAM vendor, finding that fault rates vary by more than 4x among vendors. We examine the physical location of faults in a DRAM device and in a data center; contrary to prior studies, we find no correlations with either. Finally, we study SRAM, finding that altitude has a substantial impact on SRAM faults, and that top of rack placement correlates with 20% higher fault rate.

Exploring DRAM Organizations for Energy-Efficient and Resilient Exascale Memories

Bharan Giridhar, Michael Cieslak, Deepankar Duggal, Ronald Dreslinski (University of Michigan), Hsing Min Chen (Arizona State University), Robert Patti (Tezzaron Semiconductor), Betina Hold (ARM Ltd.), Chaitali Chakrabarti (Arizona State University), Trevor Mudge, David Blaauw (University of Michigan)

The power target for exascale supercomputing is 20MW, with about 30% budgeted for the memory subsystem. Commodity DRAMs will not satisfy this requirement. Additionally, the large number of memory chips (>10M) required will result in crippling failure rates. Although specialized DRAM memories have been reorganized to reduce power through 3D-stacking or row buffer resizing, their implications on fault tolerance have not been considered. We show that addressing reliability and energy is a co-optimization problem involving tradeoffs between error correction cost, access energy and refresh power; reducing the physical page size to decrease access energy increases the energy/area overhead of error resilience. Additionally, power can be reduced by optimizing bitline lengths. The proposed 3D-stacked memory uses a page size of 4kb and consumes 5.1pJ/bit based on simulations with NEK5000 benchmarks. Scaling to 100PB, the memory consumes 4.7MW at 100PB/s which, while well within the total power budget (20MW), is also error-resilient.

Low-Power, Low-Storage-Overhead Chipkill Correct via Multi-Line Error Correction

Xun Jian, Henry Duwe (University of Illinois at Urbana-Champaign), John Sartori (University of Minnesota), Vilas Sridharan (Advanced Mirco Devices, Inc.), Rakesh Kumar (University of Illinois at Urbana-Champaign)

Due to their large memory capacities, many modern servers require chipkill correct, an advanced type of memory error detection and correction, to meet their reliability requirements. However, existing chipkill correct solutions incur high power or storage overheads, or both because they use dedicated error-correction resources per codeword to perform error correction. This requires high overhead for correction and results in high overhead for error detection. We propose a novel chipkill correct solution---multi-line error correction---that uses resources shared across multiple lines in memory for error correction to reduce the overhead of both error detection and correction. Our evaluations show that the proposed solution

reduces memory power by a mean of 27%, and up to 38% with respect to commercial solutions, at a cost of 0.4% increase in storage overhead and minimal impact on reliability.

Matrix Computations

Chair: Laura Grigori (INRIA)
10:30am-12pm
Room: 401/402/403

Parallel Reduction to Hessenberg Form with Algorithm-Based Fault Tolerance

Yulu Jia, George Bosilca, Piotr Luszczek, Jack Dongarra (University of Tennessee, Knoxville)

This paper studies the resilience of two-sided factorizations and presents a generic algorithm-based approach capable of rendering two-sided factorizations resilient. We establish the theoretical proof of the correctness and the numerical stability of the approach in the context of a Hessenberg Reduction (HR) and present the scalability and performance results of a practical implementation. Our method is a hybrid algorithm combining an Algorithm Based Fault Tolerance (ABFT) technique with diskless checkpointing to fully protect the data. We protect the trailing and the proceeding matrix with checksums, and protect finished panels in the panel scope with diskless checkpoints. Compared with the original HR (the ScaLAPACK PDGEHRD routine) our fault-tolerant algorithm introduces very little overhead, and maintains the same level of scalability. We prove that the overhead shows a decreasing trend as the size of the matrix or the size of the process grid increases.

A Computationally Efficient Algorithm for the 2D Covariance Method

Oded Green (Georgia Institute of Technology), Yitzhak Birk (Technion - Israel Institute of Technology)

The estimated covariance matrix is a building block for many algorithms, including signal and image processing. The Covariance Method is an estimator for the covariance matrix, favored both as an estimator and in view of the convenient properties of the matrix that it produces. However, the considerable computational requirements limit its use. We present a novel computation algorithm for the covariance method, which dramatically reduces the computational complexity (both ALU operations and memory access) relative to previous algorithms. It has a small memory footprint, is highly parallelizable and requires no synchronization among compute threads. On a 40-core X86 system, we achieve 1200X speedup relative to a straightforward single-core implementation; even on a single core, 35X speedup is achieved.

An Improved Parallel Singular Value Algorithm and Its Implementation for Multicore Hardware

Azzam Haidar, Jakub Kurzak, Piotr Luszczek (University of Tennessee, Knoxville)

The enormous gap between the high-performance capabilities of today's CPUs and off-chip communication has made the development of numerical software that is scalable and performant extremely challenging.

In this paper, we describe a successful methodology to address these challenges, starting from our algorithm design, kernel optimization and tuning, to our programming model in the development of a scalable high-performance singular-value-decomposition (SVD) solver. We developed a set of leading edge kernels combined with advanced optimization techniques featuring fine-grained, memory-aware kernels, a task-based approach and hybrid execution and scheduling that significantly increase the performance of the SVD solver.

Our results demonstrate an enormous performance boost compared to current available software. In particular, our software is two-fold faster than the optimized Intel Math Kernel Library when all the singular vectors are required, achieves 4-times speedup when 20% of the vectors are computed and is significantly superior (12X) if only the singular-value is required.

Memory Hierarchy

Chair: Mark Gardner (Virginia Tech)
1:30pm-3pm
Room: 405/406/407

Performance Evaluation of Intel Transactional Synchronization Extensions for High Performance Computing

Richard Yoo, Christopher Hughes, Konrad Lai, Ravi Rajwar (Intel Corporation)

Intel has recently introduced Intel Transactional Synchronization Extensions (Intel TSX) in the Intel 4th Generation Core Processors. With Intel TSX, a processor can dynamically determine whether threads need to serialize through lock-protected critical sections. In this paper, we evaluate the first hardware implementation of Intel TSX using a set of high-performance computing (HPC) workloads, and demonstrate that applying Intel TSX to these workloads can provide significant performance improvements. On a set of real-world HPC workloads, applying Intel TSX provides an average speedup of 1.41x. When applied to a parallel user-level TCP/IP stack, Intel TSX provides 1.31x average bandwidth improvement on network intensive applications. We also demonstrate the ease with which we were able to apply Intel TSX to the various workloads.

Location-Aware Cache Management for Many-Core Processors with Deep Cache Hierarchy

Jongsoo Park, Richard M. Yoo (Intel Corporation), Daya S. Khudia (University of Michigan), Christopher J. Hughes (Intel Corporation), Daehyun Kim (Intel Corporation)

As cache hierarchies become deeper and the number of cores on a chip increases, managing caches becomes more important for performance and energy. However, current hardware cache management policies do not always adapt optimally to the applications behavior: e.g., caches may be polluted by data structures whose locality cannot be captured by the caches, and producer-consumer communication incurs multiple round trips of coherence messages per cache line transferred. We propose load and store instructions that carry hints regarding into which cache(s) the accessed data should be placed. Our instructions allow software to convey locality information to the hardware, while incurring minimal hardware cost and not affecting correctness. Our instructions provide a 1.16x speedup and a 1.96x energy efficiency boost, on average, according to simulations on a 64-core system with private L1 and L2 caches. With a large shared L3 cache added, the benefits increase, providing 1.33x energy reduction on average.

Practical Nonvolatile Multilevel-Cell Phase Change Memory

Doe Hyun Yoon (IBM Thomas J. Watson Research Center), Jichuan Chang, Robert S. Schreiber, Norman P. Jouppi (Hewlett-Packard)

Multilevel-cell (MLC) phase change memory (PCM) may provide both high capacity main memory and faster-than-Flash persistent storage. But slow growth in cell resistance with time, resistance drift, can cause transient errors in MLC-PCM. Drift errors increase with time, and prior work suggests refresh before the cell loses data. The need for refresh makes MLC-PCM volatile, taking away a key advantage. Based on the observation that most drift errors occur in a particular state in four-level-cell PCM, we propose to change from four levels to three levels, eliminating the most vulnerable state. This simple change lowers cell drift error rates by many orders of magnitude: three-level-cell PCM can retain data without power for more than ten years. With optimized encoding/decoding and a wearout tolerance mechanism, we can narrow the capacity gap between three-level and four-level cells. These techniques together enable low-cost, high-performance, genuinely non-volatile MLC-PCM.

MPI Performance and Debugging

Chair: Thomas Fahringer (University of Innsbruck)
1:30pm-3pm
Room: 201/203

Distributed Wait State Tracking for Runtime MPI Deadlock Detection

Tobias Hilbrich (Technische Universität Dresden), Bronis R. de Supinski (Lawrence Livermore National Laboratory), Wolfgang E. Nagel (Technische Universität Dresden), Joachim Protze (Aachen University), Christel Baier (Technische Universität Dresden), Matthias S. Mueller (Aachen University)

The widely used Message Passing Interface (MPI) with its multitude of communication functions is prone to usage errors. Runtime error detection tools aid in the removal of these errors. We develop MUST as one such tool that provides a wide variety of automatic correctness checks. Its correctness checks can be run in a distributed mode, except for its deadlock detection. This limitation applies to a wide range of tools that either use centralized detection algorithms or a timeout approach. In order to provide scalable and distributed deadlock detection with detailed insight into deadlock situations, we propose a model for MPI blocking conditions that we use to formulate a distributed algorithm. This algorithm implements scalable MPI deadlock detection in MUST. Stress tests at up to 4,096 processes demonstrate the scalability of our approach. Finally, overhead results for a complex benchmark suite demonstrate an average runtime increase of 34% at 2,048 processes.

Globalizing Selectively: Shared-Memory Efficiency with Address-Space Separation

Nilesh Mahajan, Uday Pitambare, Arun Chauhan (Indiana University)

It has become common for MPI-based applications to run on shared-memory machines. However, MPI semantics do not allow leveraging shared memory fully for communication between processes from within the MPI library. This paper presents an approach that combines compiler transformations with a specialized runtime system to achieve zero-copy communication whenever possible by proving certain properties statically and globalizing data selectively by altering the allocation and deallocation of communication buffers. The runtime system provides dynamic optimization, when such proofs are not possible statically, by copying data only when there are write-write or read-write conflicts. We implemented a prototype compiler, using ROSE, and evaluated it on several benchmarks. Our system produces code that performs better than MPI in most cases and no worse than MPI, tuned for shared memory, in all cases.

Hybrid MPI: Efficient Message Passing for Multi-Core Systems

Andrew Friedley (Indiana University), Greg Bronevetsky (Lawrence Livermore National Laboratory), Torsten Hoefler (ETH Zurich), Andrew Lumsdaine (Indiana University)

Multi-core shared memory architectures are ubiquitous in both High-Performance Computing (HPC) and commodity systems because they provide an excellent trade-off between performance and programmability. MPI's abstraction of explicit communication across distributed memory is very popular for programming scientific applications. Unfortunately, OS-level process separations force MPI to perform unnecessary copying of messages within shared memory nodes. This paper presents a novel approach that transparently shares memory across MPI processes executing on the same node, allowing them to communicate like threaded applications. While prior work explored thread-based MPI libraries, we demonstrate that this approach is impractical and performs poorly in practice. We instead propose a novel process-based approach that enables shared memory communication and integrates with existing MPI libraries and applications without modifications. Our protocols for shared memory message passing exhibit better performance and reduced cache footprint. Communication speedups of more than 26% are demonstrated for two applications.

Sorting and Graph Algorithms

Chair: Karen Devine (Sandia National Laboratories)
1:30pm-3pm
Room: 205/207

Distributed-Memory Parallel Algorithms for Generating Massive Scale-Free Networks Using Preferential Attachment Model

Maksudul Alam, Maleq Khan, Madhav Marathe (Virginia Tech)

Recently, there has been substantial interest in the study of various random networks as mathematical models of complex systems. As these complex systems grow larger, the ability to generate progressively large random networks becomes all the more important. This motivates the need for efficient parallel algorithms for generating such networks. Naive parallelization of the sequential algorithms for generating random networks may not work due to the dependencies among the edges and the possibility of creating duplicate (parallel) edges. In this paper, we present MPI-based distributed memory parallel algorithms for generating random scale-free networks using the preferential-attachment model. Our algorithms scale very well to a large number of processors and provide almost linear speedups. The algorithms can generate scale-free networks with 50 billion edges in 123 seconds using 768 processors.

On Fast Parallel Detection of Strongly Connected Components (SCC) in Small-World Graphs

Sungpack Hong (Oracle Corporation), Nicole C. Rodia, Kunle Olukotun (Stanford University)

Detecting strongly connected components (SCCs) in a directed graph is a fundamental graph analysis algorithm that is used in many science and engineering domains. Traditional approaches in parallel SCC detection, however, show limited performance and poor scaling behavior when applied to large real-world graph instances. In this paper, we investigate the shortcomings of the conventional approach and propose a series of extensions that consider the fundamental properties of real-world graphs, e.g. the small-world property. Our scalable implementation offers excellent performance on diverse, small-world graphs resulting in a 5.01x to 29.41x parallel speedup over the optimal sequential algorithm with 16 cores and 32 hardware threads.

Algorithms for High-Throughput Disk-to-Disk Sorting

Hari Sundar, Dhairya Malhotra, Karl Schulz (University of Texas at Austin)

In this paper, we present a new out-of-core sort algorithm, designed for problems that are too large to fit into the aggregate RAM available on modern supercomputers. We analyze the performance including the cost of IO and demonstrate the fastest reported throughput using the canonical sortBenchmark on a general-purpose, production HPC resource running Lustre. By clever use of available storage and a formulation of asynchronous data transfer mechanisms, we are able to almost completely hide the computation (sorting) behind the IO latency. This latency hiding enables us to achieve comparable execution times, including the additional temporary IO required, between a large sort problem (5TB) run as a single, in-RAM sort and our out-of-core approach using 1/10th the amount of RAM. In our largest run, sorting 100TB of records using 1792-hosts, we achieved an end-to-end throughput of 1.24TB/min using our general-purpose sorter, improving on the current Daytona record holder by 65%.

Optimizing Numerical Code

Chair: Naoya Maruyama (RIKEN Advanced Institute for Computational Science)
3:30pm-5pm
Room: 401/402/403

AUGEM: Automatically Generate High Performance Dense Linear Algebra Kernels on x86 CPUs

Qian Wang, Xianyi Zhang, Yunquan Zhang (Institute of Software, Chinese Academy of Science), Qing Yi (University of Colorado Colorado Springs)

Basic liner algebra subprograms (BLAS) is a fundamental library in scientific computing. In this paper, we present a template-based optimization framework, AUGEM, which can automatically generate fully optimized assembly code for several dense linear algebra (DLA) kernels, such as GEMM, GEMV, AXPY and DOT, on varying multi-core CPUs without requiring any manual interference from developers. In particular, based on domain-specific knowledge about algorithms of the DLA kernels, we use a collection of parameterized code templates to formulate a number of commonly occurring instruction sequences within the optimized low-level C code of these DLA kernels. Then, our framework uses a specialized low-level C optimizer to identify instruction sequences that match the predefined code templates and thereby translates them into extremely efficient SSE/AVX instructions. The DLA kernels generated by our template-based approach surpass the implementations of Intel MKL and AMD ACML BLAS libraries, on both Intel Sandy Bridge and AMD Piledriver processors.

Accelerating Sparse Matrix-Vector Multiplication on GPUs using Bit-Representation-Optimized Schemes

Wai Teng Tang, Wen Jun Tan (Nanyang Technological University), Rajarshi Ray, Yi Wen Wong, Weiguang Chen (National University of Singapore), Shyh-hao Kuo, Rick Siow Mong Goh (Agency for Science, Technology and Research), Stephen John Turner (Nanyang Technological University), Weng-Fai Wong (National University of Singapore)

The sparse matrix-vector (SpMV) multiplication routine is an important building block used in many iterative algorithms for solving scientific and engineering problems. One of the main challenges of SpMV is its memory-boundedness. Although compression has been proposed previously to improve SpMV performance on CPUs, its use has not been demonstrated on the GPU because of the serial nature of many compression and decompression schemes. In this paper, we introduce a family of bit-representation-optimized (BRO) compression schemes for representing sparse matrices on GPUs. The proposed schemes, BRO-ELL, BRO-COO, and BRO-HYB, perform compression on index data and help to speed up SpMV on GPUs through reduction of memory traffic. Furthermore, we formulate a BRO-aware matrix reordering scheme as a data clustering problem and use it to increase compression ratios. With the proposed schemes, experiments show that average speedups of 1.5X compared to ELLPACK and HYB can be achieved for SpMV on GPUs.

Precimonious: Tuning Assistant for Floating-Point Precision

Cindy Rubio-González, Cuong Nguyen, Hong Diep Nguyen, James Demmel, William Kahan, Koushik Sen (University of California, Berkeley), David H. Bailey, Costin Lancu (Lawrence Berkeley National Laboratory), David Hough (Oracle Corporation)

Given the variety of numerical errors that can occur, floating-point programs are difficult to write, test and debug. One common practice employed by developers without an advanced background in numerical analysis is using the highest available precision. While more robust, this can degrade program performance significantly. In this paper we present Precimonious, a dynamic program analysis tool to assist developers in tuning the precision of floating-point programs. Precimonious performs a search on the types of the floating-point program variables trying to lower their precision subject to accuracy constraints and performance goals. Our tool recommends a type instantiation that uses lower precision while producing an accurate enough answer without causing exceptions. We evaluate Precimonious on several widely used functions from the GNU Scientific Library, two NAS Parallel Benchmarks, and three other numerical programs. For most of the programs analyzed, Precimonious reduces precision, which results in performance improvements as high as 41%.

Application Performance Characterization

Chair: Yao Zhang (Institute for Defense Analyses)
3:30pm-5pm
Room: 405/406/407

An Early Performance Evaluation of Many Integrated Core Based SGI Rackable Computing System

Subhash Saini, Haoqiang Jin, Dennis Jespersen (NASA Ames Research Center), Huiyu Feng (SGI Inc), Jahed Djomehri, William Arasin, Robert Hood (Computer Science Corp.), Piyush Mehrotra, Rupak Biswas (NASA Ames Research Center)

Intel recently introduced the Xeon Phi coprocessor based on the Many Integrated Core architecture featuring 60 cores with a peak performance of 1.0 Tflop/s. NASA has deployed a 128-node SGI Rackable system where each node has two Intel Xeon E2670 8-core Sandy Bridge processors along with two Xeon Phi 5110P coprocessors. We have conducted an early performance evaluation of the Xeon Phi. We used microbenchmarks to measure the latency and bandwidth of memory and interconnect, I/O rates, and the performance of OpenMP directives and MPI functions. We also used OpenMP and MPI versions of the NAS Parallel Benchmarks along with two production CFD applications to test four programming modes: offload, processor native, coprocessor native and symmetric (processor

plus coprocessor). In this paper we present preliminary results based on our performance evaluation of various aspects of a Phi-based system.

Predicting Application Performance using Supervised Learning on Communication Features

Nikhil Jain (University of Illinois at Urbana-Champaign),

Abhinav Bhatele (Lawrence Livermore National Laboratory), Michael Robson (University of Illinois at Urbana-Champaign), Todd Gamblin (Lawrence Livermore National Laboratory), Laxmikant Kale (University of Illinois at Urbana-Champaign)

Task mapping on torus networks has traditionally focused on either reducing the maximum dilation or average number of hops per byte for messages in an application. These metrics make simplified assumptions about the cause of network congestion, and do not provide accurate correlation with execution time. Hence, these metrics cannot be used to reasonably predict or compare application performance for different mappings. In this paper, we attempt to model the performance of an application using communication data, such as the communication graph and network hardware counters. We use supervised learning algorithms, such as randomized decision trees, to correlate performance with prior and new metrics. We propose new hybrid metrics that provide high correlation with application performance, and may be useful for accurate performance prediction. For three different communication patterns and a production application, we demonstrate a very strong correlation between the proposed metrics and the execution time of these codes.

Investigating Applications Portability with the Uintah DAG-Based Runtime System on PetaScale Supercomputers

Qingyu Meng, Alan Humphrey, John Schmidt, Martin Berzins (University of Utah)

Present trends in high performance computing present formidable challenges for applications code using multicore nodes with co-processors and reduced memory while still attaining scalability. Software frameworks that execute machine-independent applications code using a runtime system that shields users from architectural complexities offer a possible solution. The Uintah framework for example, solves a broad class of large-scale problems on structured adaptive grids using fluid-flow solvers coupled with particle-based solids methods. Uintah executes directed acyclic graphs of computational tasks with a scalable asynchronous and dynamic runtime system for CPU cores and/or accelerators/co-processors on a node. Uintah's clear separation between application and runtime code has led to scalability increases of 1000x without significant changes to application code. In this work, this methodology is tested on three leading Top500 machines; OLCF Titan, TACC Stampede and ALCF Mira using three diverse and challenging applications problems.

Posters/Scientific Visualization Showcase

Posters/Scientific Visualization Showcase

Posters provide an excellent opportunity for short presentations and informal discussions with conference attendees. This year's posters display cutting-edge, interesting research in high performance computing, storage, networking and analytics. This year we received 211 regular, electronic, and education poster submissions covering a wide variety of research topics in HPC. The committee accepted 83 posters, reflecting an average acceptance rate below 40%. Twenty-four ACM Student Research Competition posters were selected after a rigorous peer-review process---from 50 submissions---and are also on display. These posters compete for the best graduate and undergraduate student poster awards.

The Posters Exhibit kicks off at 5:15 pm. Tuesday, Nov. 19, with a reception in the posters display area of the convention center, in the corridor above the Mile High Ballroom and in the Ballrooms Pre-function area. Complimentary refreshments and appetizers are available until 7 p.m.

Scientific Visualization Showcase

The Scientific Visualization Showcase is back for a third year. Selected entries will be displayed live in a museum/art gallery format so that attendees can experience and enjoy the latest in science and engineering HPC results expressed through state-of-the-art visualization technologies.

Posters

Tuesday, November 19

Exhibit
8:30am-7pm
Reception & Exhibit
5:15pm-7pm
Room: Mile High Ballroom and Pre-function area

Research Posters
ACM Student Research Competition Posters
Electronic Posters

The Posters Exhibit, consisting of Research Posters, ACM Student Research Competition Posters, and electronic posters kicks off at 8:30am, Tuesday, November 19, with a reception at 5:15pm. The reception is an opportunity for attendees to interact with poster presenters and includes regular, electronic, and educational posters as well as the ACM Student Research Competition (SRC) posters. The reception is open to all attendees with Technical Program registration. Complimentary refreshments and appetizers are available until 7pm.

Wednesday, November 20
Thursday, November 21

Exhibit
8:30am-5pm

Research Posters

Improving the I/O Throughput for Data-Intensive Scientific Applications with Efficient Compression Mechanisms
Dongfang Zhao (Illinois Institute of Technology), Jian Yin (Pacific Northwest National Laboratory), Ioan Raicu (Illinois Institute of Technology)

Today's science is generating significantly larger volumes of data than before, making many scientific applications bounded on I/O rather than computation. To reduce the I/O time, data compression is getting more attractive and practical. Most of existing compression techniques are general purpose and not particularly crafted for scientific applications. In this context, we devise a new compression algorithm by carefully exploiting the characteristics of scientific data, namely the incremental changes to the data. Our proposed algorithm segments the data with different starting points and then stores only these increments---essentially achieving both computation and space efficiency. We design and implement a real system with the proposed algorithm at the filesystem level so that the compression and decompression is transparent to the end users and does not need any modification to the applications. The system is evaluated on a 1024-core IBM BlueGene/P supercomputer with the 128-node GPFS file system.

On the Gamification of a Graduate Course on Cloud Computing
Alexandru Iosup, Dick Epema (Delft University of Technology)
Technical universities, especially in Europe, are facing an important challenge in attracting more diverse groups of students and keeping the students they attract motivated by and engaged in the curriculum. Gamification, defined loosely as the use of social gaming elements to activate students in learning activities, can help with addressing this challenge. This poster introduces gamification as a teaching technique and discusses its (experimental) application in the area of distributed HPC education through a graduate course at TU Delft. We discuss the goal and topics of the course; and the gamification mechanics, dynamics, and content-related elements. Over the past two years, we have seen an accumulation of evidence that gamified courses can lead to a high ratio of students who pass after the first attempt, and that gamification can be promising when compared inter- and intra-cohort.

Enhancing Learning-Based Autotuning with Composite and Diagnostic Feature Vectors
Saami Rahman, Apan Qasem (Texas State University)

The use of machine learning techniques has emerged as a promising strategy for autotuning. Central to the success of such learned heuristics is the construction of a feature vector that accurately captures program behavior. Although the salient features of certain domain specific kernels are well understood, at least in principle, automatically deriving suitable features for general numerical applications remains a significant challenge. This poster presents a learning-based autotuning system that introduces two new classes of features that encapsulate key architecture-sensitive performance characteristics for a broad range of scientific applications. The first class of features are high-level, composite and derived from compiler models for data locality and parallelization. The second class of features is a series of synthesized and normalized HW performance counter values that diagnose causes of program inefficiencies. Preliminary experimental results show that the enhanced feature vectors increase prediction accuracy by as much as 23% for several learning algorithms.

Preparing Computer Science Students for an Increasingly Parallel World: Teaching Parallel Computing Early and Often

Martin Burtscher, Wuxu Peng, Apan Qasem, Hongchi Shi, Dan Tamir (Texas State University)

The widespread deployment of multicore systems over the last decade has brought about drastic changes in the computing landscape. However, most undergraduate computer science (CS) curricula have yet to embrace the pervasiveness of parallel computing. In their first years, students are typically exclusively trained to think and program sequentially. Too firm a root in sequential thinking can be a non-trivial barrier for parallel thinking and computing. Thus, there is an urgent need of including parallel computing concepts earlier in the CS program. This poster describes an innovative early-and-often approach of systematically integrating parallel computing into existing curricula. Through a series of self-contained, context-aware course modules, parallel concepts are introduced in early CS courses; reinforced throughout the curriculum, and collated in an upper-level capstone course. The project bears special significance as it leverages Texas State's diverse student body to broaden the participation of Hispanic students; a traditionally underrepresented group in STEM.

Optimizing Shared Resource Contention in HPC Clusters

Sergey Blagodurov, Alexandra Fedorova (Simon Fraser University)

Contention for shared resources in HPC clusters occurs when jobs are concurrently executing on the same multicore node (there is a contention for allocated CPU time, shared caches, memory bus, memory controllers, etc). This contention incurs severe degradation to workload performance and stability and hence must be addressed. The state-of-the-art HPC clusters, however, are not contention-aware, with no virtualization supported to mitigate the contention effects through live job migration. The goal of this work is the design, implementation and evaluation of a virtualized HPC scheduling framework that is contention aware.

Petascale General Solver for Semidefinite Programming Problems with Over Two Million Constraints

Katsuki Fujisawa (Chuo University and JST CREST), Toshio Endo, Hitoshi Sato (Tokyo Institute of Technology and JST CREST), Yuichiro Yasui (Chuo University and JST CREST), Naoki Matsuzawa (Tokyo Institute of Technology), Hayato Waki (Kyushu University)

SemiDefinite Programming (SDP) problem is one of the most central problems in mathematical optimization. We have developed the new version of SDPARA, which is a parallel implementation on multiples CPUs and GPUs for solving large-scale SDP problems. SDPARA can attain high scalability using a large number of CPU cores and some techniques of processor affinity and memory interleaving. SDPARA can also perform the parallel Cholesky factorization using thousands of GPUs and techniques to overlap the computation and communication if an SDP problem has over a million constraints. We demonstrate that SDPARA is a petascale general solver for SDP problems in various application fields through numerical experiments on the TSUBAME 2.0 supercomputer and we solved the largest SDP problem (which has over 2.33 million constraints), thereby creating a new world record. Our implementation also achieved 1.018 PFlops in double precision for large-scale Cholesky factorization using 2,720 CPUs and 4,080 GPUs.

Massive Parallelization of a Linear Scaling DFT Code OpenMX

Truong Vinh Truong Duy (Japan Advanced Institute of Science and Technology and University of Tokyo), Taisuke Ozaki (Japan Advanced Institute of Science and Technology)

OpenMX is an open-source first-principles calculation code based on density functional theory for explaining and predicting materials' properties. We massively parallelize OpenMX by developing a domain decomposition method for atoms and grids. In the atom decomposition, we develop a modified recursive bisection method based on the moment of inertia tensor for reordering the atoms from 3D to 1D along a principal axis so that the atoms that are close in real space are also close on the axis to ensure data locality. The atoms are then divided into sub-domains depending on their projections onto the principal axis in a balanced way among the processes. In the grid decomposition, we define four data structures to make data locality consistent with that of the clustered atoms. Benchmark results show that the parallel efficiency at 131,072 cores is 67.7% compared to the baseline of 16,384 cores on the K computer.

Parallel Computing In High Energy Astrophysics: Combining Microlensing and Black Hole Ray-Tracing

Bin Chen, Xinyu Dai, Eddie Baron, Ronald Kantowski (University of Oklahoma)

Gravitational lensing is an important probe in astrophysics. In particular, quasar microlensing strongly constrains the geometry of the high-energy X-ray emitting regions of active galactic nuclei (AGNs). The X-ray emission of a microlensed quasar is lensed by both the background supermassive black hole powering the AGN and the thousands (to millions) of random stars in the foreground lens galaxy. Ray-tracing and parallel computing are needed to compute the image of the accretion disk lensed by a Kerr black hole and the microlensing magnification pattern caused by foreground random stars. We present two parallel ray-tracing codes recently developed by our group to study strong lensing of AGN X-ray emission by Kerr black holes and microlensing by foreground lens galaxies, and a numerical scheme to combine microlensing with Kerr black hole lensing. We discuss the unique constraints on AGN X-ray emission sizes obtained by our computer programs.

Towards Tera-Scale Performance for Longest Common Subsequence Using Graphics Processors

Adnan Ozsoy, Arun Chauhan, Martin Swany (Indiana University Bloomington)

GPUs have been attracting more and more high-performance users recently. However, the computation and memory access patterns in certain classes of algorithms do not lend themselves to hardware optimizations on GPUs, without which they fall far short of the promised performance. One such class of algorithms is longest common subsequence (LCS). In this paper, we describe a novel technique to optimize LCS for one-to-many matching problem on GPUs by transforming the computation into bit-wise operations and a post-processing step. The former can be highly optimized and achieves more than a trillion operations (cell updates) per second (CUPS)---a first for LCS algorithms. The latter is more efficiently done on CPUs, in a fraction of the bit-wise computation time. The bit-wise step promises to be a foundational step and a fundamentally new approach to developing algorithms for increasingly popular heterogeneous environments that could dramatically increase the applicability of hybrid CPU-GPU environments.

Algorithmic Choice in Optimization Problems: A Performance Study

Hammad Rashid, Clara Novoa, Apan Qasem (Texas State University)

To harness the full potential of emerging manycore platforms, the compiler needs to play a key role in exploiting the available on-chip parallelism. Since the amount of extracted parallelism is directly influenced by the selection of the algorithm, algorithmic choice plays a critical role in achieving scalable high performance. This research investigates the impact of algorithmic choice on performance of parallel implementations of optimization problems. The study implements several algorithmic variants of the integral knapsack problem and evaluates each one based on a range of performance characteristics including thread affinity, task granularity and data reuse. Experimental results reveal that selection of the algorithm does have a significant impact on parallel performance. Furthermore, the study provides insight into the relationship of algorithm selection and specific aspects of performance including HW prefetch activity and exploited data locality through favorable cache sharing.

Scibox: Online Sharing of Scientific Data via the Cloud

Jian Huang, Xuechen Zhang, Greg Eisenhauer, Karsten Schwan, Matthew Wolf (Georgia Institute of Technology), Stephane Ethier (Princeton Plasma Physics Laboratory), Scott Klasky (Oak Ridge National Laboratory)

Collaborative science demands global sharing of scientific data. But it cannot leverage universally accessible cloud-based infrastructures like DropBox, as those offer limited interfaces and inadequate levels of access bandwidth. We present the Scibox cloud facility for online sharing scientific data. It uses standard cloud storage solutions, but offers a usage model in which high end codes can access the cloud via the same ADIOS APIs they already use for their I/O actions, thereby naturally coupling data generation with subsequent data analytics. Extending current ADIOS IO methods, with Scibox, data upload/download volumes are controlled via D(ata)R(eduction)-functions stated by end users and applied at the data source, before data is moved, with further gains in efficiency obtained by combining DR-functions to move exactly what is needed by current data consumers. Scibox is evaluated with science applications, e.g., GTS demonstrating the potential for ubiquitous data access with substantial reductions in network traffic.

PUNO: Predictive Unicast and Notification to Mitigate the Mismatch Between Coherence Protocol and Conflict Detection in HTM

Lihang Zhao (Information Sciences Institute), Lizhong Chen (University of Southern California), Jeffrey Draper (Information Sciences Institute)

Hardware Transactional Memory (HTM) usually piggybacks onto the cache coherence protocol to detect data access conflicts between transactions. We identify an intrinsic mismatch between the coherence scheme and transaction execution which leads to a pathological behavior called false aborting. False aborting increases the amount of wasted computation and on-chip communication, manifesting itself as a performance and energy pitfall in HTM designs. For the TM applications we examined, 41% of the transactional write requests incurred false aborting.

To combat false aborting, we propose Predictive Unicast and Notification (PUNO), a novel hardware mechanism to 1) replace the wasteful coherence multicast with a unicast scheme to prevent transactions from being disrupted needlessly; 2) restrain transaction polling through proactive notification to further suppress false aborting. PUNO reduces transaction aborting by 61% in workloads representative of future TM workloads. The network traffic is reduced by 32%. These improvements are achieved with a marginal hardware overhead.

Vectorization of Multi-Center, Highly-Parallel Rigid-Body Molecular Dynamics Simulations

Wolfgang Eckhardt, Alexander Heinecke, Wolfgang Hölzl, Hans-Joachim Bungartz (Technische Universität München)

The molecular dynamics (MD) code ls1 mardyn is being developed for large-scale molecular simulation of e.g. nano- scale flows for fluids composed of simple, small molecules. In contrast to most other MD codes, ls1 mardyn treats molecules as completely rigid bodies with multiple interaction sites, instead of treating all atoms of a molecule individually and applying a constraint motion algorithm. Having previously demonstrated very good results for scalabilty, we present our efficient implementation of the compute kernel in this poster. We focus on the vectorized implementation of the interaction computation for multi-centered molecules using the AVX instruction set and show the performance of the new kernel on SuperMUC hosted at LRZ, Germany. In addition, we include early results for the Intel MIC architecture. On Sandy Bridge, we are able to gain a 3x runtime improvement, while our implementation on KNC achieves the performance of one SuperMUC node.

Quantifying the Dominance of Leakage Energy in Large-Scale System Caches

Aditya M. Deshpande, Jeffrey T. Draper (University of Southern California)

Energy consumption is becoming a critical metric in the design and use of large-scale high-performance systems. With large on-chip caches and advances in chip fabrication technologies, on-chip caches account for a large proportion of total leakage power losses. In this work, we quantify on-chip cache leakage power losses across a wide set of parallel applications and compare the leakage and dynamic energy consumption for various levels of on-chip caches. Our scheme profiles an application to measure cache accesses in order to estimate the energy consumption for various levels of caches. Our study indicates that for various levels of cache, leakage-energy consumption significantly dominates dynamic-energy consumption, and this trend of leakage domination in on-chip caches is expected to increase with every new generation of semiconductor process. Our study indicates that the problem of leakage in on-chip caches cannot be neglected in attacking the energy barrier for building exascale systems.

Vlasiator: Hybrid-Vlasov Simulation Code for the Earth's Magnetosphere

Sebastian von Alfthan, Dimitry Pokhotelov, Yann Kempf, Ilja Honkonen, Sanni Hoilijoki, Arto Sandroos, Minna Palmroth (Finnish Meteorological Institute)

Here we present Vlasiator, the first global magnetospheric simulation code based on a hybrid-Vlasov description of plasma. Ions are represented by a six-dimensional distribution function, while electrons are modeled as a charge neutralizing fluid. To propagate the distribution function we use a finite volume method, coupled to an upwind constrained transport method for propagating electromagnetic fields. We have parallelized the code with a two-level MPI and OpenMP scheme, and it scales well to tens of thousands of cores. The hybrid-Vlasov approach is enabled through a novel sparse representation of the distribution function that reduces the problem size by more than two orders of magnitude. The capabilities of the code are demonstrated by reproducing key features of the collisionless bowshock in a global five-dimensional magnetospheric simulation. Vlasiator produces the first uniformly discretized ion velocity distribution functions with quality comparable to spacecraft measurements, opening up new opportunities for studying the Earth's magnetosphere.

A Highly Scalable Three-Level Schwarz Method with Geometry Preserving Coarse Space for 3D Linear Elasticity Problems on Domains with Complex Geometry

Fande Kong, Xiao-Chuan Cai (University of Colorado Boulder)

The multilevel version of domain decomposition methods is theoretically scalable in the sense that the number of iterations doesn't grow when the number of processors is increased for shorter compute time. However, the theoretically optimal scalability doesn't translate into linear scalability in total compute time, especially when the number of processors is large, because the coarse level solvers aren't scalable in terms of the compute time. We introduce a new way to construct coarse level spaces that preserve the geometric feature of the computational domain, but give up accuracy in the interior of the domain. As it turns out the tradeoff in accuracy provides the high scalability in terms of the total compute time. We show numerically that such a new preconditioner is highly scalable for solving linear elasticity equation discretized on unstructured 3D meshes with hundreds of millions of unknowns on a supercomputer with over 10,000 processors.

Extreme-Scale Space-Time Parallelism

Daniel Ruprecht (University of Lugano), Robert Speck (Forschungzentrum Juelich), Matthew Emmett (Lawrence Berkeley National Laboratory), Matthias Bolten (Universität Wuppertal), Rolf Krause (University of Lugano)

Time-parallel methods for the integration of time-dependent problems introduce concurrency in the temporal direction in addition to typically used parallelization in space. A recently developed algorithm of this type is the "parallel full approximation scheme in space and time" (PFASST). PFASST performs "sweeps" with a low-order time stepper along a hierarchy of space-time meshes, iteratively computing a collocation solution on multiple time-slices simultaneously. For improved parallel efficiency, an FAS correction allows to use space-time coarsening strategies and thus to reduce the computational

cost of sweeps on higher levels. The poster presents novel scaling results of a combination of PFASST with a parallel multi-grid (PMG) solver in space on up to the full IBM Blue Gene/Q system JUQUEEN with 448K cores. It is demonstrated that for large core counts the space-time parallel solver consisting of PFASST plus PMG features significantly better strong scaling capabilities than the pure space-parallel PMG code alone.

Matrix Factorization Routines on Heterogeneous Architectures

Nikita Shustrov, Nadya Mozartova, Tamara Kashevarova, Konstantin Arturov (Intel Corporation)

In this work we consider a method for parallelizing matrix factorization algorithms on the on systems with Intel(R) Xeon Phi(TM) coprocessors. Performance results of matrix factorization routines implementing this approach and available from Intel(R) Math Kernel Library (Intel MKL) on the Intel(R) Xeon(R) platform with Xeon Phi(TM) coprocessors are provided. The implementation of our method is DAG-based and uses panel factorization kernels that were redesigned and rewritten for new Intel(R) Xeon Phi(TM) architectures. The proposed novel method provides a high degree of parallelism while minimizing synchronizations and communications. The algorithm enables adaptable workload distribution between CPUs and coprocessors to improve load balancing. The main features of the algorithm are: adaptive data/task distribution on the fly between CPUs and coprocessors to improve load balancing; efficient utilization of all available computational units in heterogeneous systems; support of heterogeneous systems with unlimited number of coprocessors; and scalability.

Towards Benchmarking Distributed Graph-Processing Platforms

Yong Guo, Marcin Biczak (Delft University of Technology), Ana Lucia Varbanescu (University of Amsterdam), Alexandru Iosup (Delft University of Technology), Claudio Martella (VU University Amsterdam), Theodore L. Willke (Intel Corporation)

Graph-processing platforms are increasingly used in a variety of domains. Although both industry and academia are developing and tuning graph-processing algorithms and platforms, the performance of graph-processing platforms has never been explored or compared in-depth. Thus, users face the daunting challenge of selecting an appropriate platform for their specific application. To alleviate this challenge, we propose an empirical method for benchmarking graph-processing platforms. We define a comprehensive process, and a selection of representative metrics, datasets, and algorithmic classes. We implement a benchmarking suite of five classes of algorithms and seven diverse graphs. Our suite reports on basic (user-lever) performance, resource utilization, scalability, and various overhead. We use our benchmarking suite to analyze and compare six platforms. We gain valuable insights for each platform and present the first comprehensive comparison of graph-processing platforms.

Framework for Optimizing Power, Energy, and Performance

Prasanna Balaprakash (Argonne National Laboratory), Ananta Tiwari (San Diego Supercomputer Center), Stefan Wild (Argonne National Laboratory)

Code optimization in the high-performance computing realm has traditionally focused on reducing execution time. The problem, in mathematical terms, has been expressed as a single objective optimization problem. The expected concerns of next-generation systems, however, demand a more detailed analysis of the interplay among execution time and other metrics. Metrics such as power, performance, energy, and resiliency may all be targeted together and traded against one another. We present a multi-objective formulation of the code optimization problem. Our proposed framework helps one explore potential tradeoffs among multiple objectives and provides a significantly richer analysis than can be achieved by treating additional metrics as hard constraints. We empirically examine a variety of metrics, architectures, and code optimization decisions and provide evidence that such tradeoffs exist in practice.

Optimized Kernels for Large Scale Earthquake Simulations with SeisSol, an Unstructured ADER-DG Code

Alexander Heinecke, Alexander Breuer, Sebastian Rettenberger, Michael Bader (Technical University of Munich), Alice Gabriel, Christian Pelties (Ludwig-Maximilians Universität München)

The software package SeisSol is one of the leading codes for earthquake scenarios, in particular for simulating dynamic rupture processes and for problems that require discretization of very complex geometries. SeisSol uses the discontinuous Galerkin (DG) method for spatial and Arbitrary high order DERivatives (ADER) for time discretization.We present optimizations for the inner, cell-local kernel routines of SeisSol, where SeisSol spends most of the total computing time. The kernels are implemented as a sequence of matrix-matrix-multiplications of a relatively small size and varying sparsity patterns (ranging from “very” sparse to dense). When running the LOH.1 benchmark, employing a 6th order space-time-discretization, in a strong scaling setting on up to 24,000 cores of SuperMUC, we achieved a maximum performance of 95.27 TFLOPS, which corresponds to 18.4% of the theoretical peak performance. By comparing our optimized version to the classic version of SeisSol we obtained a total application speed-up of roughly 3.1-3.9X

X10 at Petascale

Olivier Tardieu, Benjamin Herta, David Cunningham, David Grove, Prabhanjan Kambadur, Vijay A. Saraswat, Avraham Shinnar, Mikio Takeuchi, Mandana Vaziri (IBM Research)

X10 is a high-performance, high-productivity programming language aimed at large-scale distributed and shared-memory parallel applications. It is based on the Asynchronous Partitioned Global Address Space (APGAS) programming model, supporting the same fine-grained concurrency mechanisms within and across nodes.

We demonstrate that X10 delivers solid performance at petascale by running (weak scaling) eight application kernels on an IBM Power 775 supercomputer utilizing up to 55680 Power7 cores (1.7 Pflop/s). We sketch advances in distributed termination detection, distributed load balancing, and use of high-performance interconnects that enable X10 to scale out to thousands of nodes.

Research Activities at Fermilab for Big Data Movement

Parag Mhashilkar, Wenji Wu, Hyun Woo Kim, Gabriele Garzoglio, David Dykstra, Marko Slisz, Phil Demar (Fermi National Accelerator Laboratory)

Fermilab is the US Tier-1 Center for the Large Hadron Collider's (LHC) Compact Muon Solenoid (CMS) experiment and the central data center for several other large-scale research collaborations. These experiments generate a significant amount of data every year which need to be transferred, stored and analyzed. To deal with the scaling and wide-area distribution challenges of the data, the Laboratory is now connected to next generation of networking infrastructure, ESnet's 100G-backbone. Current Fermilab R&D efforts include optimization of network I/O handling in multi-core systems, identification of gaps in the middleware, modification of middleware to improve application performance in 100G network environments, and end-to-end data and network aware storage selection mechanism. This poster describes the multi-prong approach in network-related R&D activities at Fermilab to facilitate big data processing and movement.

Network Traffic Monitoring and Analysis with GPUs

Wenji Wu, Phil Demar (Fermi National Accelerator Laboratory)

In high-speed networks (10 Gbps and above), network traffic monitoring and analysis applications that require scrutiny on a per-packet basis typically demand immense computing power and very high I/O throughputs. These applications case extreme performance and scalability challenges. At Fermilab, we have prototyped a GPU-accelerated architecture for network traffic capturing, monitoring, and analyzing. With a single NVIDIA M2070 GPU, our system can handle 11 million+ packets per second without packet drops. In this paper, we will describe our architectural approach in developing a generic GPU-assisted packet capture and analysis capability.

GMT: Enabling Easy Development and Efficient Execution of Irregular Applications on Commodity Clusters

Alessandro Morari (Pacific Northwest National Laboratory), Oreste Villa (NVIDIA Corporation), Antonino Tumeo, Daniel Chavarria (Pacific Northwest National Laboratory), Mateo Valero (Barcelona Supercomputing Center)

In this poster we introduce GMT (Global Memory and Threading library), a custom runtime library that enables efficient execution of irregular applications on commodity clusters. GMT only requires a cluster with x86 nodes supporting MPI. GMT integrates the Partititioned Global Address Space (PGAS) locality-aware global data model with a fork/join control model common in single node multithreaded environments. GMT supports lightweight software multithreading to tolerate latencies for accessing data on remote nodes, and is built around data aggregation to maximize network bandwidth utilization.

Modules to Teach Parallel Computing Using Python and the LittleFe Cluster

Jose Ortiz-Ubarri, Rafael Arce-Nazario (University of Puerto Rico)

The ability to design effective solutions using parallel processing should be a required competency for every computing student. However, teaching parallel concepts is sometimes challenging and costly. For such reasons here we present a set of modules to teach parallel computing paradigms using python MPI4py and disco in the LittleFe educational cluster. Compiler Independent Strategy for Data Locality Optimization Jinxin Yang, Abid Malik, Barbara Chapman (University of Houston)

Data locality is an important optimization for loop oriented kernels. Auto tuning techniques are used to find the best strategy for data locality optimization. However, auto tuning techniques are expensive and not independent of computing frameworks. Porting an application from one framework to another requires the whole auto tuning process to be repeated, in order to get an optimal solution for the new one. A global strategy will help in expediting the porting process for an application. In this work, we present a framework, consisting of OpenUH transformation directives and CHiLL framework, which provides an optimal strategy for the data locality problem which is independent of compilers. Our results show that the strategies given by our framework clearly out class the default optimization levels of OpenUH, GCC, Intel and PGI compilers.

Task Profiling through OpenMP Runtime API and Tool Support

Ahmad R. Qawasmeh, Abid M. Malik, Deepak Eachempati, Barbara M. Chapman (University of Houston)

The introduction of tasks in the OpenMP programming model brings a new level of parallelism. This also creates new challenges with respect to its applicability through an event-based performance profiling. The OpenMP Architecture Review Board (ARB) approved an interface specification known as the "OpenMP Runtime API (ORA) for Profiling" to enable performance tools to interact with OpenMP programs.

With negligible overheads, we propose new extensions to the ORA for profiling task-level parallelism. We implement these extensions in the open-source OpenUH compiler. We present an efficient method to distinguish individual task instances to capture their associated events. These events illustrate the behavior of OpenMP tasks and identify the overheads associated with the OpenMP tasking model.

Additionally, we integrate our ORA into the performance tool TAU to visualize task measurements at the construct-level. The proposed extensions are in line with the specification recently proposed by the OpenMP tools committee for task profiling.

Performance Analysis of Hybrid BFS Approach Using Semi-External Memory

Keita Iwabuchi, Hitoshi Sato (Tokyo Institute of Technology, JST CREST), Yuichiro Yasui, Katsuki Fujisawa (Chuo University, JST CREST)

NVRAM devices can help processing of extremely large-scale graphs with over DRAM capacity on a single node; however, performance studies and access pattern analyses of graph kernels using both DRAM and NVRAM devices are limited. In order to address the issue, we propose a graph data offloading technique by using NVRAM for the hybrid BFS algorithm and conduct performance analysis of the Hybrid BFS implementation with our data offloading technique. Experimental results show that our approach can achieve 2.8 GTEPS at the maximum and reduce half the size of DRAM with 47.1% performance degradation. The poster also includes performance analyses of our Hybrid BFS approach, which suggests that we can process large-scale graphs with minimum performance degradation using NVRAM by carefully considering the data structures and the access patterns.

A Scalable Hamiltonian Hybrid DFT/PMM-MD Approach for Accurately Simulating Biomolecules on SuperMUC

Magnus Schwörer, Paul Tavan (Ludwig-Maximilians Universität München), Ferdinand Jamitzky, Helmut Satzger (Leibniz Supercomputing Center), Gerald Mathias (Ludwig-Maximilians Universität München)

In this poster, we present an Hamiltonian approach for hybrid molecular dynamics (MD) simulations, in which the forces acting on the atoms are calculated by grid-based density functional theory (DFT) for a solute molecule and by a polarizable molecular mechanics (PMM) force field for a large solvent environment. It guarantees energy conservation and treats the long-range electrostatic interactions within the hybrid simulation system in a linearly scaling fashion using hierarchically nested fast multipole expansions.

Our implementation links the MPI/OpenMP-parallel PMM-MD code IPHIGENIE (developed by us under GPL) to the well-established DFT code CPMD and offers a unique tool which greatly enhances the accuracy of the description of the condensed phase environment of a DFT molecule compared to conventional unpolarizable MM force fields. It scales very well even for small simulation systems and makes studies of large DFT molecules solvated in accurately modeled condensed phase feasible.

Structural Comparison of Parallel Applications

Matthias Weber (Technical University Dresden), Kathryn Mohror, Martin Schulz (Lawrence Livermore National Laboratory), Holger Brunst (Technical University Dresden), Bronis R. de Supinski (Lawrence Livermore National Laboratory), Wolfgang E. Nagel (Technical University Dresden)

With the rising complexity of both architectures and applications, performance analysis and optimization has become essential in the development of scalable applications. Trace-based analysis has proven to be a powerful approach. However, a manual comparison of traces is extremely challenging and time consuming because of the large volume of detailed data and the need to correctly line up trace events. Our solution is a set of techniques that automatically align traces so they can be compared, along with novel metrics that quantify the differences between traces, both in terms of differences in the event stream and timing differences across events. We evaluate clustering algorithms for trace comparison using a novel definition of similarity between traces. We demonstrate the effectiveness of our solution by showing automatically detected performance and code differences across different versions of two real-world applications.

Implementation of Neighbor Communication Algorithm Using Multi-NICs Effectively by Extended RDMA Interfaces

Yoshiyuki Morie (Kyushu University), Takeshi Nanri (Kyushu University)

A neighbor communication algorithm for large messages in K computer was proposed. The key idea of this algorithm was to divide messages into fragments according to the number of neighboring processes and the number of NICs available on each node, so that the bandwidth of each NIC was fully used. To show the effectiveness of the proposed algorithm, an implementation of the algorithm for six neighbors with four NICs by MPI functions was examined. Experimental results showed that, for large messages, the performance of the proposed algorithm was two times faster than the existing one in mpich-3.0.4. However, in medium message size, proposed algorithm is lower performance than existing one. Therefore, the proposed algorithm is implemented by extended RDMA interfaces instead of MPI functions. This implementation became faster than existing one in medium message size too. This result showed the proposed neighbor communication was widely effective.

High Performance CPU/GPU Multiresolution Poisson Solver

Wim van Rees, Babak Hejazialhosseini, Diego Rossinelli, Panagiotis Hadjidoukas, Petros Koumoutsakos (ETH Zurich)

We present a multipole-based N-body solver for 3D multi-resolution, block-structured grids. The solver is designed for a single heterogeneous CPU/GPU compute node, which evaluates the multipole expansions on the CPU while offloading the compute-heavy particle-particle interactions to the GPU. The regular structure of the destination points is exploited for data parallelism on the CPU, to reduce data transfer to the GPU and to minimize memory accesses during evaluation of the direct and indirect interactions. The algorithmic improvements together with HPC techniques lead to 81% and 96% of the predicted peak performance for the CPU and GPU parts, respectively.

Caranx: Scalable Social Image Index Using Phylogenetic Tree of Hashtags

Yusheng Xie (Northwestern University), Zhuoyuan Chen (Adobe Research), Ankit Agrawal (Northwestern University), Wei-keng Liao (Northwestern University), Alok Choudhary (Northwestern University)

Most existing image indexing techniques rely on Scale Invariant Feature Transformation (SIFT) for extracting local point features. Applied to individual image, SIFT extracts hundreds of numerical vectors. The vectors are quantized and stored in tree-like data structures for fast search. SIFT-based indexing can exhibit weakness under certain non-rigid transformations, which are common among real world applications. For example, SIFT often cannot recognize a face as the same with different expressions (e.g. giggling vs. crying). Non-Rigid Dense Correspondence (NRDC) addresses such drawbacks of SIFT. However, directly using NRDC incurs an impractical amount of computation in large-scale image indexing. We present a novel idea here that uses social hashtags to organize the images into a phylogenetic tree (PT). We provide an efficient algorithm to build/search the PT, and show that using PT structure can effectively avoid unnecessary NRDC computation. The resulting image index provides more accurate and diversified search results.

Optimizing User Oriented Job Scheduling within TORQUE

Dalibor Klusácek, Václav Chlumský (CESNET), Hana Rudová (Masaryk University)

We present a major extension of the widely used TORQUE Resource Manager. Unlike common resource managers that use queuing approach, our solution uses planning (job schedule construction) and schedule optimization in order to achieve better predictability, performance and fairness with respect to common queue-based approaches.

Moreover, additional important features are supported, e.g., so called multi-resource fairness that is used to fairly prioritize users subject to their (highly) heterogeneous demands concerning various system resources. Also, new textual and graphical interfaces allow users to better control and plan computation of their jobs.

Our solution is currently undergoing experimental deployment in the Czech national Grid "MetaCentrum."

Extreme Scaling of Real World Applications to >130,000 Cores on SuperMUC

Helmut Satzger, Momme Allalen, Christoph Bernau, Nicolay Hammer, David Brayford, Carmen Navarrete, Ferdinand Jamitzky, Anupam Karmakar (Leibniz Supercomputing Centre), Wolfram Schmidt, Jan Frederik Engels (University of Göttingen), Gurvan Bazin (Universitäts-Sternwarte, Ludwig-Maximilians Universität München), Jürg Diemand (University of Zürich), Klaus Dolag (Universitäts-Sternwarte, Ludwig-Maximilians Universität München), Carsten Kutzner (Max Planck Institute for Biophysical Chemistry, Göttingen), Andreas Marek (Rechenzentrum der Max-Planck-Gesellschaft am Max-Planck-Institut für Plasmaphysik, Garching), Philipp Trisjono (RWTH Aachen)

In July 2013, the Leibniz Supercomputing Centre held the first workshop to test extreme scaling on SuperMUC, the 3 PFLOP/s system with 147,456 Intel Sandy Bridge CPU cores. Groups from 15 international projects came to the LRZ with codes that had could scale up to 4 islands (32,768 cores). During the workshop, the participants tested the scaling capabilities on the whole system. Application experts from the LRZ, Intel and IBM were on site to resolve issues and assist in the perfor-

mance optimization. New techniques like fast startup were successfully tested which helped to reduce the startup time by a factor of 2-3. At the end of the workshop, 6 applications were successfully running on the full machine, while the other 8 applications managed to run on half of the system. The LRZ is already planning a follow-up workshop where the improvements and feedback from the experts will be tested.

Nerstrand: Fast Multi-Threaded Graph Clustering

Dominique W. LaSalle, George Karypis (University of Minnesota)

In this work we apply the multilevel paradigm to optimizing the modularity of a graph clustering on parallel shared memory architectures. We improve upon the state of the art by introducing new methods for effectively and efficiently coarsening graphs with power-law degree distributions, detecting an unknown number of communities, and for performing greedy modularity refinement in parallel. Finally, we present the culmination of this research, the clustering tool Nerstrand. In serial mode, Nerstrand runs in a fraction of the time of current methods and produces results of similar quality. When run with multiple threads, Nerstrand exhibits significant speedup without any degradation of clustering quality. Nerstrand works well on large graphs, clustering a graph with over 18 million vertices and 261 million edges in 18.3 seconds.

Loop Cluster Monte Carlo Simulation of Quantum Magnets Based on Global Union-Find Algorithm

Synge Todo (University of Tokyo), Haruhiko Matsuo (Research Organization for Information Science and Technology), Hideyuki Shitara (Fujitsu Limited)

A large-scale parallel loop cluster algorithm quantum Monte Carlo simulation on the K computer is presented. On 24,576 nodes of the K computer, one Monte Carlo update of the world-line configuration of the spin-1/2 antiferromagnetic Heisenberg chain with 2,621,440 spins at inverse temperature 310,690 is executed in about 8.62 seconds, in which global union-find cluster identification on a graph of about 1.1 trillion vertices and edges is performed. By combining the nonlocal global updates and the large-scale parallelization on the K computer, we have virtually achieved about 10^17-fold speedup. By using the highly parallelized quantum Monte Carlo algorithm, we have estimated successfully the magnitude of the first excitation gap and the antiferromagnetic correlation length of the spin-4antiferromagnetic Heisenberg chain for the first time as 0.000799 +/- 0.000005 and 10,400 +/- respectively.

Scalable Performance Analysis of Exascale MPI Programs through Signature-Based Clustering Algorithms

Amir Bahmani, Frank Mueller (North Carolina State University)

Exascale computing pose a number of challenges to application performance. Developers need to study application behavior by collecting detailed information with the help of tracing toolsets. But not only applications are scalability challenged, current tracing toolsets also fall short of exascale requirements for low background overheads since trace collection for each execution entity is becoming infeasible. One effective solution is to cluster processes with the same behavior into groups. Instead of collecting performance information from all individuals, this information can be collected from just a set of representatives. This work proposes a fast, scalable, signature-based clustering algorithm that clusters processes that exhibit the same execution behavior. Instead of prior work for statistical clustering metrics, it produces precise results without loss of events or accuracy. The proposed algorithm combines log(P) time complexity, low overhead at the clustering level, and it splits the merge process to make tracing suitable for exascale computing.

Xevolver: An XML-based Programming Framework for Software Evolution

Hiroyuki Takizawa, Shoichi Hirasawa, Hiroaki Kobayashi (Tohoku University)

In this work, we propose an extensible programming framework, named Xevolver, to separate system-specific optimizations from the original application code. The proposed framework exposes an abstract syntax tree (AST) in an XML data format to programmers. Hence, the programmers can adopt various XML-related technologies to transform, analyze, and visualize the application code. In this work, we use XSLT to define custom compiler directives for application-specific code transformations. By incrementally inserting the user-defined directives, a real application can be migrated to another system without significantly modifying the original code because custom code translation rules are written in an external XSLT file. We can change the behaviors of user-defined directives for individual systems by changing XSLT rules in the external file. Accordingly, we can evolutionally improve the application so as to have a high performance portability without messing up the original code.

LittleFe - The High Performance Computing Education Appliance

Charles Peck, Kristin Muterspaw (Earlham College), Mobeen Ludin, Aaron Weeden (Shodor Foundation), Skylar Thompson (University of Washington), Tom Murphy (Contra Costa College), Ivan Babic, Elena Sergienko (Earlham College)

Many institutions have little to no access to parallel computing platforms for in-class computational science or parallel and

distributed computing education. Key concepts, motivated by science, are taught more effectively and memorably on an actual parallel platform. LittleFe is a complete six node Beowulf style portable cluster. The entire package weighs less than 50 pounds, travels easily, and sets up in five minutes. LittleFe hardware includes multi--core processors and GPGPU capability, which enables support for shared and distributed memory parallelism, GPGPU parallelism, and hybrid models. By leveraging the Bootable Cluster CD project, LittleFe is an affordable, powerful, and ready-to-run computational science, parallel programming and distributed computing educational appliance.

Improving Utilization and Application Performance in High-Performance GPGPU Clusters with GPGPU Assemblies

Alexander Merritt, Naila Farooqui (Georgia Institute of Technology), Vishakha Gupta (Intel Corporation), Magda Slawinska, Ada Gavrilovska, Karsten Schwan (Georgia Institute of Technology)

High-end computing systems are increasingly characterized by machine configurations with nodes comprised of multiple CPUs and GPGPUs. Challenges using such heterogeneous machines include potential mismatches between applications' needs for and capacities in using these resources versus the fixed hardware configuration: (1) codes must be tuned and configured to match underlying hardware; (2) schedulers must map parallel jobs to efficiently use heterogeneous resources; and (3) mismatched workloads and machine configurations may leave hardware under-utilized.

Our work introduces `GPGPU Assemblies,' a software abstraction enabling the construction and maintenance of logical machine configurations, managed as `slices' of high-performance clusters. A novel method for extracting and applying application profiles enables our `Shadowfax' runtime to efficiently manage these abstractions. On Keeneland we demonstrate SHOC/S3D scaling up to 111 GPGPUs from a single machine, and 36%-72% reductions in queue completion times using both CPU and GPGPU LAMMPS codes on up to 64 nodes.

Optimizations of a Spectral/Finite Difference Gyrokinetic Code for Improved Strong Scaling Toward Million Cores

Shinya Maeyama (Japan Atomic Energy Agency), Tomohiko Watanabe (National Institute for Fusion Science), Yasuhiro Idomura, Motoki Nakata (Japan Atomic Energy Agency), Masanori Nunami, Akihiro Ishizawa (National Institute for Fusion Science)

Plasma turbulence is inherently multi-scale physics including electron and ion scales, which requires huge computations. In this work, we present optimizations of the gyrokinetic Vlasov simulation code GKV, which performs five-dimensional CFD calculation by using spectral and finite difference methods, on massively parallel platforms. First, segmented rank mapping on the three-dimensional torus network is advantageous to utilize the bi-section bandwidth and to reduce collisions of the messages, and therefore, to reduce communication costs. Second, uses of the collective communication optimized for the K computer and of the simultaneous communications on the Tofu interconnect further reduce the costs. Finally, computation-communication overlaps with pipelining methods effectively mask the communication cost. Thanks to the optimizations, GKV achieves excellent strong scaling up to ~600k cores with high parallel efficiency ~99.99994% on the K computer, and enables us to address first multi-scale plasma turbulence simulations employing the real ion-to-electron mass ratio.

Optimizing Built-To-Order BLAS while Considering Matrix Order

Robert Crimi (University Of Colorado Boulder)

Linear Algebra has always been a very useful tool in the field of Computer Science. Many techniques have been developed to help in the writing of graphics, internet search, and scientific computing algorithms. Such applications often use sequences of Basic Linear Algebra Subprograms (BLAS), and highly efficient implementations of those routines enable scientists to achieve high performance at little cost. The BLAS evolved from basic problems, such as dot products, to a large database including more complex problems, such as a matrix-matrix multiplication. Computer scientists apply tuning techniques to improve data locality and create highly efficient implementations of the Basic Linear Algebra Subprograms and LAPACK, enabling scientists to build high-performance software at reduced cost. However, because the BLAS are individually optimized, when piecing together different BLAS routines, users may not see the performance they desire. Thus, BTO is being developed to handle these situations.

Optimal Placement of Retry-Based Fault Recovery Annotations in HPC Applications

Ignacio Laguna, Martin Schulz, Jeff Keasler, David Richards, Jim Belak (Lawrence Livermore National Laboratory)

As larger high-performance computing systems are built, hardware fault recovery becomes a fundamental capability to make use of these systems. Retry-based recovery techniques have been proposed in which a code region is simply re-executed when a fault occurs. Although this technique requires code annotations, no previous work has investigated the optimal placement of these annotations in a program. In this poster, first, we will present our annotation-based model for fault recovery. Then, we will show our experimental evaluation (via fault injection) of how to place optimally annotations in a hydrodynamics mini application. We found that, contrary to our expectations, a simple scheme of protecting the main function works well for low fault rates. We also found that the optimal recovery method is rolling a few iterations back in the application's main loop and that the entire application can be protected with a maximum slowdown of 1.1.

Multi-Core Optimizations for Synergia and ART

Qiming Lu, James Amundson, Nick Gnedin (Fermi National Accelerator Laboratory)

We describe our recent work in optimizing the performance and scaling of Synergia and ART for multi-socket multi-core architectures including BlueGene/Q and GPUs. We show multiple hybridization and optimization options, including communication avoidance, interchangeable multi-threading kernel using OpenMP or CUDA for different hardware architectures, customized FFT, etc., each demonstrating much better scaling behavior than the pre-optimization code. By implementing the optimization techniques, we have extend strong scaling and peak performance by at least a factor of 2. We expect different optimization schemes to be optimal on different architectures. We have further tailored the code for BG/Q with optimized communication divider, redundant field solver, and FFT methods. The final code of Synergia scales up to 128K cores with over 90% efficiency running on Mira (BG/Q at Argonne).

Scalable Parallel Debugging via Loop-Aware Progress Dependence Analysis

Subrata Mitra (Purdue University), Ignacio Laguna, Dong H. Ahn, Todd Gamblin, Martin Schulz (Lawrence Livermore National Laboratory), Saurabh Bagchi (Purdue University)

Debugging large-scale parallel applications is challenging, as this often requires extensive manual efforts to isolate the origin of errors. For many bugs in a scientific application, where its tasks progress forward in a coordinated fashion, finding those tasks that progressed the least can significantly reduce the time to isolate the root- cause. In our poster, we present a novel run-time technique, the loop-aware progress-dependence analysis, which can improve the accuracy of identifying the least-progressed (LP) task(s). We extend AutomaDeD to detect LP task(s) even when the error arises on code with complex loop structures. Our evaluation shows that it accurately finds LP task(s) on several hangs on which the baseline technique fails. During the poster session, we will begin with a case that illustrates some of the challenges in accurately analyzing progress dependencies of MPI tasks executing within a loop, and then present key techniques needed to address these challenges.

Federated Computing for the Masses - Tackling Large-Scale Engineering Problems

Javier Diaz-Montes (Rutgers University), Yu Xi (Iowa State University), Ivan Rodero (Rutgers University), Jaroslaw Zola (Rutgers University), Baskar Ganapathysubramanian (Iowa State University), Manish Parashar (Rutgers University)

The complexity of many problems in science and engineering requires computational capacity greatly exceeding what average user can expect from a single computational center. This work explores the use of aggregated HPC resources to solve large-scale scientific problems. We show that it is possible to build a computational federation that is easy to use by end-users, elastic, resilient and scalable. We argue that the fusion of federated computing and real-life problems can be brought to average user if relevant middleware infrastructure is provided. To support our claims, we report on the use of federation of 10 geographically distributed HPC resources to perform a large-scale interrogation of the parameter space in the microscale fluid-flow problem. This problem is of great importance in many domains such as biological processing, guiding chemical reactions, and creating structured materials. As a result, we obtained the most comprehensive data on the effect of pillars on microfluid channel-flow.

Hybrid MPI/OpenMP/GPU Parallelization of XGC1 Fusion Simulation Code

Eduardo F. D'Azevedo (Oak Ridge National Laboratory), Jianying Lang (Princeton Plasma Physics Laboratory), Patrick H. Worley (Oak Ridge National Laboratory), Stephane A. Ethier (Princeton Plasma Physics Laboratory), Seung-Hoe Ku (Princeton Plasma Physics Laboratory), Choong-Seock Chang (Princeton Plasma Physics Laboratory)

By exploiting MPI, OpenMP, and CUDA Fortran, the FORTRAN fusion simulation code XGC1 achieves excellent weak scalability out to at least 18,624 GPU-CPU XK7 nodes, enabling science studies that have not been possible before.

XGC1 is a full-f gyrokinetic particle-in-cell code designed specifically for simulating edge plasmas in tokamaks. XGC1 was recently ported to and optimized on the 18,688 node Cray XK7 sited in the Oak Ridge Leadership Computing Facility, making use of both the 16-core AMD processor and the NVIDIA Kepler GPU on each node.

XGC1 uses MPI for internode and intranode parallelism, OpenMP for intranode parallelism, and CUDA Fortran for implementing key computational kernels on the GPU. XGC1 also uses the CPU and GPU simultaneously for these computational kernels. The optimized version achieves a four times speed-up over the original CPU-only version.

The Heuristic Static Load-Balancing Algorithm Applied to CESM

Yuri Alexeev, Sheri Mickelson, Sven Leyffer, Robert Jacob (Argonne National Laboratory), Anthony Craig (National Center for Atmospheric Research)

We propose to use the heuristic static load-balancing algorithm for solving load balancing problems in the Community Earth System Model (CESM), using fitted benchmark data, as an alternative to the current manual approach. The problem of allocating the optimal number of CPU cores to CESM components is formulated as a mixed-integer nonlinear optimization problem which is solved by using an optimization solver imple-

mented in the MINLP package MINOTAUR. Our algorithm was tested for the 1° and 1/8° resolution simulations on 163,840 cores of IBM Blue Gene/P where we consistently achieved well load balanced results. This work is a part of a broader effort by NCAR and ANL scientists to eliminate the need for manual tuning of the code for each platform and simulation type, improve the performance and scalability of CESM, and develop automated tools to achieve these goals.

Making GMRES Resilient to Single Bit Flips

James J. Elliott (North Carolina State University), Mark Hoemmen (Sandia National Laboratories), Frank Mueller (North Carolina State University)

The collective surface area and increasing density of components has led to an increase in the number of observed bit flips. This effort works towards rigorously quantifying the impact of bit flips on floating point arithmetic. We exploit norm bounds to ensure the upper Hessenberg entries constructed in the GMRES algorithm do not exceed theoretical limits. We also combine analytical modeling to show how scaling the inputs to these algorithms mitigate silent faults should they occur. We then combine our upper Hessenberg check, analytical model of bit flip perturbations, and the concept of sandbox reliability to show how scaling data can enhance existing fault tolerant algorithms, namely Fault Tolerant GMRES.

Designing and Auto-Tuning Parallel 3-D FFT for Computation-Communication Overlap

Sukhyun Song, Jeffrey K. Hollingsworth (University of Maryland)

We present a method to design and auto-tune a new parallel 3-d FFT code using non-blocking MPI all-to-all operations. Preliminary results show that we maximize computation-communication overlap, and execute 3-d FFT faster than the MPI-enabled FFTW library by up to 1.76x.

Scalable Ensemble Ab Initio Calculations Using the Fragment Molecular Orbital Method in GAMESS

Yuri Alexeev, Maricris Mayes (Argonne National Laboratory), Spencer Pruitt (Iowa State University), Graham Fletcher (Argonne National Laboratory), Dmitri Fedorov (National Institute of Advanced Industrial Science & Technology), Mark Gordon (Iowa State University)

A practical solution to achieving scalability on millions of cores is to employ 'ensemble' calculations in which multiple related jobs are executed concurrently. Molecular dynamics (MD) is a powerful technique that allows ensemble methods to be employed while providing detailed and accurate information on structural and dynamical properties of chemical systems. However, classical molecular dynamics simulations use empirical force fields that are obtained by parameterization. Thus, the accuracy of predicted properties depends on the system. To solve this problem we use ab initio quantum chemistry dynamics implemented in the software package, GAMESS. In order to make ab initio dynamics computationally tractable, supercomputers and linear scaling methods are necessary. Specifically, we use the Blue Gene/Q supercomputer, and the linear-scaling Fragment Molecular Orbital (FMO) method, together with ensemble methods, to expedite the calculations. In this poster, we describe how the ensemble strategy can be employed with FMO in GAMESS to do MD.

Introducing HPC & Multithreaded Computing to Middle School Girls Using Scratch

Russell Feldhausen, Scott Bell, Daniel A. Andresen (Kansas State University)

As part of a summer outreach program, we created an activity using Scratch that demonstrated HPC concepts. Our goal was to give students experience working with a multithreaded computer simulation by allowing them to adjust the number of threads used and the model resolution of the simulation. Students observed how those changes affect the time it takes the simulation to complete. Based on those results, they were able to determine the optimal number of threads for the hardware configuration used.

After participating in our activity, 27 of 37 students felt they can learn how to write computer programs, and 22 of 41 students showed interest in a job involving HPC. We plan to expand on this activity by giving students more opportunity to customize the simulation and gain experience working with simple programming concepts. We also plan to refine our survey questions to glean more information from students.

High Performance Computations Enable Scientific Discoveries: Adapting Computational Physics Codes to New Hardware

James F. Amundson, Philippe G. Canal, Donald J. Holmgren, Jim B. Kowalkowski, Qiming Lu, Marc Paterno, Saba Sehrish (Fermi National Accelerator Laboratory)

Physics research as we know it today would not be possible without computers. Evaluating fundamental physical theories at the smallest and largest physical scales requires the production of enormous quantities of numerical data. Designing, building, and commissioning highly complex accelerators and detectors requires the simulation of enormous event samples similar to those collected for an experiment. Breakthroughs in physics research depend on the use of leading-edge computing technologies. The Fermilab Scientific Computing Division together with university and laboratory collaborators is advancing the use of state-of-the-art computing platforms that include GP-GPUs and Blue Gene to enable scientific discoveries by evolving the software of four application toolkits. Here we

show the purpose, goals, and progress of this work within accelerator modeling (Synergia), computational cosmology (ART - Adaptive Refinement Tree), detector simulation (Geant4), and LQCD.

Real-Time Stochastic Optimization of Complex Energy Systems on High Performance Computers

Cosmin Petra (Argonne National Laboratory), Olaf Schenk (Universita della Svizzera Italiana), Mihai Anitescu (Argonne National Laboratory)

We present a scalable approach that computes in operationally-compatible time the energy dispatch under uncertainty for electrical power grid systems of realistic size with thousands of scenarios. We propose several algorithmic and implementation advances in our parallel solver PIPS-IPM for stochastic optimization problems. The new developments include a novel incomplete augmented multicore sparse factorization implemented within PARDISO linear solver and new multicore- and GPU-based dense matrix implementation. We also adapt and improve the interprocess communication strategy. PIPS-IPM is used to solve 24-hour horizon power grid problems with up to 1.95 billion decision variables and 1.94 billion constraints on "Titan" (Cray XK7) and "Piz Daint" (Cray XC30), where we observe very good parallel inefficiencies and solution times within a operationally defined time interval. To our knowledge, "real-time"-compatible performance on a broad range of architectures for this class of problems has not been possible prior to present work.

Building Scalable Data Analysis Platform for Metagenomics

Wei Tang, Jared Wilkening, Narayan Desai, Wolfgang Gerlach, Andreas Wilke, Folker Meyer (Argonne National Laboratory)

With the advent of the high-throughput sequencing technology, the management and analysis of genomic and metagenomic data has become a tremendous challenge. For example, MG-RAST, a metagenome annotation system serving large scientific communities worldwide, is experiencing an increasingly large amount of sequence data being submitted for data analysis—a situation that threatens to overwhelm efficient production. To address this situation, we have developed a data management system (Shock) and a workflow engine (AWE), which can collectively build a scalable data analysis platform for running bioinformatics applications in a distributed fashion. In this poster, we present the design and implementation of our new scalable data analysis platform and how they help to address the data deluge problem experienced by MG-RAST.

Mapping and Characterizing Global-Scale Human Settlements Using HPC

Dilip Patlolla, Anil Cheriyadat, Harini Sridharan, Vincent Paquit, Jeanette Weaver, Mark Tuttle (Oak Ridge National Laboratory)

Analytics derived from large-scale analysis of satellite image data is a key driver for many geospatial models including population mapping, risk modeling, and critical infrastructure assessment. Making informed national-level decisions based on analytics such as the spatial distribution of mobile home parks or the development rate of new constructions requires repeated processing of Peta-scale high resolution image data. The process often comprises computationally expensive steps like extracting, representing, and identifying pixel patterns that corresponds to the physical location and man-made structures. Such processes can easily saturate conventional CPUs.

Our settlement mapping system consists of computer vision and machine-learning techniques implemented on a multi-GPU architecture to extract, model, and interpret image data to characterize the spatial, structural, and semantic attributes of human settlements. It scales linearly with increase in the number of GPUs while delivering significant speedup enabling us to produce world scale human settlement mapping products at every 1-2 year's.

A Portable Approach for Repeatability Testing

Quan Pham, Tanu Malik, Ian Foster (University of Chicago)

Provenance-To-Use (PTU) is a tool from SOLE Framework that minimizes computation time during repeatability testing of computation-based experiments. PTU allows authors to assemble code, data, environment, and provenance of an initial reference execution into a single package that can be distributed to testers. The resulting package allows testers to view the provenance graph and specify, at a process and file level, nodes of the graph that they want for a partial deterministic replay-based, for example, on their compute, memory and I/O utilization as measured during the reference execution. Using the provenance trace, PTU guarantees that events are processed in the same order using the same data from one execution to the next. We show the efficiency of PTU for conducting repeatability testing of workflow-based scientific programs outweighs the overhead incurred during the reference executions.

PGAS Models Using an MPI Runtime: Design Alternatives and Performance Evaluation

Jeffrey Daily, Abhinav Vishnu, Bruce Palmer, Hubertus van Dam (Pacific Northwest National Laboratory)

In this poster, we explore the suitability of using MPI in designing a scalable PGAS communication subsystem. We focus on the RMA communication in PGAS models which typically includes get, put, and atomic memory operations. We perform an in-depth exploration of design alternatives based on MPI including using a semantically-matching interface such as MPI-RMA, as well as not-so-intuitive MPI interfaces such as two-sided, multi-threading, and dynamic process management. We propose a novel design which leverages a combination

of highly-tuned MPI two-sided semantics and an automatic, user-transparent split of MPI communicators to provide asynchronous progress. We implement the asynchronous progress ranks (PR)-based approach and other approaches within ComEx - Communication runtime for Exascale, a communication subsystem for Global Arrays. Our performance evaluation includes numerous benchmarks as well as a computational chemistry application. Our PR-based approach achieves a 2.17x speed-up on 1008 processors over the other MPI-based designs.

Implementation of the Colorado State University Global Cloud Resolving Model on the NSF Blue Waters System

Ross P. Heikes, David A. Randall, Celal S. Konor (Colorado State University)

We have ported key components of the Colorado State University (CSU) Global Cloud Resolving Model (GCRM) to the NSF Blue Waters system. With the model on Blue Waters there are two levels of parallelism to explore. The first is a coarse grain MPI based communication between computational cores. This parallelism works in conjunction with our global domain decomposition. We show scaling characteristics of the Blue Waters system to 40K cores, and include comparisons with other computer systems. The second level of parallelism is a fine scale parallelism which utilizes the NVIDIA Kepler accelerators. This loop based parallelism directly modifies the numerical operators used within the model using openACC directives. We show that the parallel efficiency of the accelerator strongly depends on the problem size, and have devised modifications to the model to better utilize the accelerators.

Pattern-Driven Node-Level Performance Engineering

Jan Treibig, Georg Hager, Gerhard Wellein (Erlangen Regional Computing Center)

Performance engineering has become an addition to traditional software engineering. However, in practice the chaotic approaches used in what was formerly called “performance optimization” are often transferred without change. We propose a systematic performance engineering process constructed around a diagnostic performance model. The performance model is a tool for the developer to acquire knowledge about the interaction of his code with a specific system architecture. In the poster we want to combine the process with the notion of “performance patterns”. Patterns provide concrete plans how to understand the observed performance, how to formulate a suitable model, and help decide on optimizations. We use a simplified view on computer architecture containing a suitable level of abstraction for getting an intuitive access to observed performance. Examples are presented for common performance patterns and how they can be detected. Finally we show the application of our performance engineering process on several case studies.

Highly Scalable Linear Time Estimation of Spectrograms - A Tool for Very Large Scale Data Analysis

Onkar Bhardwaj (Rennselaer Polytechnic Institute), Yves Ineichen, Costas Bekas, Alessandro Curioni (IBM Research - Zurich)

In many situations data analysis transform into eigenvalue problems. However, the era of big data means that dataset sizes render this problem practically intractable. The cubic complexity of dense methods and the limitation of iterative techniques to look deep into the interior of the spectrum at an acceptable cost, call for a new approach.

We present a close to linear cost method to estimate the spectrogram of a matrix, that is the density of eigenvalues in a certain unit of space. The spectrogram creates a compact graphical illustration of data matrices that can foster easier interpretation. This can be achieved by approximating the cdf of the eigenvalues and subsequently estimating the trace with help of a stochastic diagonal estimator.

We have designed and implemented a highly scalable implementation of our method, taking advantage of nested levels of parallelism that ultimately allow us to scale to massively parallel machines.

Test-Driven Parallelization of a Legacy Fortran Program

Damian W. I. Rouson (Stanford University), Hari Radhakrishnan (University of Cyprus), Karla Morris (Sandia National Laboratories), Sameer Shende (University of Oregon), Stavros C. Kassinos (University of Cyprus)

This poster describes the steps involved in modernizing a Fortran 77 turbulent flow model using the object-oriented (OO) and coarray parallel programming features of Fortran 2003 and 2008, respectively. OO programming (OOP) facilitates the construction of an extensible suite of model verification and performance tests that drive the development. Fortran’s new coarray parallel programming feature set facilitates evolving rapidly from a serial application to a parallel application capable of running on multicore processors and manycore accelerators in shared and distributed memory.

The modernization strategy involves 17 steps from automating the building and testing process to adding version control, transforming the source code, and analyzing the performance with the Tuning and Analysis Utilities (TAU) [1]. The verification tests ensure program correctness after each source-transformation step. The performance tests guard against execution-time regressions.

The ultimate steps of parallelization via the coarray syntax demonstrates the approach’s minimal impact on the source code.

High-Performance Design Patterns for Modern Fortran

Damian W. I. Rouson (Stanford University), Karla Morris (Sandia National Laboratories), Magne Haveraaen (University of Bergen), Jim Xia (IBM Canada), Sameer Shende (University of Oregon)

HPC hardware is experiencing a flood of new technologies that promise higher performance and better energy utilization. Continuously adapting HPC codes for multicore processors and manycore accelerators, however, drains human resources and prevents utilization of more cost-effective hardware.

We propose some design patterns for Fortran 2008 addressing this problem. Fortran's new coarray distributed data structures provide our backdrop along with the Fortran's side-effect-free, pure function capability. These features empower compilers to organize the computation on many cores with efficient communication. The presented design patterns support asynchronous evaluation of expressions comprised of parallel operations on distributed data structures. We demonstrate the design patterns on the tensor calculus expressions common in computational physics. A solver for the Burgers partial differential equation illustrates the approach's performance and scalability.

A Lightweight Data Location Service for Nondeterministic Exascale Storage Systems

Zhiwei Sun, Anthony Skjellum (University of Alabama at Birmingham), Lee Ward, Matthew Curry (Sandia National Laboratories)

We present LWDLS, a lightweight data location service designed for Exascale storage systems and geo-distributed storage systems. LWDLS provides a search-based data location solution, and enables free data placement, movement, and replication. In LWDLS, probe and prune protocols reduce topology mismatch, and a heuristic flooding search algorithm achieves higher search efficiency than pure flooding search while achieving comparable search speed and coverage.

LWDLS is lightweight and scalable in terms of having low overhead, high search efficiency, avoiding global state or periodic messages. LWDLS is fully distributed, and can be used in nondeterministic storage systems and in deterministic storage systems.

Extensive simulations modeling large-scale HPC storage environments provide representative performance outcomes, which are evaluated by metrics including search scope,search efficiency, and average neighbor distance. Results showed LWDLS is able to locate data efficiently with low cost of state maintenance in arbitrary network environments.We demonstrated the effectiveness of protocols and search algorithms.

Parallelizing Irregular Applications through the YAPPA Compilation Framework

Silvia Lovergine, Antonino Tumeo (Pacific Northwest National Laboratory), Oreste Villa (NVIDIA Corporation), Fabrizio Ferrandi (Politecnico di Milano)

Modern High Performance Computing (HPC) clusters are optimized for floating point intensive applications, and regular, localizable data structures. However, many emerging classes of scientific applications are irregular. Irregular applications are characterized by data sets difficult to partition, unpredictable memory accesses, unbalanced control flow and fine grained communication. Designing irregular applications results complex and requires significant programming effort. This work describes YAPPA (Yet Another Parallel Programming Approach), a compilation framework for the automatic parallelization of irregular applications on modern HPC systems, based on LLVM. We start by considering an efficient parallel programming approach for irregular applications on distributed memory systems. We then propose a set of transformations to reduce development and optimization effort.

Optimizing the Barnes-Hut Algorithm for Multicore Clusters

Junchao Zhang, Babak Behzad, Marc Snir (University of Illinois at Urbana-Champaign)

Nowadays all supercomputers are multicore clusters with one-sided communication support. It is important to explore new programming models for them. Partitioned global address space (PGAS) languages have drawn much attention in recent years. However, there are few application mapping studies. In this paper, through the Barnes-Hut algorithm, we demonstrate a PGAS + X programing paradigm specifically targeting multicore clusters.

The novelty is that we integrate intranode multithreading with internode one-sided communication. we decompose computation into tasks and hide network latency by descheduling tasks that are blocked on a remote access and reusing the core to run another, ready to execute task. We spawn multiple threads per process, which do either communication or computation. We show how to manage tasks and do thread synchronization. We are in progress of designing runtime abstractions.

We show details of the design, particularly the task management. We use experiment results to show benefits gained.

Generating Customized Eigenvalue Solutions Using Lighthouse

Luke Groeninger, Ramya Nair (University of Colorado Boulder), Sa-Lin Cheng Bernstein (University of Chicago), Javed Hossain, Elizabeth R. Jessup (University of Colorado Boulder), Boyana Norris (University of Chicago)

Eigenproblems are a fundamental part of science and engineering. Indeed, at least one thing you have done today was made possible in part by solving an eigenproblem. Whether it was using a search engine or speech recognition software, driving across a bridge, or flying on an aircraft, eigenproblems are now responsible for many everyday services on which you depend. Scientists and engineers across many disciplines must thus rely on high-performance numerical libraries to solve eigenproblems quickly and efficiently. However, exploiting the latest high-performance computational techniques for solving eigenproblems remains difficult. Our work addresses this problem by enabling users to find and use customized routines for solving eigenproblems without the steep learning curve traditionally associated with HPC libraries. We present new contributions in three areas: taxonomies of dense and sparse eigensolvers available in LAPACK and SLEPc; user interfaces for searching the taxonomies; and performance-based recommendation of solvers based on machine-learning analysis.

Predictions of Large-scale QMCPack I/Os on Titan Using Skel

Matthew Wezowicz, Michael Matheny, Stephen Herbein (University of Delaware), Jeremy Logan, Jeongnim Kim, Jaron Krogel, Scott Klasky (Oak Ridge National Laboratory), Michela Taufer (University of Delaware)

Large-scale applications such as the Quantum Monte Carlo code QMCPack face a major challenge to preserve their scalability when fine-grained data are gathered and stored to disk or used for in-situ analysis. Using an IO framework such as ADIOS allows us to address the trade-off by deploying different IO methods with little code modifications. Still the search for the most suitable methods and settings can be challenging. Ideally such a search can be conducted on a tool such as Skel that allows us to decouple the IO from the computation in real applications and to tune the IO on real supercomputers. To allow tuning for applications with unbalanced generation of IO, we extended Skel to integrate the IO variability shown real world applications. We validated the Skel results against actual QMCPack IO performance data and showed that Skel results can be used to predict IO and steer choices for applications.
Speedup and Numerical Evaluation of Multiple-Precision Krylov Subspace Method Using GPU Cluster for Large-Sparse Linear System

Yuta Hirokawa, Taku Itoh (Tokyo University of Technology), *Hiroto Tadano (University of Tsukuba), Soichiro Ikuno (Tokyo University of Technology)*

Implementing multiple-precision Krylov subspace method on GPU cluster for a large-sparse linear system is investigated, and the method is numerically evaluated. It is well known that a number of iteration of Krylov subspace method depend on accumulation errors, and the error may affect the calculation results. In order to settle these issues, the multiple precision operation Krylov subspace method is implemented on GPU cluster using GNU Multiple Precision Arithmetic Library (GMP) and CUDA Multiple Precision Arithmetic Library (CUMP), and the method is parallelized to get high performance. The result of computation shows that the variable preconditioned Bi-CGSTAB on GPU cluster is up to 16.38 times faster than that of CPU with OpenMP.

Mitigating System Noise With Simultaneous Multi-Threading

Eli Rosenthal (Brown University), Edgar A. Leon, Adam T. Moody (Lawrence Livermore National Laboratory)

System noise on commodity clusters significantly impacts the scalability of scientific applications. As the number of nodes in a cluster grows, this impact outweighs the benefits of the increased system capability. In this work, we propose using Simultaneous Multi-Threading (SMT) to mitigate system noise, and we demonstrate its effectiveness using Fixed Work Quantum (FWQ) micro-benchmarks. This technique relies on using the floating-point and integer units of a core (shared SMT resources) for system processing when an application is stalled waiting for an out-of-core operation. Unlike core-specialization, which dedicates a core for system processing, our approach allows an application to use all cores of a node and proves significantly more effective in reducing system noise.

Ball Computer: a 3D Wireless Grid for Massively Parallel Computers

Amir Mansoor Kamali Sarvestani, Christopher Chrispin-Bailey, Jim Austin (University of York)

To date, wireless connections have not been good choices for the interconnect network of a massively parallel computer. This is mainly because they are considered to have both more energy demands and have lower data rates compared to wireline solutions. This paper investigates if (and to what extent) that assumption is still valid. Capacitive coupling, inductive coupling and radio waves are reviewed. They are compared with each other and with wireline technologies on factors including data rate, energy consumption and occupied area. The question is how close we are to an efficient wireless parallel computer. For an evaluation case we utilize a simulated 3D wireless grid for a parallel computer. The model is used to test the viability of a concept 3D wireless parallel computer called A Ball Computer (ABC). Some simulation results on a range of task models, data sizes, network parameters and network sizes are presented.

A Transactional Model for Fault-Tolerant MPI for Petascale and Exascale Systems

Amin Hassani, Anthony Skjellum (University of Alabama at Birmingham), Ron Brightwell (Sandia National Laboratories)

Fault-Aware MPI (FA-MPI) is a novel approach to provide fault-tolerance through a set of extensions to the MPI Standard. It employs a transactional model to address failure detection, isolation, mitigation, and recovery via application-driven policies. This approach allows applications to employ different fault-tolerance techniques, such as algorithm-based fault tolerance (ABFT) and multi-level checkpoint/restart methods. The goal of FA-MPI is to support fault-awareness in MPI objects and enable applications to run to completion with higher probability than running on a non-fault-aware MPI. FA-MPI leverages non-blocking communication operations combined with a set of TryBlock API extensions that can be nested to support multi-level failure detection and recovery. Managing fault-free overhead is a key concern as well.

Prov-Vis: Large-Scale Scientific Data Visualization using Provenance

Felipe Horta, Jonas Dias, Renato Elias (COPPE/UFRJ), Daniel Oliveira (Fluminense Federal University), Alvaro Coutinho, Marta Mattoso (COPPE/UFRJ)

Large-scale experiments on computational engineering and science rely on compute-intensive tasks chained through a dataflow. These experiments may be modeled as scientific workflows, to ease the experiment management and take advantage of provenance data. Monitoring workflow execution is the task of verifying the status of the execution at specific points to discover if anything odd has happened. Traversing provenance data at runtime can support this monitoring, so that users can just stop or re-execute some tasks. However, most of the workflow systems execute workflows in an "offline" way. Prov-Vis is a scientific data visualization tool for large-scale workflows that uses runtime provenance queries to organize and aggregate data helping to follow the steps of the workflow and the produced partial results. A parallel numerical simulation workflow was executed on a supercomputer while Prov-Vis displayed, on a tiled-wall, visualizations of simulation steps chosen based on runtime provenance queries.

Task Mapping for Non-Contiguous Allocations

David P. Bunde, Johnathan Ebbers, Stefan P. Feer (3M Health Information Systems), Vitus J. Leung (Sandia National Laboratories), Nickolas W. Price (Knox College), Zachary D. Rhodes (Allstate Corporation), Matthew Swank (Knox College)

We examine task mapping algorithms for systems that allocate jobs non-contiguously, such as the Cray X series. Several studies have shown that task placement affects job running time for both contiguously and non-contiguously allocated jobs. Our focus is on jobs with a stencil communication pattern. We use experiments on a Cray XE to evaluate novel task mapping algorithms as well as some adapted to this setting. This is done with the miniGhost miniApp which mimics the performance of CTH, a shock physics application. Our strategies improve running time by as much as 35% over a baseline strategy.

Solution of the Time-Dependent Acoustic-Elastic Wave Equation on a Heterogeneous, Coprocessor-Enabled Supercomputer

Jesse Kelly, Hari Sundar, Omar Ghattas (University of Texas at Austin)

Modern supercomputers increasingly utilize accelerators and coprocessors. A key challenge in using such heterogeneous systems has been achieving load balance such that neither the CPU nor the coprocessor is left idle. Traditional approaches have offloaded entire computations to the coprocessor, resulting in an idle CPU, or have opted for task-level parallelism requiring large data transfers between the CPU and the coprocessor. True work-parallelism has been hard as coprocessors cannot directly communicate with other CPUs (besides the host) or other coprocessors. We present a new nested partition scheme to overcome this problem. By partitioning the work assignment on a given node asymmetrically into boundary and interior work, and assigning the interior to the coprocessor, we are able to achieve excellent efficiency while ensure proper utilization of both the CPU and the coprocessor. The problem used for evaluating such partitioning is an hp-discontinuous Galerkin spectral element method for coupled acoustic-elastic wave propagation.

Running A Seismic Workflow Application on Distributed Resources

Scott Callaghan, Philip Maechling (University of Southern California), Karan Vahi, Gideon Juve, Ewa Deelman (Information Sciences Institute), Yifeng Cui, Efecan Poyraz (University of California, San Diego), Thomas H. Jordan (University of Southern California)

In this poster, we present an approach to running workflow applications on distributed resources, including systems without support for remote job submission. We show how this approach extends the benefits of scientific workflows, such as job and data management, to large-scale applications on open-science HPC resources such as Blue Waters, Stampede, and USC HPCC. We demonstrate this approach with SCEC CyberShake, a physics-based seismic hazard application, to run over 470 million tasks via 18,000 jobs submitted to Blue Waters and Stampede.

A Synthesis Approach for Mapping Irregular Applications on Reconfigurable Architectures

Vito Giovanni Castellana, Antonino Tumeo (Pacific Northwest National Laboratory), Fabrizio Ferrandi (Politecnico di Milano)

Emerging applications such as bioinformatics and knowledge discovery algorithms are irregular. They generate unpredictable memory accesses and are mostly memory bandwidth bound. Several efforts are looking at accelerating these applications on hybrid architectures, which integrate general purpose processors with reconfigurable devices. Some solutions include custom-hand tuned accelerators on the reconfigurable logic. Hand crafted accelerators provide great performance benefits, but their development time often discourages their adoption. We propose a novel High Level Synthesis approach, for the automatic generation of adaptive custom accelerators, able to manage multiple execution flows. Our approach supports multiple, multi-ported and distributed memories, and atomic operations. It features a memory interface controller, which maps unpredictable memory access requests to the corresponding memory ports, while managing concurrency. We present a case study on a typical irregular kernel, the Graph Breadth First search, evaluating performance tradeoffs when varying the number of memories and the number of concurrent flows.

An Improved Parallel Iterative Algorithm for Stable Matching

Colin R. White (Amherst College), Enyue Lu (Salisbury University)

We propose a parallel algorithm for finding a stable matching that converges in linear average time using n^3 processors. The algorithm is based on the Parallel Iterative Improvement (PII) algorithm, which finds a stable matching with approximately a 90% success rate, with an average time of O(n log n) using n^2 processors. Our algorithm uses a smart initialization method that decreases the number of iterations to find a stable matching, and also applies a cycle detection method to find a stable matching based on patterns in the preference lists. Both methods decrease the number of times it fails to find a stable matching by three orders of magnitude, and when combined, the chance of failure is less than 10^-7.

Asynchronous PDE Solver for Computing at Extreme Scales

Aditya Konduri, Diego Donzis (Texas A&M University), Torsten Hoefler (ETH Zurich)

Scalability of codes at extreme scales is a critical issue of current relevance. At extreme levels of parallelisms, communication between processing elements(PEs) could represent a substantial running time, resulting in substantial waste in computing cycles. We investigate a novel approach based on common finite-difference schemes for partial differential equations(PDEs) in which computations are done locally with values already available at each PE, without waiting for updated data from neighboring PEs. No synchronization among cores is enforced and computations proceed regardless of messages status. This drastically reduces processor idle times, improving computation rates and scalability. We show that accuracy of common numerical schemes is reduced when asynchronicity is present. We derive new schemes that can maintain the accuracy under asynchronous conditions. These new schemes were implemented and tested. Performance is evaluated through statistics of different measures of asynchronicity. A new implementation of RDMA communications is shown to provide significant advantages.

Efficient Data Compression of Time Series of Particles' Positions for High-Throughput Animated Visualization

Katsumi Hagita (National Defense Academy), Takaaki Takeda (VASA Entertainment, Inc), Tsunehiko Kato (Hiroshima University), Hiroaki Ohtani (National Institute for Fusion Science), Seiji Ishiguro (National Institute for Fusion Science)

We tried to improve data-reading throughput by data compression on animated visualization of time series data generated from simulations of huge particle systems. It is considered that data compression is expected to reduce problems related to relatively slow reading from storage device. In the present paper, we examined efficient data compression schemes for huge time series data of plasma particle simulation. For example, to read positions of 300,000 particles with double precision text, required Read Throughput is almost equal to that of recent SSD. Thus, we considered data compression enables visualization of larger system. For smooth animations, compression of trajectories in time order is effective, because correlation of positions of a certain particle in time orders is higher than spatial correlation. Especially, lossy compression with polynomial functions in time order shows good performances. We called this scheme as TOKI (Time-Order, Kinetic, and Irreversible) compression.

ACM Student Research Competition Posters

The ACM Student Research Competition Posters include 24 student-developed posters where students compete for the best graduate and undergraduate student poster awards. Finalists will present their work on Wednesday, November 20, at 3:30pm in Room 404 of the convention center, and the winners will be announced at the SC13 award ceremony. The first place graduate and undergraduate students will go on to compete in the 2014 ACM SRC Grand Finals, sponsored by Microsoft Research.

High Performance Computing for Irregular Algorithms and Applications

Oded Green (Georgia Institute of Technology)

Irregular algorithms such as graph algorithms, sorting, and sparse matrix multiplication present numerous programming challenges that include scalability, load balancing, and efficient memory utilization. The advent of Big Data increases the challenges of analyzing massive data sets such as Facebook and Twitter which change constantly. Thus, the continuous updating of the analytics of such data sets proves to be a big challenge. One such solution has been through the usage of special HPC systems like the Cray XMT. My current research focuses upon creating new algorithms for irregular algorithms with a particular emphasis on graph analytics which allow monitoring of different types of networks. For example, Betweenness Centrality (BC) is an important analytic used to find key players in social networks. However, computing BC usually requires more time than the characteristic time in which the graph is updated thus making the entire process virtually useless.

Power and Performance Modeling for High Performance Computing Algorithms

Jee Choi (Georgia Institute of Technology)

The overarching goal of this research is to enable algorithm designers to create more energy efficient algorithms by providing the means of analyzing the relationship between time, energy and power on real systems.

Firstly, we provide a simple analytical cost model for energy and power. This model expresses energy and power for algorithms using a small set of simple parameters (e.g., memory bandwidth, number of FLOPs). We validate our model using highly optimized microbenchmarks and a fine-grained power measurement tool on state-of-the-art systems.

Secondly, we iteratively refined our model and extend its validation on over a dozen "candidate compute-node building blocks" for future high performance systems, including the latest x86, GPUs, and mobile SoCs. The purpose of this study is to provide a precise analytical characterization of abstract algorithmic regimes where one building block may be preferable to others.

Together, they provide the means to re-design algorithms for energy efficiency.

Parallel Algorithms for Large-Scale Graph Clustering on Distributed Memory Architectures

Inna Rytsareva (Washington State University)

Clustering is one of the advanced analytical functions that has an immense potential to transform the knowledge space when applied particularly to a data-rich domain such as computational biology. The computational hardness of the underlying theoretical problem necessitate use of heuristics in practice. In this study, we evaluate the application of a randomized sampling-based heuristic called shingling on unweighted biological graphs and propose a new variant of this heuristic that extends its application to weighted graph inputs and is better positioned to achieve qualitative gains on unweighted inputs as well. We also present parallel algorithms for this heuristic using the MapReduce paradigm. Experimental results on subsets of a medium-scale real world input (containing up to 10.3M vertices and 640M edges) demonstrate significant qualitative improvements in the reported clustering, both with and without using edge weights. Furthermore, performance studies indicate near-linear scaling on up to 4K cores of a distributed memory supercomputer.

local_malloc(): malloc() for OpenCL __local memory

John Kloosterman (Calvin College)

One of the complexities of writing kernels in OpenCL is managing the scarce per-workgroup—local memory on a device. For instance, temporary blocks of—local memory are necessary to implement algorithms like non-destructive parallel reduction. However, all—local memory must be allocated at the beginning of a kernel, and programmers are responsible for tracking which buffers can be reused in a kernel. We propose and implement an extension to OpenCL that provides a malloc()-like interface for allocating workgroup memory. This extension was implemented by using an extension to the Clang compiler to perform a source-to- source transformation on OpenCL C programs.

Handling Datatypes in MPI-3 One Sided

Robert Gerstenberger (Chemnitz University of Technology)

With the rise of modern interconnects offering direct remote memory access (RDMA), one sided programming becomes increasingly popular. The MPI-3.0 standard specifies a programming interface that leverages these features directly. Fast one-sided MPI (foMPI) was introduced as a reference implementation to provide highest performance and minimal overheads.

Many applications have to communicate non-consecutive data. In message passing each partner specifies their own MPI datatype (a memory layout), while for one sided communication one process defines both types (for the origin and the target process), which requires different kinds of optimizations for truly one sided implementations. In this work three different strategies are proposed to handle the combination of those MPI datatypes and implement them in foMPI. Those strategies were evaluated with DDTBench, a set of data access pattern micro-applications, which were extended to handle one sided communication. Some cases improved up to a magnitude.

Towards Co-Evolution of Auto-Tuning and Parallel Languages

Ray S. Chen (University of Maryland)

There exists a gap between the massive parallelism available from today's HPC systems and our ability to efficiently develop applications that utilize such parallelism. Several emerging programming languages aim to narrow this gap through high-level abstractions of data and/or control parallelism. However, this requires the compiler to make performance affecting decisions based on information unavailable at compile time, such as workload complexity or thread distribution. By evolving auto-tuning along with these languages, such decisions could be deferred until run time when it's possible to test for optimal values.

In this poster, we present our vision of co-evolution by pairing the Chapel programming language with the Active Harmony auto-tuning framework. We include results of our co-evolution experiments which demonstrate the ability to automatically improve the performance of a proxy application over the default by nearly 25%.

Graphical Analysis of I/O in High Performance Computing

Denver Coker (Oak Ridge National Laboratory)

In HPC, full simulation application jobs can be time consuming and expensive to run. To analyze I/O performance a skeletal application is used to monitor only the I/O, thereby making the simulation less costly in many ways. After conversion, the output is then imported into graphical analysis software such as R for analysis. The script I have written takes all possible combinations of the variables from the skeletal application, and prints only the significant graphical combinations. Also the script will automatically color the graph intelligently as well. This method of data analysis is much more efficient since every graphical combination is not outputted but only the ones the user is interested in.

Fast Prediction of Network Performance: k-packet Simulation
Nikhil Jain (University of Illinois at Urbana-Champaign)
Prediction of application communication performance on supercomputer networks, without doing actual runs, is useful for a variety of what-if analyses: how does the performance change with different task mappings, studying the performance of future networks etc. A significant amount of research is conducted to study these aspects because the scalability of many applications is adversely affected by communication overheads. Flit-level and packet-level simulations have been shown to be useful for prediction, but they are inefficient and such simulations are very slow even for moderate-sized networks. We propose a new simulation methodology that relies on increasing the granularity of simulation to k-packets and use of simple heuristics for predicting the state of the network. Preliminary results show that the proposed approach accurately models network behavior, and is orders of magnitude faster than the previous methods: up to two minutes per prediction for 16,384 cores of Blue Gene/Q using the given benchmarks.

UDT Based Application Performance for High Speed Data Transfer

Joshua Miller (University of Chicago)

Over high performance wide area networks, TCP has been shown to reach a data transfer bottleneck before UDP. UDT is a reliable UDP based application level protocol designed to circumvent this limit. Despite increased data transfer speeds, the lack of UDT based file transfer tools has prevented its wide spread use within the scientific community.

Our purpose has been to create and test UDT based applications for bulk data transfer. This suite (UDStar) includes a wrapper for the stand alone file transfer tools rsync and scp as well as a UDT based netcat tool. Inclusion as a prefix to standard rsync and scp commands, UDStar minimizes the modification of existing work flows to aid its adoption into the scientific community.

Currently, UDStar is used to sync NASA satellite and ENCODE data sets. This poster presents the performance increases for transfers over 10G wide area networks.

Enabling Low I/O-Cost Fine-Grained Checkpointing in QMCPack

Michael Matheny (University of Delaware)

As simulation size increases, the likelihood of a node crash and the cost of checkpointing increases. This is more of a problem as flops increase exponentially while I/O speeds increase at a slower rate. The cost of saving data is even higher when checkpointed data is used for in-situ or in-transit analyses such as for visualization or for atom proximity identifications.

QMCPack is one such example that is limited by its checkpointing implementation. QMCPack uses serial HDF5 to write checkpointing data to disk. As the number of processes increase, the associated I/O cost grows significantly. Preliminary results showed that even for simple simulations, the time spent in I/O associated to a checkpointing could make up 90% of the execution time on systems with large numbers of nodes. In this poster we use ADIOS to tackle the problem of enabling more efficient fine-grained checkpointing in QMCPack while limiting the I/O overhead.

Active-Testing and Large Scale Non-determinism Control Studies with Programmable Schedulers

Bruce Bolick (University of Utah School)

Active-testing methods for hybrid concurrency models are essential, but poorly understood, and not yet demonstrated at scale. Today's active-testing approaches run tests on specific platforms for specific configurations. They are hard to modify once they have been written, and offer little flexibility.

This research investigates how active-testing approaches can be created to be effective in the modern setting of hybrid concurrency. We are re-implementing a scheduler for a relatively mature active-tester for MPI in Scheme/Racket, a modern higher order functional language, in order to provide elegance, simplicity and extensibility.

I will demonstrate how such a scheduler empowers large scale non-determinism control studies, facilitates search-bounding, and provides flexibility in formal verification of HPC systems.

Enabling Fine-grained Gathering of Scientific Data in QMCPack Simulations on Titan

Stephen N. Herbein (University of Delaware)

Traditional petascale applications, such as QMCPack, can scale their computation to completely utilize modern supercomputers like Titan but cannot scale their I/O. Because of this lack of scalability, scientists cannot save data at the granularity level they need to enable scientific discovery.

In this work, we tackle the problem of increasing the granularity of data collected from QMCPack simulations without increasing I/O overhead or compromising the simulation's scalability. Our solution relies on the redesign of the QMCPack algorithms to gather fine-grained information and the integration of the ADIOS API to help select effective IO methods without major code changes.

Results presented in the poster outline how we can increase the quality of the scientific knowledge extracted from Titian simulations of QMCPack while keeping the I/O overhead below 10%.

Theory, Meet Practice: Challenges in Applying Failure Prediction on Large Systems

Ana Gainaru (University of Illinois at Urbana-Champaign)

As the size of supercomputers increases, so does the probability of a single component failure within a time frame. Checkpoint-Restart, the classical method to survive application failures, faces many challenges in the Exascale era due to frequent and large rollbacks. A complement to this approach is failure avoidance, by which the occurrence of a fault is predicted and proactive measures are taken. With the growing complexity of extreme scale supercomputers, the act of predicting failures in real time becomes cumbersome and presents a couple of challenges not encountered before. This work is nearly complete and presents key issues I have encountered when applying online failure prediction on the Blue Waters system. The overhead of combining fault prediction and checkpointing on smaller and large scale systems will also be reported. The results give insights on the challenges in achieving an effective fault prevention mechanism for current and future HPC systems.

Structure-Aware Parallel Algorithm for Solution of Sparse Triangular Linear Systems

Ehsan Totoni (University of Illinois at Urbana-Champaign)

Solution of sparse triangular systems of linear equations is a performance bottleneck in many methods for solving more general sparse systems. In both direct methods and iterative preconditioners, it is used to solve the system or refine the solution, often across many iterations. Triangular solution is notoriously resistant to parallelism, however, and existing parallel linear algebra packages appear to be ineffective in exploiting much parallelism for this problem. We develop a novel parallel algorithm based on various heuristics that adapts to the structure of the matrix and extracts parallelism that is unexploited by conventional methods. By analysis and reordering operations, our algorithm can even extract parallelism in some cases where most of the nonzero matrix entries are near the diagonal. We describe the implementation of our algorithm in

Charm++ and MPI and present promising results on up to 512 cores of BlueGene/P, using numerous sparse matrices from real applications.

A Distributed-Memory Fast Multipole Method for Volume Potentials

Dhairya Malhotra (University of Texas, Austin)

We present a Fast Multipole Method (FMM) for computing volume potentials and use them to construct spatially-adaptive solvers for the Poisson, Stokes and Helmholtz problems. Conventional N-body methods apply to discrete particle interactions. With volume potentials, one replaces the sums with volume integrals. We present new near interaction traversals and incorporate a cache optimized for interaction traversal. Finally, we use vectorization, including the AVX instruction set on the Intel Sandybridge architecture to get over 50% of peak floating point performance. We use task parallelism to employ the Xeon Phi on the Stampede platform at the Texas Advanced Computing Center (TACC). We achieve over 550Gflop/s of double precision performance on a single node. Our largest run on Titan at ORNL took 7.8 secs on 16K nodes for a problem with 74E+9 unknowns for a highly nonuniform particle distribution.

Distributed Algorithms for Aligning Massive Networks

Arif Khan (Purdue University)

Given two graphs, the network alignment (NA) problem is to find the best one-to-one mapping between the vertices of one graph to those in the other by maximizing the number of overlapped edges. NA is an important problem with several applications in bioinformatics, computer vision and ontology matchings. It is an NP-hard problem, and solutions are heuristic and iterative in nature. Our algorithm is based on belief propagation and approximate weighted matching. Memory intensive requirements necessitate a distributed-memory implementation to solve large problems. A combination of sparse matrix operations and combinatorial algorithms makes network alignment a challenging problem. Our implementation is based on a special runtime system, Graph Multi-Threaded (GMT), designed to address irregular computation and memory accesses on distributed architectures. We demonstrate the utility of this approach by solving large problems that were previously not possible on shared-memory systems, and show strong scaling with 8x improvement on 64 nodes.

Evaluating Named Data Networking For Large Scientific Data

Susmit Shannigrahi (Colorado State University and Lawrence Berkeley National Laboratory)

With growing volume of scientific data, data management and retrieval are becoming increasingly complicated. Named Data Networking (NDN) is a potential next generation Internet architecture. In this preliminary work, we show that Named Data Networking (NDN) offers unique properties that can reduce complexities involving large scientific data. We discuss these built in properties of NDN for reducing the need for catalogs, complex middleware or applications. We then support these observations with results from experiments done using the ESNet testbed and real scientific data.

We have used a comparatively small dataset and simple testbed topology to demonstrate NDN's capabilities. We plan to extend this work for real scientific analysis involving larger datasets and different network topologies.

This poster focuses on demonstrating how NDN is useful for large scientific data. It also shows NDN's ability to reduce middleware and complex applications and at the same time, offer improved performance.

Inter-Application Coordination for Reducing I/O Interference
Sagar Thapaliya (University of Alabama at Birmingham)
Scientific applications running on HPC systems share common I/O resources, including parallel file systems. This sharing often results in inter-application I/O interference, leading to degraded application I/O performance. In this work, we study the benefits of coordinating between applications to reduce I/O interference and develop methods to perform this coordination. Our results with micro-benchmarks show complete reduction of I/O interference with small cost for applications.

Analysis and Optimization of an OrangeFS Block Driver for QEMU/KVM

Timothy J. Scott (Clemson University)

QEMU/KVM is an important part of many cloud computing stacks and is a critical element in the OpenStack consortium. This project evaluates the block driver written using a standard POSIX interface that leverages libofs, a POSIX interface that uses optimized calls in the OrangeFS usrint library. Analysis of performance includes measures and analysis of performance variation with different cache and striping configurations. This project seeks to meet two primary goals. First is to measure performance gains that will allow OrangeFS to be positioned as a backend to a datacenter or cloud infrastructure and rapidly serve images as new Virtual Machine (VM) instances are launched. Second is to serve as a proof-of-concept usage of new functionality in libofs and the performance gains available when using libofs to bypass the OrangeFS Linux kernel module. Performance will ultimately be compared between other QEMU networked filesystem block drivers including GlusterFS and NBD.

Fusion Active Storage for Write-Intensive Big Data Applications

Gregory A. Thorsness (Texas Tech University)

Many HPC applications have become highly data intensive due to the substantial increase of both simulation data generated from scientific computing models and instrument data collected from increasingly large-scale sensors and instruments.

These applications transfer large amounts of data between compute nodes and storage nodes, which is a costly and bandwidth consuming process. The data movement often dominates the applications[1] run time. This study investigates a new Fusion Active Storage System (FASS) to address the data movement bottleneck issue specifically for write-intensive big data applications. The FASS enables a paradigm that moves write-intensive computations to storage nodes, generates and writes data in place to storage devices. It moves computations to data and avoids the data movement bottleneck on the data path. The FASS has an advantage of minimizing data movements and can have an impact on big data applications.

Enabling Efficient Intra-Warp Communication for Fourier Transforms in a Many-Core Architecture

Carlo del Mundo (Virginia Tech)

Shuffle, a new mechanism in NVIDIA GPUs that allows for direct register-to-register data exchange within a warp, aims to reduce the shared memory footprint for data communication. Despite vendor claims on its efficacy, the mechanism is poorly understood with few works demonstrating performance improvement. Therefore, we seek to characterize the behavior of shuffle and provide insight into optimizing applications with intra-warp communication.

We evaluated the efficacy of the shuffle mechanism in the context of matrix transpose as part of the communication stage in a 1D FFT code. Our study indicates that refactoring algorithms to fit the shuffle paradigm requires careful co-design between software and hardware. In particular, algorithmic decisions should avoid CUDA local memory allocation and usage at all costs. Overall, our optimized shuffle version accelerates matrix transpose by up to 44% with an overall application speedup of 1.17-fold for a 256-point FFT.

Accurate Change Point Detection in Electricity Market Analysis

Will Gu (Lawrence Berkeley National Laboratory)

Electricity is a vital part of our daily life; therefore it is important to detect any irregularities such as those during the California Electricity Crisis of 2000 and 2001. Many advanced machine learning algorithms exist to detect anomalies, but they are computationally expensive and could not be applied on large data sets such as those from electricity markets. In this work, we develop a strategy to accelerate the computation of the Gaussian Process (GP) for financial time series within the framework of a Change Point Detection (CPD) process, reducing its computational complexity from O(N^5) to amortized O(N^2). We apply this fast algorithm to correlate the change points detected with the known events during the California Electricity Crisis. This calculation would have been impossible without the fastest supercomputers.

Scalable Client-Server Visualization of Callpath Traces for Large-Scale Parallel Executions

Philip A. Taffet (Rice University)

Scalable performance analysis tools are important in producing large-scale parallel executions. Sampling, typically far more scalable than instrumentation, still produces quantities of data that are unmanageable by current post-mortem analysis tools when analyzing large-scale executions. Hpctraceserver, a new addition to the HPCToolkit performance tools, improves the scalability of hpctraceviewer, a GUI post-mortem analysis tool, by extending it with a client-server model. Hpctraceserver, an MPI application, computes projections in parallel of the traces to fit the client's window size with sampling and streams the results to the client over a TCP connection. Using the server allows hpctraceviewer to render views from multi-gigabyte databases in fewer than 10 seconds, time reductions of nearly 90%. Using hpctraceserver also removes the need to transfer large quantities of data in order to analyze it. The increased scalability and performance allow interactive analysis and bottleneck detection in large applications.

Efficient Datarace Detection in a Structured Programming Language

Trudy Lynn Firestone (University of Utah)

Race detection is increasingly important in the world of Computer Science as non-specialized computer scientists are required to write more parallel programs because of slower CPU clocks with only an increasing number of cores to compensate. As incorrect parallel programs often contain dataraces, detection of these conflicts is a necessary debugging tool. Due to well-studied difficulty involving race detection in an arbitrary program, guaranteed use of a structured programming language provides a great asset when developing a race detection system. To this end, we look to expand the race detector in the structured language, Habanero Java (HJ). While this language is not alone in its structured parallelism, its Author have already developed a partial race detector that detects races in the async and finish constructs as well as in isolated sections. This research builds upon this race detector by extending it to include the phaser construct as well.

Scientific Visualization Showcase

Monday, November 18

6pm-9pm
Room: Mile High Pre-Function

Tuesday, November 19

8:30am-5pm
Room: Mile High Pre-Function

Reception & Showcase
5:15pm-7pm
Room: Mile High Pre-Function

Wednesday-Thursday, November 20-21

8:30am-5pm (closes at Noon on Thursday)
Room: Mile High Pre-Function

Visualizing Simulated Volcanic Eruptions
Amit Chourasia (San Diego Supercomputer Center, University of California, San Diego), Darcy Ogden (Scripps Institution of Oceanography, University of California, San Diego)

Eruptive conduits feeding volcanic jets and plumes are connected to the atmosphere through volcanic vents that, depending on their size and 3D shape, can alter the dynamics and structure of these eruptions. The host rock comprising the vent, in turn, can collapse, fracture, and erode in response to the eruptive flow field. This project uses cutting edge visualization to illustrate and analyze results from fully coupled numerical simulations of high speed, multiphase volcanic mixtures erupting through erodible, visco-plastic host rocks.

The visualizations explore the influence of different host rock rheologies and eruptive conditions on the development of simulated volcanic jets.

Understanding the Physics of the Deflagration-to-Detonation Transition
Alexei Poludnenko, Elaine Oran (Naval Research Laboratory), Christopher Lewis (Lockheed Martin / HPCMP Data Analysis and Assessment Center), Miguel Valenciano (Data Management Consultants / HPCMP Data Analysis and Assessment Center)

An astronomical phenomenon known as the Type Ia supernovae (SN Ia) results from the thermonuclear incineration of compact, degenerate, white dwarf stars. When a close-by stellar companion is present, white dwarfs can end their life in one of the most powerful explosions in the Universe. The resulting energy release and the associated fluid expansion produce turbulent motions, which wrinkle and fold the flame, accelerating burning significantly. Turbulent flame acceleration alone, however, is often not sufficient to explain the power of these explosions. The missing piece could be the deflagration-to-detonation transition (or DDT), in which a subsonic flame develops a supersonic shock-driven reaction wave. Elucidating the physics of DDT, as well as the conditions that can lead to it, would be important for a broad range of problems from the safety of fuel storage and chemical processing facilities to the nature of the SN Ia phenomenon and of the enigmatic dark energy. The study of detonation wave concepts for propulsion applications is promising. An increase in fuel efficiency of up to 25% is possible. Since the conditions required for the onset of DDT were not known, a survey of a large parameter space was required. The largest calculations had the computational grid size of up to 1 billion cells. The overall number of time-steps per calculation reached ~1014 cell-steps. The total CPU cost of each calculation ranged from 100,000 to 500,000 CPU hours.

Formation of the Cosmic Web
Ralf Kaehler (SLAC National Accelerator Laboratory), Oliver Hahn (ETH Zurich), Tom Abel (Stanford University)

This animation is based on a cosmological N-body dark matter simulation of the formation of the large-scale structure in the Universe. It shows how gravity amplifies small dark matter density fluctuations generated shortly after the Big Bang, leading to the formation of galaxies in the cosmic web. Features of the web, like sheets, filaments and halos, the latter hosting galaxies made out of normal, baryonic matter, become clearly visible.

The visualization was generated using a highly accurate GPU-based rendering method that employs full phase-space information to generate a tetrahedral tessellation of the computational domain, with mesh vertices defined by the simulation's dark matter tracer particle positions. Over time the mesh is deformed by gravitational forces as computed by the N-body code, which causes the tetrahedral cells to warp and overlap. About 4 billion tetrahedral elements contributed to each time step in this visualization.

In-Silico Modeling for Fracture Fixation in Osteoporotic Bone

Juri Steiner (ETH Zurich), G. Harry van Lenthe (KU Leuven), Stephen J. Ferguson (ETH Zurich), Jean M. Favre (Swiss National Supercomputing Center)

Osteoporosis has become more prevalent in our aging population, making fracture treatment more difficult because of an impaired peri-implant bone microstructural quality.

Biomechanical tests are usually conducted to investigate the mechanical bone-screw competence. However, in vitro mechanical testing is expensive and time consuming because human bone material is sparse and shows a high variability.

Alternatively, computational specimen-specific models make it possible to run different mechanical tests on the same specimen. A micro-CT based Finite Element (µFE) Analysis is used to investigate the mechanical competence of bone-screw models in silico. A cylindrical trabecular bone specimen is instrumented with a titanium screw and scanned using micro-computed tomography at a nominal isotropic resolution of 20 µm. µFE models were created for two cases: a two-component bone-screw model, and a three-component model including a soft interface layer (10 voxels thick). Young's moduli for bone tissue, soft layer and titanium screw were 10 GPa, 2 GPa and 100 GPa, respectively. The models were solved using the parallel solver ParaSol, running with 384 parallel tasks for a model of 142 million hexahedral cells. Visualization was performed with a ParaView server with 64 tasks.

The resulting force was 31% lower in model 2 compared to model 1. In model 2, high displacements occur only in the vicinity of the implant, better mimicking the deformation pattern of in vitro bone-screw constructs.

This work was funded by the Swiss Commission for Technology and Innovation (CTI 14067.1 PFLS-LS) and by Synthes, Solothurn, Switzerland.

Interactive Volume Rendering of Block Copolymer's Defect Formation and Evolution

Shuxia Zhang, David Porter, Jeffrey Mcdonald, Jorge Viñals (University of Minnesota)

The study of defect formation and evolution in large aspect ratio samples of a block copolymer has great importance for practical designs in emerging technologies because they usually require coherent and uniform patterns over sufficiently large length scales. Many numerical computations of transient pattern dynamics to date have documented the existence and motion of topological defects. However, they are either two dimensional (so that defect lines are simply points) or consider a relatively small system in which it is difficult to make quantitative statements about defect statistics. We have developed a new method that combines traditional CPUs for computation, and GPUs for real time visualization to examine defect formation and motion as given by the Swift Hohenberg equation governing the evolution of a scalar order parameter field ψ. Volume-rendering technique is used to visualize the location and motion of singularities of ψ. Two different approaches are considered: one based on the local rotation of the order parameter called MFP and another, GP2P2, based on a combination of ψ and its gradient so as to extract the local, slowly varying amplitude of the pattern, but eliminate its order one periodicity. We find that both methods are capable of identifying the location of defects, and yield essentially the same large-scale structures. The animation shows the defect structures measured by MFP (the left) and GP2P2 (the right) as well as their evolution in 3D domain.

Early Evolution of a Star Cluster in the Galactic Tidal Field
David Reagan (Indiana University), Enrico Vesperini (Indiana University), Anna Lisa Varri (Indiana University), Patrick Beard (Indiana University), Chris Eller (Indiana University)
This animation shows the dynamical evolution of a star cluster as it evolves toward an equilibrium configuration. The cluster is initially 'cold' (the initial stellar velocities are not large enough to support the stellar system against its own self-gravity) and collapses. As it collapses, the cluster is also moving on its orbit around the Galactic center. The presence of the external tidal field along with the Coriolis effect due to its orbital motion in the Galaxy significantly affects the collapse process and the resulting structural and kinematic properties. In particular, the animation presented shows a deeper collapse along the vertical direction (z-axis; perpendicular to the orbital plane) leading to a final system characterized by a significant flattening. As a result of the Coriolis effect, stars acquire a significant rotation around the z-axis. During the phases of collapse and rebound shown in the animation, different parts of the system acquire rotation in opposite directions; the final system is characterized by a dense central core with counter-clockwise rotation and a low-density outer halo with clockwise rotation.

Nuclear Pasta

David Reagan, Andre S. Schneider, Charles J. Horowitz, Joseph Hughto, Don K. Berry, Eric A. Wernert, Chris Eller (Indiana University)

A supernova is a dramatic event that in a fraction of a second transforms the 10^55 separate nuclei that form the core of a massive star into a single large nucleus, a neutron star. Between the crust and the core of a neutron star, matter reaches such large densities, 10^13 to 10^14 g/cm³, that what were initially spherical nuclei merge and rearrange themselves into exotic shapes such as sheets, cylinders and others. Because of the resemblance of some of these shapes to spaghetti and lasagna these phases of matter are collectively known as nuclear pasta. We use molecular dynamics (MD) simulations to study the transitions between the different pasta shapes as

the density of matter decreases from uniform to much lower densities where nuclei become spherical again.

Organic Photovoltaic Simulation

Michael Matheson, Michael Brown, Jan-Michael Carrillo (Oak Ridge National Laboratory)

Organic Photovoltaics (OPV) offer the promise of renewable energy in the future when used in solar cells. Supercomputer simulations utilizing coarse grain molecular dynamics allow for the study of materials and the important morphology which determines the efficiency of such devices. Overcoming the low efficiency compared with inorganic photovoltaic cells is an important area of research.

Visualization and Analysis of Coherent Structures, Intermittent Turbulence, and Dissipation in High-Temperature Plasmas

Burlen Loring (Lawrence Berkeley National Laboratory), Homa Karimabadi, Vadim Rortershteyn (University of California San Diego), Minping Wan (University of Delaware), William Matthaeus, William Daughton (Los Alamos National Laboratory), Pin Wu, Michael Shay (University of Delaware), Joe Borovsky (Space Science Institute), Ersilia Leonardis, Sandra Chapman (University of Warwick), Takuma Nakamura (Los Alamos National Laboratory)

An unsolved problem in plasma turbulence is how energy is dissipated at small scales. Particle collisions are too infrequent in hot plasmas to provide the necessary dissipation. Simulations either treat the fluid scales and impose an ad-hoc form of dissipation (e.g., resistivity) or consider dissipation arising from resonant damping of small amplitude disturbances where damping rates are found to be comparable to that predicted from linear theory. Here, we report kinetic simulations that span the macroscopic fluid scales down to the motion of electrons. We find that turbulent cascade leads to generation of coherent structures in the form of current sheets that steepen to electron scales, triggering strong localized heating of the plasma. The dominant heating mechanism is due to parallel electric fields associated with the current sheets, leading to anisotropic electron and ion distributions. The motion of coherent structures also generates waves that are emitted into the ambient plasma in form of highly oblique compressional and shear Alfven modes. In 3D, modes propagating at other angles can also be generated. This indicates that intermittent plasma turbulence will in general consist of both coherent structures and waves. However, the current sheet heating is found to be locally several orders of magnitude more efficient than wave damping and is sufficient to explain the observed heating rates in the solar wind. In this work the visualization and analysis; new visualization software; and the use of supercomputing visualization resources lead to breakthroughs in the understanding of fundamental physical processes responsible for turbulent heating of the solar wind.

Ultra-High Resolution Simulation of a Downburst-Producing Thunderstorm

Robert Sisneros (University of Illinois at Urbana-Champaign), Leigh Orf (Central Michigan University), George Bryan (National Center for Atmospheric Research)

In this work we investigate simulation data from the CM1 cloud model. Specifically, typical thunderstorms are simulated under environmental conditions known to produce downbursts, such as those in the High Plains. A downburst is created by a column of sinking air that, upon impinging the ground, can create an outburst of damaging winds. Understanding the specific capabilities of downbursts is critically important to the maintenance of structures susceptible to these intense downdrafts, such as power transmission lines. In this video we first provide a cloud- and precipitation- based overview of a downburst-producing thunderstorm. We then provide the context to transition from the more recognizable and general features of the thunderstorm to those pertaining to downbursts. In high-detail we illustrate the significant small-scale characteristics of the downburst.

Clustering of Inertial Cloud Droplets in Isotropic Turbulence

John Clyne (National Center for Atmospheric Research), Peter J. Ireland, Lance R. Collins (Cornell University)

Inertial particles (i.e., particles with densities greater than that of the carrier fluid) are centrifuged out of vortex cores in a turbulent flow and accumulate in low vorticity regions. Such accumulation, referred to as “clustering” or “preferential concentration,” has been linked to the acceleration of precipitation in cumulus clouds by increasing the collision frequency and the coalescence rate of the cloud droplets. The videos show inertial droplets clustering in a direct numerical simulation of a turbulent flow performed on over 8 billion grid points, using 16,384 processors on the Yellowstone supercomputer at NCAR. The cloud droplets are visualized in white, and high vorticity regions of the flow are shown in yellow. The left image shows intermediate-sized droplets (75 microns with a Kolmogorov-scale Stokes number of 1) and the right image shows rain-size drops (400 microns with a Stokes number of 30). The former are very responsive to the underlying flow and form small, dense clusters, whereas the latter have a damped response to the flow and form larger, much less dense clusters. The radial distribution function (RDF) is a statistical measure of clustering. As you can see from the RDF plots included in the video, the degree of clustering for the smaller droplets is much higher than for the larger droplets. Indeed, clustering peaks near a Stokes number of unity.

Simulated Wave Propagation for the Mw5.4 Chino Hills, CA, Earthquake

Kim B. Olsen, William Savran (San Diego State University), Yifeng Cui, Efecan Poyraz (San Diego Supercomputer Center), Philip Maechling, Thomas H. Jordan (University of Southern California), Tim Scheitlin, Perry Domingo (National Center for Atmospheric Research)

As ever-increasing computational resources allow earthquake scientists to push the frequency limits of deterministic ground motion estimates higher, understanding small-scale, near-surface heterogeneities becomes paramount. SCEC researchers are working to improve ground motion simulations for California by developing more realistic small-scale models of the earth's near-surface structure.

This visualization compares two simulation results showing velocity magnitude at the earth's surface for the 29 July 2008, M5.4 Chino Hills, California earthquake. The earthquake was felt by many people throughout the Los Angeles basin and the surrounding areas, although there was very little damage.

For the two simulations shown, all differences can be attributed to the impact of the small-scale heterogeneities as well as anelastic attenuation. The animation on the right shows a Chino Hills simulation with unmodified SCEC Community Velocity Model (CVM-S v11.2). The animation on the left shows a Chino Hills simulation that uses a modified version of CVM-S v11.2 that contains more realistic small-scale complexities. The animations show that the more complex velocity structure used in the left simulation, clearly impacts the ground motion distribution, the levels of peak ground motion, and the duration of shaking.

The next scientific step is to compare both simulation results against observed data for this event to determine which velocity model most closely reproduces the observed ground motions for this earthquake.

Meso-to Planetary Scale Processes in a Global Ultra- High Resolution Climate Model

Justin Small, Julio Bacmeister, David Bailey, Frank Bryan, Julie Caron, David Lawrence, Bob Tomas, Joe Tribbia, Allison Baker, John Dennis, Jim Edwards, Andy Mai, Mariana Vertenstein, Tim Scheitlin, Perry Domingo (National Center for Atmospheric Research)

Community Earth System Model (CESM) output depicting hourly time steps of total column integrated water vapor (TMQ) across the globe for an entire year. The TMQ variable is represented by a 3D surface where the surface shape/height, color (black to white), and opacity (transparent to opaque) all redundantly represent the range of the variable's value from low to high. Time evolving patterns of circulation, seasonal changes, and proxies for hurricanes, cyclones, and typhoons are made apparent.

The CESM is a fully coupled, global climate model that provides state-of-the-art computer simulations of the Earth's past, present, and future climate states. This ultra-high resolution global model shows developing weather patterns in rich detail, and helps researchers understand the effect of smaller, more localized weather on the dynamics of climate at the planetary scale.

The visualization is produced from model data generated from the Advanced Scientific Discovery awards on Yellowstone, the petascale computing resource in the NCAR-Wyoming Supercomputing Center (NWSC), which opened in October 2012 in Cheyenne, Wyoming.

Visualization Of Deterministic High-Frequency Ground Motions From Simulations Of Dynamic Rupture Along Rough Faults With And Without Medium Heterogeneity Using Petascale Heterogeneous Supercomputers

Amit Chourasia, Yifeng Cui (San Diego Supercomputer Center), Efecan Poyraj (University of California, San Diego), Kim B. Olsen (San Diego State University), Jun Zhou (University of California, San Diego), Kyle Withers (San Diego State University), Scott Callaghan (University of Southern California), Jeff Larkin (NVIDIA), Clark C. Guest (University of California, San Diego), Dong J. Choi (San Diego Supercomputer Center), Philip J. Maechling (University of Southern California), Thomas H. Jordan (University of Southern California)

The accuracy of earthquake source descriptions is a major limitation in high-frequency (~>1 Hz) deterministic ground motion prediction, which is critical for performance-based design by building engineers. Here, we attempt to quantify the contributions to high-frequency (up to ~10 Hz) ground motion from small-scale complexities in both fault geometry and media. We perform wave propagation simulations using a complex kinematic source from a simulation of dynamic rupture along a rough fault. Calculation of ground motions is extended to a distance of 35 km from the fault with a highly scalable fourth-order staggered-grid finite difference method (GPU-based AWP). These animations show the strike-parallel component of the surface velocity wavefield for crustal models with and without a statistical model of small-scale heterogeneities for a Mw 7.2 earthquake. These simulations demonstrate the importance of including small-scale heterogeneities in ground motion models, in that they can cause the rupture front to break up and cause both defocusing and focusing of the wavefield. The result is a modified distribution of predicted peak ground motion, an important parameter used by engineers to establish building codes. There is also a significant amount of backscatter behind the initial waves, lengthening the duration of ground shaking. Both the rough fault and medium heterogeneities will become increasingly important in the field of earthquake hazard prediction as more powerful computers enable ground motion prediction to higher frequencies.

Secrets of the Dark Universe: Simulating the Sky on the Blue Gene/Q, The Outer Rim Simulation

Joseph A. Insley, Salman Habib, Katrin Heitmann (Argonne National Laboratory)

An astonishing 99.6% of our Universe is dark. Observations indicate that the Universe consists of 70% of a mysterious dark energy and 25% of a yet unidentified dark matter component, and only 0.4% of the remaining ordinary matter is visible. Understanding the physics of this dark sector is the foremost challenge in cosmology today. Sophisticated simulations of the evolution of the Universe play a crucial task in this endeavor. This movie shows a large simulation of the distribution of matter in the Universe, the so-called cosmic web, which evolved under the influence of dark energy. The simulation is evolving 1.1 trillion particles using HACC, a new computational framework developed to overcome the challenges posed by future supercomputing architectures. It is currently running on 32 racks of Mira, the Blue Gene/Q system at the Argonne Leadership Computing Facility. The visualization was performed on Argonne's Tukey cluster.

Visualization of Three-Dimensional Temperature Distribution in Google Earth

Fumiaki Araki (Japanese Agency for Marine-Earth Science and Technology), Tooru Sugiyama (Japanese Agency for Marine-Earth Science and Technology), Shintaro Kawahara (Japanese Agency for Marine-Earth Science and Technology), Keiko Takahashi (Japanese Agency for Marine-Earth Science and Technology)

The graphical contents described on Google Earth are mainly classified according to a number of spatial dimensions from zero to two (2D): the point source of an earthquake-epicenter, the curves of cruising-track, and the surface of satellite image are examples on the geo-scientific topics. This classification below 2D comes from that Keyhole Markup Language (KML) used in Google Earth does not support three-dimensional objects except the statistical charts or the model buildings. Therefore, a new idea is expected to describe the three-dimensional contents, especially obtained from numerical simulations in which the data are defined on spatially three-dimensional grids. Here we propose a novel technique to represent the three-dimensional scalar data on Google Earth. The key is opacity. Appropriate value of the opacity is assigned to data. Resultantly, we can easily see multiple data aligned the view direction. We select PNG image-format because the alpha channel for opacity is supported. Attached is the visualized movie on Google Earth. The simulation is performed on Earth Simulator using the urban-scale atmospheric simulation model of MSSG. The shown dataset is the three-dimensional, time-developing temperature distribution in the area of about 4km-square in Minato Mirai 21 district, Yokohama, Japan, The spatial resolution is about 5 meter (~ 0.000045 degree) with the grid number of 968 (longitude) x 792 (latitude) x 82 (layers). The calculated time is 30 minutes with interval of 10 seconds.

Tutorials

Tutorials

The Tutorials program is always one of the highlights of the SC Conference, offering attendees a variety of short courses on key topics and technologies relevant to high performance computing, networking, storage, analysis. Tutorials also provide the opportunity to interact with recognized leaders in the field and to learn about the latest technology trends, theory, and practical techniques. As in years past, tutorial submissions were subjected to a rigorous peer review process. Of the 74 submissions, the 30 members of the Tutorials Committee selected 30 presented here.

Tutorials

Sunday, November 17

A "Hands-On" Introduction to OpenMP

8:30am-5pm
Room: 402

Presenters: Tim Mattson (Intel Corporation), Mark Bull (Edinburgh Parallel Computing Center), 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 3.1 standard.

This will be a hands-on tutorial. 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.

Advanced MPI Programming

8:30am-5pm
Room: 406

Presenters: Pavan Balaji (Argonne National Laboratory), James Dinan (Intel Corporation), 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 approximately 1.6 million cores) and achieving 12 to 14 petaflops/s of sustained performance. At the same time, the MPI standard itself is evolving (MPI-3 was released late last year) 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 data types, one-sided communication, hybrid (MPI + shared memory) programming, topologies and topology mapping, and neighborhood and non-blocking 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.

Debugging MPI and Hybrid/Heterogeneous Applications at Scale

8:30am-5pm
Room: 404

Presenters: Ganesh Gopalakrishnan (University of Utah), David Lecomber (Allinea Software), Matthias S. Mueller (RWTH Aachen University), Bronis R. de Supinski (Lawrence Livermore National Laboratory), Tobias Hilbrich (Dresden University of Technology)

MPI programming is error prone due to the complexity of MPI semantics and the difficulties of parallel programming. Increasing heterogeneity (e.g., MPI plus OpenMP/CUDA), scale, nondeterminism, and platform dependent bugs exacerbate these difficulties. This tutorial covers the detection/correction of errors in MPI programs at small and large scale, as well as for heterogeneous/hybrid programs. We will first introduce our main tools: MUST, that detects MPI usage errors at runtime with a high degree of automation; ISP/DAMPI, that detects interleaving-dependent MPI deadlocks and assertion violations through application replay; and DDT, a highly scalable parallel debugger. Attendees will be encouraged to explore our tools early during the tutorial to appreciate their strengths/limitations better. We will present best practices and a cohesive workflow for comprehensive application debugging with all our tools.

We dedicate the afternoon session to advanced use-cases, tool deployment on Leadership-Scale systems, updates on new tool functionality, and for the debugging of hybrid/heterogeneous programming models. The latter includes debugging approaches for MPI, OpenMP, and CUDA and is especially crucial for systems such as Titan (ORNL) and Sequoia (LLNL). DDT's capabilities for CUDA/OpenMP debugging will be presented, in addition to a short introduction to GKLEE, a new symbolic verifier for CUDA applications.

Globus Online and the Science DMZ as Scalable Research Data Management Infrastructure for HPC Facilities

8:30am-5pm
Room: 403

Presenters: Steve Tuecke (University of Chicago), Raj Kettimuthu (Argonne National Laboratory), Eli Dart (Lawrence Berkeley National Laboratory), Vas Vasiliadis (University of Chicago)

The rapid growth of data in scientific research endeavors is placing massive demands on campus computing centers and HPC facilities. Computing facilities must provide robust data services built on high-performance infrastructure, while continuing to scale as needs increase. Traditional research data management (RDM) solutions are typically difficult to use and error-prone, and the underlying networking and security

infrastructure is often complex and inflexible, resulting in user frustration and sub-optimal use of resources.

An increasingly common solution in HPC facilities is Globus Online deployed in a network environment built on the Science DMZ model. Globus Online is software-as-a-service for moving, syncing, and sharing large data sets. The Science DMZ model is a set of design patterns for network equipment, configuration, and security policy for high-performance scientific infrastructure. The combination of user-friendly, high-performance data transfer tools, and optimally configured underlying infrastructure results in enhanced RDM services that increase user productivity and lower support overhead.

Guided by two case studies from national supercomputing centers (NERSC and NCSA), attendees will explore the challenges such facilities face in delivering scalable RDM solutions. Attendees will be introduced to Globus Online and the Science DMZ, and will learn how to deploy and manage these systems.

Hands-On Practical Hybrid Parallel Application Performance Engineering

8:30am-5pm
Room: 301

Presenters: Markus Geimer (Juelich Supercomputing Center), Sameer Shende (University of Oregon), Bert Wesarg (Dresden University of Technology), Brian Wylie (Juelich Supercomputing Center)

This tutorial presents state-of-the-art performance tools for leading-edge HPC systems founded on the Score-P community instrumentation and measurement infrastructure, demonstrating how they can be used for performance engineering of effective scientific applications based on standard MPI, OpenMP, hybrid MPI+OpenMP, and increasingly common usage of accelerators. Parallel performance evaluation tools from the VI-HPS (Virtual Institute – High Productivity Supercomputing) are introduced and featured in hands-on exercises 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. Using their own notebook computers with a provided Linux Live-ISO image containing the tools (booted from DVD/USB or within a virtual machine) will help to prepare participants to locate and diagnose performance bottlenecks in their own parallel programs.

Hybrid MPI and OpenMP Parallel Programming

8:30am-12pm
Room: 407

Presenters: Rolf Rabenseifner (High Performance Computing Center Stuttgart, HLRS), Georg Hager (Erlangen Regional Computing Center), Gabriele Jost (Supersmith)

Most HPC systems are clusters of shared memory nodes. Such systems can be PC clusters with single/multi-socket and multi-core SMP nodes, but also “constellation” type systems with large SMP nodes. Parallel programming may combine the distributed memory parallelization on the node interconnect with the shared memory parallelization inside of each node. 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 MPI message passing and remote memory access on the cluster interconnect. 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 on different platforms are presented. Numerous case studies demonstrate the performance-related aspects of hybrid programming, and application categories that can take advantage of this model are identified. Tools for hybrid programming such as thread/process placement support and performance analysis are presented in a “how-to” section. Details: https://fs.hlrs.de/projects/rabenseifner/publ/SC2013-hybrid.html.

InfiniBand and High-speed Ethernet for Dummies

8:30am-12pm
Room: 201/203

Presenters: Dhabaleswar K. Panda, Hari Subramoni (Ohio State University)

InfiniBand (IB) and High-Speed Ethernet (HSE) technologies are generating a lot of excitement towards building next generation High-End Computing (HEC) systems including clusters, datacenters, file systems, storage, cloud computing and Big Data (Hadoop, 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. Next, an overview of the emerging OpenFabrics stack which encapsulates IB, HSE and RoCE in a unified manner will be presented.

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.

Large Scale Visualization with ParaView

8:30am-12pm
Room: 405

Presenters: Kenneth Moreland, W. Alan Scott (Sandia National Laboratories), David DeMarle (Kitware, Inc.), Li-Ta Lo (Los Alamos 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.

OpenACC: Productive, Portable Performance on Hybrid Systems Using High-Level Compilers and Tools

8:30am-5pm
Room: 401

Presenters: Luiz DeRose, Alistair Hart, Heidi Poxon, James Beyer (Cray Inc.)

Portability and programming difficulty are two critical hurdles in generating widespread adoption of accelerated computing in high performance computing. The dominant programming models for accelerator-based systems (CUDA and OpenCL) offer the power to extract performance from accelerators, but with extreme costs in usability, maintenance, development, and portability. To be an effective HPC platform, hybrid systems need a high-level programming environment to enable widespread porting and development of applications that run efficiently on either accelerators or CPUs. In this hands-on tutorial we present the high-level OpenACC parallel programming model for accelerator-based systems, demonstrating compilers, libraries, and tools that support this cross-vendor initiative. Using personal experience in porting large-scale HPC applications, we provide development guidance, practical tricks, and tips to enable effective and efficient use of these hybrid systems.

Parallel Computing 101

8:30am-5pm
Room: 303

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, 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 programming, and Hadoop on big data. It discusses numerous parallelization 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 primary 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.

Programming for the Intel Xeon Phi

8:30am-5pm
Room: 205/207

Presenters: Lars Koesterke, Luke Wilson, John McCalpin, Kent Milfeld (Texas Advanced Computing Center)

The Innovative Technology component of the recently deployed XSEDE Stampede supercomputer at TACC provides access to 8 Petaflops of computing power in the form of the new Intel Xeon Phi Coprocessor, also known as a MIC. While the MIC is x86 based, hosts its own Linux OS, and is capable of running most user codes with little porting effort, the MIC architecture has significant features that are different from that of present x86 CPUs, and optimal performance requires an understanding of the possible execution models and basic details of the architecture. This tutorial is designed to introduce attendees to the MIC architecture in a practical manner. Multiple lectures and hands-on exercises will be used to acquaint attendees with the MIC platform and explore the different execution modes as well as parallelization and optimization through example testing and reports.

Structured Parallel Programming with Patterns

8:30am-5pm
Room: 302

Presenters: Michael D. McCool (Intel Corporation), James R. Reinders (Intel Corporation), Arch Robison (Intel Corporation), Michael Hebenstreit (Intel Corporation)

Parallel programming is important for performance, and developers need a comprehensive set of strategies and technologies for tackling it. This tutorial is intended for C++ programmers who want to better grasp how to envision, describe and write efficient parallel algorithms at the single shared-memory node level.

This tutorial will present a set of algorithmic patterns for parallel programming. Patterns describe best known methods for solving recurring design problems. Algorithmic patterns in particular are the building blocks of algorithms. Using these patterns to develop parallel algorithms will lead to better structured, more scalable, and more maintainable programs. This course will discuss when and where to use a core set of parallel patterns, how to best implement them, and how to analyze the performance of algorithms built using them. Patterns to be presented include map, reduce, scan, pipeline, fork-joint, stencil, tiling, and recurrence. Each pattern will be demonstrated using working code in one or more of Cilk Plus, Threading Building Blocks, OpenMP, or OpenCL. Attendees also will have the opportunity to test the provided examples themselves on an HPC cluster for the time of the SC13 conference.

Advanced Topics in InfiniBand and High-Speed Ethernet for Designing High-End Computing Systems

1:30pm-5pm
Room: 201/203

Presenters: Dhabaleswar K. Panda, Hari Subramoni (Ohio State University)

As InfiniBand (IB) and High-Speed Ethernet (HSE) technologies mature, they are being used to design and deploy different kinds of High-End Computing (HEC) systems: HPC clusters with accelerators (GPGPUs and MIC) supporting MPI and PGAS (UPC and OpenSHMEM), Storage and Parallel File Systems, Cloud Computing with Virtualization, Big Data systems with Hadoop (HDFS, MapReduce and HBase), Multi-tier Datacenters with Web 2.0 (memcached) and Grid Computing systems. These systems are bringing new challenges in terms of performance, scalability, and portability. Many scientists, engineers, researchers, managers and system administrators are becoming interested in learning about these challenges, approaches being used to solve these challenges, and the associated impact on performance and scalability. This tutorial will start with an overview of these systems and a common set of challenges being faced while designing these systems. Advanced hardware and software features of IB and HSE and their capabilities to address these challenges will be emphasized. Next, case studies focusing on domain-specific challenges in designing these systems (including the associated software stacks), their solutions and sample performance numbers will be presented. The tutorial will conclude with a set of demos focusing on RDMA programming, network management infrastructure and tools to effectively use these systems.

Practical Fault Tolerance on Today's HPC Systems

1:30pm-5pm
Room: 405

Presenters: Kathryn Mohror (Lawrence Livermore National Laboratory), Nathan Debardeleben (Los Alamos National Laboratory), Eric Roman (Lawrence Berkeley National Laboratory), Laxmikant Kale (University of Illinois at Urbana-Champaign)

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 coordinated checkpoint and restart mechanisms and programming model-based approaches for MPI applications.

Scaling I/O Beyond 100,000 Cores using ADIOS

1:30pm-5pm
Room: 407

Presenters: Qing Liu (Oak Ridge National Laboratory), Norbert Podhorszki (Oak Ridge National Laboratory), Scott Klasky (Oak Ridge National Laboratory)

As concurrency continues to increase on high-end machines, from both the number of cores and storage devices, we must look for a revolutionary way to treat I/O. As a matter of fact, one of the major roadblocks to exascale is how to write and read big datasets quickly and efficiently on high-end machines. On the other hand applications often want to process data in an efficient and flexible manner, in terms of data formats and operations performed (e.g., files, data streams). In this tutorial we will show how users can do that and get high performance with ADIOS on 100,000+ cores. Part I of this tutorial

will introduce parallel I/O and the ADIOS framework to the audience. Specifically we will discuss the concept of ADIOS I/O abstraction, the binary-packed file format, and I/O methods along with the benefits to applications. Since 1.4.1, ADIOS can operate on both files and data streams. Part II will include a session on how to write/read data, and how to use different I/O componentizations inside of ADIOS. Part III will show users how to take advantage of the ADIOS framework to do compression/indexing. Finally, we will discuss how to run in-situ visualization using VisIt/Paraview+ADIOS.

Monday, November 18

Advanced OpenMP: Performance and 4.0 Features

8:30am-5pm
Room: 406

Presenters: Christian Terboven (RWTH Aachen University), Michael Klemm (Intel Corporation), Ruud van der Pas (Oracle), 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 and will easily grasp those basics. 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 features recently added to OpenMP such as tasking and cancellation. We close with an overview of the new OpenMP 4.0 directives for attached compute accelerators.

Advanced PGAS Programming in UPC

8:30am-12pm
Room: 303

Presenters: Katherine Yelick (University of California, Berkeley), Yili Zheng, Costin Iancu (Lawrence Berkeley National Laboratory)

Partitioned Global Address Space (PGAS) languages combine the convenience of shared memory programming with the locality control needed for scalability. There are several PGAS languages based on a variety of serial languages (e.g., Fortran, C, and Java) and different parallel control constructs, but all with a similar model for building shared data structures. This tutorial will focus on performance programming in the UPC language, which has a strict superset of ISO C. Following a short introduction to the basic constructs in UPC, the tutorial will cover effective parallel programming idioms used in UPC applications and performance optimization techniques. This will include the design of pointer-based data structures, the proper use of locks, and the memory consistency model. The tutorial will also cover the latest extensions of the language, both those in the current specification and others being considered for the next version of the language specification.

Asynchronous Hybrid and Heterogeneous Parallel Programming with MPI/OmpSs and its Impact in Energy Efficient Architectures for Exascale Systems

8:30am-5pm
Room: 401

Presenters: Jesus Labarta, Eduard Ayguade, Alex Ramirez, Rosa M. Badia (Barcelona Supercomputing Center)

Due to its asynchronous nature and look-ahead capabilities, MPI/OmpSs is a promising programming model approach for future exascale systems, with the potential to exploit unprecedented amounts of parallelism, while coping with memory latency, network latency and load imbalance. Many large-scale applications are already seeing very positive results from their ports to MPI/OmpSs (see EU projects Montblanc, DEEP, TEXT). We will first cover the basic concepts of the programming model. OmpSs can be seen as an extension of the OpenMP model. Unlike OpenMP, however, task dependencies are determined at runtime thanks to the directionality of data arguments. The OmpSs runtime supports asynchronous execution of tasks on heterogeneous systems such as SMPs, GPUs and clusters thereof. The integration of OmpSs with MPI facilitates the migration of current MPI applications and improves the performance of these applications by overlapping computation with communication between tasks. The tutorial will also cover the performance tools available for the programming model: Paraver performance analysis tool. Examples of benchmarks and applications parallelized with MPI/OmpSs will also be presented. The tutorial will present the impact of the programming model to address the limitations of using low-end devices to build power-efficient parallel platforms. The tutorial will also include hands-on.

Debugging and Optimizing MPI and OpenMP™ Applications Running on CUDA™, OpenACC®, and Intel® Xeon Phi™ Coprocessors with TotalView®

8:30am-5pm
Room: 405

Presenters: Mike Ashworth (Science and Technology Facilities Council), Vince Betro (National Institute for Computational Sciences), Sandra Wienke (RWTH Aachen University), Nikolay Piskun (Rogue Wave Software), Chris Gottbrath (Rogue Wave Software)

With HPC trends heading towards increasingly heterogeneous solutions, scientific developers face challenges adapting software to leverage these new systems. For instance, 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™ coprocessor. In order to utilize these systems, scientific programmers need to leverage as much parallelism in applications as possible. Developers also need to juggle technologies including MPI, OpenMP, CUDA, and OpenACC. While troubleshooting, debugging, and optimizing applications are an expected part of porting, they become even more critical with the introduction of so many technologies.

This tutorial provides an introduction to parallel debugging and optimization. Debugging techniques covered include: MPI and subset debugging, process and thread sets, reverse and comparative debugging, and techniques for CUDA, OpenACC, and Intel Xeon Phi coprocessor debugging. Participants will have the opportunity to do hands-on CUDA and Intel Xeon Phi coprocessor debugging using TotalView on a cluster at RWTH Aachen University and on Keeneland and Beacon at NICS. Therefore, it is recommended that participants bring a network-capable laptop to the session. Optimization techniques will include profiling, tracing, and cache memory optimization. Examples will use ThreadSpotter and vendor-supplied tools.

Effective Procurement of Supercomputers

8:30am-12pm
Room: 302

Presenters: Terry Hewitt (WTH Associates Limited), Andrew Jones (Numerical Algorithms Group), Jonathan Follows (Science and Technology Facilities Council)

In this tutorial we will guide you through the process of purchasing a cluster, HPC system or supercomputer. We will take you through the whole process from engaging with major stakeholders in securing the funding, requirements capture, market survey, specification of the tender/request for quote documents, engaging with suppliers, and evaluating proposals. You will learn how to specify what you want, yet enable the suppliers to provide innovative solutions beyond your specification both in technology and in the price, and then how to demonstrate to stakeholders the solution you select is value for money.

The tutorial will be spilt into 3 major parts: the procurement process, the role of benchmarks and market surveys, and specification & evaluation.

The presenters have been involved in major national and international HPC procurements since 1990 both as bidders and as customers. Whether you are spending $100k or $100M you will benefit from this tutorial.

Ensuring Network Performance with perfSONAR

8:30am-5pm
Room: 402

Presenter: Jason Zurawski (Energy Sciences Network)

A key component to “Super Computing” is “Super Networking.” Scientific data sets continue to increase in number, size, and importance for numerous research and education (R&E) communities, and rely on networks to facilitate sharing. Solicitations, such as the NSF’s CC-NIE, have recognized this; new paradigms in networking, increased capacities, and emphasis on monitoring were specifically mentioned as areas of potential funding.

A user’s network experience must be reliable—free of architectural flaws, and physical limitations. Operational staffs are limited in the support they can deliver, normally within a domain; innovative tools are required to solve the “end-to-end” performance problems that hamper network use. End users should be familiar with these tools; to protect their own interests and expectations, and to assist operations in debugging exercises.

We will present an overview of network performance tools and techniques, focusing on the deployment of the pS-Performance Toolkit. This “all-in-one”, community developed, monitoring solution allows local control, while providing a global view of performance that will directly impact the use of networks. Goals include familiarizing attendees on ways these tools may aid in debugging networks, hosts, and applications, as well as the proper way to install and configure software for personal use.

How to Analyze the Performance of Parallel Codes 101

8:30am-12pm
Room: 301

Presenters: Martin Schulz (Lawrence Livermore National Laboratory), Jim Galarowicz (Krell Institute), Jennifer Green (Los Alamos National Laboratory), Matthew LeGendre (Lawrence Livermore National Laboratory), Don Maghrak (Krell Institute)

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.

In this tutorial 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 in HPC codes. We will cover both on-node performance and communication optimization and will also touch on threaded and accelerator-based architectures. Throughout this tutorial, we will show live demos using Open|SpeedShop, a comprehensive and easy to use performance analysis tool set, to demonstrate the individual analysis steps. All techniques will, however, apply broadly to any tool and we will point out alternative tools where useful.

Linear Algebra Libraries for HPC: Scientific Computing with Multicore and Accelerators

8:30am-5pm
Room: 407

Presenters: Jack Dongarra (University of Tennessee, Knoxville), James Demmel (University of California, Berkeley), Michael Heroux (Sandia National Laboratories), Jakub Kurzak (University of Tennessee, Knoxville)

Today, desktops 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 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 sub-problems. 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. PLASMA, MAGMA 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.

OpenCL: A Hands-On Introduction

8:30am-5pm
Room: 403

Presenters: Tim Mattson (Intel Corporation), Alice Koniges (Lawrence Berkeley National Laboratory), Simon McIntosh-Smith (University of Bristol)

OpenCL is an open standard for programming heterogeneous parallel computers composed of CPUs, GPUs and other processors. OpenCL consists of a framework to manipulate the host CPU and one or more compute devices (CPUs, GPUs or accelerators), and a C-based programming language for writing programs for the compute devices. Using OpenCL, a programmer can write parallel programs that harness all of the resources of a heterogeneous computer.

In this hands-on tutorial, we will introduce OpenCL. For ease of learning we will focus on the easier to use C++ API, but attendees will also gain an understanding of OpenCL's C API.

The format will be a 50/50 split between lectures and exercises. Students will use their own laptops (Windows, Linux or OS/X) and log into a remote server running an OpenCL platform on a range of different processors. Alternatively, students can load OpenCL onto their own laptops prior to the course (Intel, AMD and NVIDIA provide OpenCL SDKs. Apple laptops with X-code include OpenCL by default).

By the end of the course, attendees will be able to write and optimize OpenCL programs, and will have a collection of example codes to help with future OpenCL program development.

Parallel I/O in Practice

8:30am-5pm
Room: 205/207

Presenters: Robert Latham, Robert Ross (Argonne National Laboratory), Brent Welch (Panasas), 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 from parallel file systems at the lowest layer, to intermediate layers (such as MPI-IO), and finally high-level I/O libraries (such as HDF-5). We emphasize ways to use these interfaces that result in high performance. Benchmarks on real systems are used throughout to show real-world results.

This tutorial first discusses parallel file systems in detail (PFSs). We cover general concepts and examine four examples: GPFS, Lustre, PanFS, and PVFS. We examine the upper layers of the I/O stack, covering POSIX I/O, 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 I/O best practice.

Python in HPC

8:30am-5pm
Room: 201/203

Presenters: Andy Terrel (University of Texas at Austin), Travis Oliphant (Continuum Analytics), Aron Ahmadia (Continuum Analytics), Kurt Smith (Enthought, Inc.)

The Python ecosystem empowers the HPC community with a stack of tools that are not only powerful but a joy to work with. It is consistently one of the top languages in HPC with a growing vibrant community of open source tools. Proven to scale on the world's largest clusters, it is a language that has continued to innovate with a wealth of new data tools.

This tutorial will survey the state of the art tools and techniques used by HPC Python experts throughout the world. The first half of the day will include an introduction to the standard toolset used in HPC Python and techniques for speeding Python and using legacy codes by wrapping Fortan and C. The second half of the day will include discussion on using Python in a distributed workflow via MPI and tools for handling Big Data analytics.

Students should be familiar with basic Python syntax, we recommend the Python 2.7 tutorial on python.org. We will include hands-on demonstrations of building simulations, wrapping low-level code, executing on a cluster via MPI, and use of big data tools. Examples for a range of experience levels will be provided.

The Practitioner's Cookbook for Good Parallel Performance on Multi- and Many-Core Systems

8:30am-5pm
Room: 404

Presenters: Georg Hager, Jan Treibig (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. Also, the potential of node-level improvements is widely underestimated, thus it is vital to understand the performance-limiting factors on modern hardware. We convey the architectural features of current processor chips, multiprocessor nodes, and accelerators, as well as the performance properties of the dominant MPI and OpenMP programming models, 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 is introduced as a powerful tool that helps the user assess the impact of possible code optimizations by establishing models for the interaction of the software with the hardware.

An Overview of Fault-Tolerant Techniques for HPC

1:30pm-5pm
Room: 301

Presenters: Thomas Herault (University of Tennessee, Knoxville), Yves Robert (ENS Lyon)

Resilience is a critical issue for large-scale platforms. This tutorial provides a comprehensive survey on fault-tolerant techniques for high-performance computing. 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) Application-specific techniques, such as ABFT for grid-based algorithms or fixed-point convergence for iterative applications; (iii) General-purpose techniques, which include several checkpoint and rollback recovery protocols, possibly combined with replication; and (iv) Relevant execution scenarios will be evaluated and compared through quantitative models (from Young's approximation to Daly's formulas and recent work).

The tutorial is open to all SC13 attendees who are interested in the current status and expected promise of fault-tolerant approaches for scientific applications. There are no audience prerequisites; background will be provided for all protocols and probabilistic models. Only the last part of the tutorial devoted to assessing the future of the methods will involve more advanced analysis tools.

Effective HPC Visualization and Data Analysis using VisIt

1:30pm-5pm
Room: 303

Presenters: Cyrus Harrison (Lawrence Livermore National Laboratory), Jean M. Favre (Swiss National Supercomputing Center), Hank Childs (University of Oregon and Lawrence Berkeley National Laboratory), Harinarayan Krishnan (Lawrence Berkeley National Laboratory), Brad Whitlock (Lawrence Livermore National Laboratory)

Visualization is an essential component of the scientific discovery process. Scientists and businesses running HPC simulations leverage visualization tools for data exploration, quantitative analysis, visual debugging, and communication of results. This half-day tutorial will provide attendees with a practical introduction to mesh-based HPC visualization 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 creating visualizations.

This tutorial includes: (1) An introduction to visualization techniques for mesh-based simulations; (2) A guided tour of VisIt; (3) Hands on demonstrations of end-to-end visualizations of both a fluid simulation and climate simulation.

This tutorial builds on the past success of VisIt tutorials, updated and anchored with new concrete use cases. Attendees will gain practical knowledge and recipes to help them effectively use VisIt to analyze data from their own simulations.

Introducing R: from Your Laptop to HPC and Big Data

1:30pm-5pm
Room: 302

Presenters: Drew Schmidt (University of Tennessee, Knoxville), George Ostrouchov (Oak Ridge National Laboratory)

The R language has been called the lingua franca of data analysis and statistical computing and is quickly becoming the de facto standard for analytics. As such, R is the tool of choice for many working in the fields of machine learning, statistics, and data mining. This tutorial will introduce attendees to the basics of the R language with a focus on its recent high performance extensions enabled by the ``Programming with Big Data in R'' (pbdR) project. Although R has a reputation for lacking scalability, our initial experiments with pbdR have easily scaled to 12 thousand cores. No background in R is assumed but even R veterans will benefit greatly from the session. We will cover only those basics of R that are needed

for the HPC portion of the tutorial. The tutorial is very much example-oriented, with many opportunities for the engaged attendee to follow along. Examples will utilize common data analytics techniques, such as principal components analysis and cluster analysis.

Workshops

SC13 includes nearly 30 full-day and half-day workshops that complement the overall Technical Program events, expand the knowledge base of its subject area, and extend its impact by providing greater depth of focus. These workshops are geared toward providing interaction and in-depth discussion of stimulating topics of interest to the HPC community.

New programs have been added that highlight innovative technologies in HPC, alongside the traditional program elements that our community relies upon to stay abreast of the complex, changing landscape of HPC. Together these are the elements that make SC the most informative, exciting, and stimulating technical program in HPC!

Workshops

Workshops

Sunday, November 17

1st Workshop on Sustainable Software for Science: Practice and Experiences (WSSSPE)

9am-5:30pm
Room: 503

Daniel S. Katz (National Science Foundation), Gabrielle Allen (Skolkovo Institute of Science and Technology), Neil Chue Hong (Software Sustainability Institute, University of Edinburgh), Manish Parashar (Rutgers University), David Proctor (National Science Foundation)

Progress in scientific research is dependent on the quality and accessibility of software at all levels and it is now critical to address many new challenges related to the development, deployment, and maintenance of reusable software. In addition, it is essential that scientists, researchers, and students are able to learn and adopt a new set of software-related skills and methodologies. Established researchers are already acquiring some of these skills, and in particular a specialized class of software developers is emerging in academic environments who are an integral and embedded part of successful research teams. This workshop will provide a forum for discussion of the challenges, including both positions and experiences. The short papers and discussion will be archived as a basis for continued discussion, and we intend the workshop to feed into the collaborative writing of one or more journal publications. For more information go to: *wssspe.researchcomputing.org.uk.*

4th SC Workshop on Petascale (Big) Data Analytics: Challenges and Opportunities

9am-5:30pm
Room: 501

Ranga Raju Vatsavai, 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 they are archiving every year. On the other hand, computing infrastructures are becoming more heterogeneous. The first three workshops held with SC10, SC11, and SC12 were 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 facilitates 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.

6th Workshop on Many-Task Computing on Clouds, Grids, and Supercomputers (MTAGS) 2013

9:30am-5:30pm
Room: 502

Ioan Raicu (Illinois Institute of Technology), Ian Foster (Argonne National Laboratory), Yong Zhao (University of Electronic Science and Technology of China), Justin Wozniak (Argonne National Laboratory)

This workshop 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 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. For more information go to: *datasys.cs.iit.edu/events/MTAGS13/*.

'Building' Energy Efficient High Performance Computing: 4th Annual EE HPC WG Workshop

9am-5:30pm
Room: 603

Anna Maria Bailey (Lawrence Livermore National Laboratory), James Rogers (Oak Ridge National Laboratory), James Laros (Sandia National Laboratories), Susan Coghlan (Argonne National Laboratory), Josip Loncaric (Los Alamos National Laboratory), William Tschudi (Lawrence Berkeley National Laboratory), Ralph Wescott (Pacific Northwest National Laboratory), Natalie Bates (Lawrence Berkeley National Laboratory), Steven Hammond (National Renewable Energy 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 300 members from approximately 20 different countries. Speakers from prior years include John Shalf, Lawrence Berkeley National Laboratory, Satoshi Matsuoka, Tokyo Institute of Technology,

Herbert Huber, Leibniz Supercomputing Center, Steve Hammond, National Renewable Energy Laboratory, Michael Patterson, Intel and Nic Dube, Hewlett Packard. They are well known leaders in energy efficiency for supercomputing and

delivered a lively and informative session. Similar speakers can be expected for SC13. For more information see: *eehpcwg.lbl.gov/conferences.*

IA^3 2013 - 3rd Workshop on Irregular Applications: Architectures & Algorithms

9am-5:30pm
Room: 504

John Feo, Antonino Tumeo, Oreste Villa (Pacific Northwest National Laboratory), Simone Secchi (Università di Cagliari)

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 them demands a substantial effort and often leads to poor performance. However, solving these applications efficiently will be a key requirement for future 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. For more information go to: *cassmt.pnnl.gov/irregularworkshop.aspx.*

The 3rd International Workshop on Network-Aware Data Management

9am-5:30pm
Room: 601

Mehmet Balman (Lawrence Berkeley National Laboratory), Brian L. Tierney (Energy Sciences Network), Surendra Byna (Lawrence Berkeley National Laboratory)

In addition to increasing data volumes, future scientific collaborations require cooperative work at the extreme scale. As the number of multidisciplinary teams and experimental facilities increase, data sharing and resource coordination among distributed centers are becoming significant challenges every passing year. In the age of extraordinary advances in communication technologies, there is a need for efficient use of the network infrastructure to address increasing data requirements of today's applications. Traditional network and data management techniques are unlikely to scale in future collaborative data-intensive systems. We require novel data access mechanisms and intelligent network middleware to envision future design principles of network-aware data management. This workshop will seek contributions from academia, government, and industry to discuss emerging trends and current technological developments in dynamic resource provisioning, intelligent data-flow and resource coordination, end-to-end processing of data, network-aware application design issues, and cutting-edge network performance problems. For more information go to: *sdm.lbl.gov/ndm.*

The 4th International Workshop on Data-Intensive Computing in the Clouds

9am-5:30pm
Room: 507

Yong Zhao (University of Electronic Science and Technology of China), Ziming Zheng (University of Chicago), Hui Jin (Oracle)

Applications and experiments in all areas of science are becoming increasingly complex and more demanding in terms of their computational and data requirements. As scientific applications become more data intensive, the management of data resources and dataflow between storages 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.

The 4th international workshop on Data-intensive Computing in the Clouds (DataCloud 2013) will provide the scientific community a dedicated forum for discussing new research, development, and deployment efforts in running data-intensive computing workloads on Cloud Computing infrastructures. The DataCloud 2013 workshop will focus on the use of cloud-based technologies to meet the new data intensive scientific challenges that are not well served by the current supercomputers, grids or compute-intensive clouds. For more information go to: *datasys.cs.iit.edu/events/DataCloud2013/.*

The 8th Workshop on Ultrascale Visualization

9am-5:30pm
Room: 505

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 and experiments is so voluminous and complex that advanced visualization techniques are necessary to interpret the calculated results. 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. This workshop 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 the scientific supercomputing and discovery process. For more information go to: *vis.cs.ucdavis.edu/Ultravis13.*

The 8th Workshop on Workflows in Support of Large-Scale Science (WORKS13)

9am-5:30pm
Room: 506

Johan Montagnat (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. For more information go to: *works.cs.cardiff.ac.uk/index.ph.*

Monday, November 18

'Building' Energy Efficient High Performance Computing: 4th Annual EE HPC WG Workshop

9am-12pm
Room: 603

Anna Maria Bailey (Lawrence Livermore National Laboratory), James Rogers (Oak Ridge National Laboratory), Susan Coghlan (Argonne National Laboratory), Josip Loncaric (Los Alamos National Laboratory), William Tschudi (Lawrence Berkeley National Laboratory), Ralph Wescott (Pacific Northwest National Laboratory), Natalie Bates (Lawrence Berkeley National Laboratory), Steven Hammond (National Renewable Energy Laboratory), James Laros (Sandia 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 300 members from approximately 20 different countries. Speakers from prior years include John Shalf, Lawrence Berkeley National Laboratory, Satoshi Matsuoka, Tokyo Institute of Technology, Herbert Huber, Leibniz Supercomputing Center, Steve Hammond, National Renewable Energy Laboratory, Michael Patterson, Intel and Nic Dube, Hewlett Packard. They are well known leaders in energy efficiency for supercomputing and delivered a lively and informative session. Similar speakers can be expected for SC13. For more information go to: *eehpcwg.lbl.gov/conferences.*

4th International Workshop on Performance Modeling, Benchmarking and Simulation of HPC Systems (PMBS13)

9am-5:30pm
Room: 502

Stephen Jarvis (University of Warwick), Simon Hammond (Sandia National Laboratories), Todd Gamblin (Lawrence Livermore National Laboratory), Darren Kerbyson (Pacific Northwest National Laboratory), Rolf Riesen (IBM Research), Yunquan Zhang (Chinese Academy of Sciences)

This workshop is concerned with the comparison of HPC systems through performance modeling, benchmarking or through the use of tools such as simulators. We are particularly interested in the ability to measure and make tradeoffs in software/hardware co-design to improve sustained application performance. We are also concerned with the assessment of future systems to ensure continued application scalability through peta- and exascale 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 HPC systems. The coverage of the term 'performance' has broadened to include power consumption and reliability, and that performance modeling is practiced through analytical methods and approaches based on software tools and simulators. For more information go to: *www.pmbsworkshop.org.*

4th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (ScalA)

9am-5:30pm
Room: 507

Vassil Alexandrov (Barcelona Supercomputing Center-ICREA), 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. This is especially true for the current tier of leading petascale machines and the road to exascale computing 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, and have no synchronization points. Scientific algorithms for multi-petaflop and exa-flop systems also need to be fault tolerant and fault resilient, since the probability of faults increases with scale. With the advent of heterogeneous compute nodes that employ standard processors and GPGPUs, scientific algorithms need to match these architectures to extract the most performance. Key science

applications require novel mathematical models and system software that address the scalability and resilience challenges of current- and future-generation extreme-scale HPC systems. For more information go to: *www.csm.ornl.gov/srt/conferences/Scala/2013/.*

Energy Efficient Supercomputing (E2SC)

8:30am-5:30pm
Room: 210/212 (Morning)
Room: 603 (Afternoon)

Darren Kerbyson, Andres Marquez (Pacific Northwest National Laboratory), Kirk Cameron (Virginia Polytechnic Institute), Dimitrios Nikolopoulos (Queens University Belfast), Sudhakar 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 this, 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 various workloads. At the same time, high fidelity measurement techniques, typically ignored in data-center level measurement, are of importance for scalable and energy efficient inter-play in different layers of application, system software and hardware. This workshop will bring together researchers from different communities working on challenging problems in this area for a dynamic exchange of ideas. For more information go to: *hpc.pnl.gov/conf/e2sc/2013/.*

Extreme-Scale Programming Tools

9am-5:30pm
Room: 501

Michael Gerndt (Technical University Munich)

As we approach exascale, architectural complexity and severe resource limitations with respect to power, memory and I/O make tools support in debugging and performance optimization more critical then ever before. However, the challenges mentioned above also apply to tool development and, in particular, raise the importance of topics such as automatic tuning and methodologies for exascale tools-aided application development. This workshop will serve as a forum for application, system, and tool developers to discuss the requirements for future exascale-enabled tools and the roadblocks that need to be addressed on the way. We also highly encourage application developers to share their experiences with using the tools.

The workshop is organized by the Virtual Institute - High Productivity Supercomputing, an international initiative of HPC programming-tool builders. The event will also focus on the community-building process necessary to create an integrated toolsuite ready for an exascale software stack. For more information go to: *www.vi-hps.org/events/other/espt-sc13.html.*

Python for High Performance and Scientific Computing (PyHPC 2013)

9am-5:30pm
Room: 505

Andreas Schreiber (German Aerospace Center (DLR)), William Scullin (Argonne National Laboratory)

Python is an established, high-level programming language with a large community of users in academia and industry. It is a general-purpose language adopted by many scientific applications such as computational fluid dynamics, bio-molecular simulation, artificial intelligence, statistics, data analysis, scientific visualization etc. More and more industrial domains are turning towards it as well, such as aeronautics, robotics, semiconductor manufacturing, automotive solutions, telecommunication, finance, and games. The use of Python for scientific, high performance parallel, and distributed computing is increasing, especially, the use of Python for data analysis and big data processing. Traditionally, system administrators are using Python for automating tasks. Since Python is extremely easy to learn with a very clean syntax, it is well-suited for education in scientific computing, mathematics, and other disciplines. For more information go to: *www.dlr.de/sc/pyhpc2013.*

The 6th Workshop on HPC Finance

9am-5:30pm
Room: 601

Matthew Dixon (University of San Francisco), Andrew Sheppard (Fountainhead), David Daly (IBM Research), Jose Moreira (IBM Research)

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 learnings in the application of high performance computing to computational finance among researchers, scholars, vendors and practitioners, and to encourage and highlight collaborations between these groups in

addressing high performance computing research challenges. For more information go to: *ewh.ieee.org/conf/whpcf/.*

The 8th Parallel Data Storage Workshop (PDSW)

9am-5:30pm
Room: 503

Robert B. Ross (Argonne National Laboratory), Garth Gibson (Carnegie Mellon University), John Bent (EMC 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 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. This 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, wide area file systems, metadata intensive workloads, autonomics for HPC storage, virtualization for storage systems, archival storage advances, resource management innovations, storage systems for big data and analytics, and incorporation of emerging storage technologies. For more information go to: *www.pdsw.org.*

WOLFHPC: Workshop on Domain-Specific Languages and High-Level Frameworks for HPC

9am-5:30pm
Room: 506

Organizers: Sriram Krishnamoorthy (Pacific Northwest National Laboratory), J. 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. For example, very few high-level frameworks can make effective use of accelerators like GPUs and FPGAs. 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. For more information go to: *hpc.pnl.gov/conf/wolfhpc/2013/.*

The 2nd International Workshop on Data Intensive Scalable Computing Systems (DISCS)

9am-6pm
Room: 504

Yong Chen (Texas Tech University), Philip C. Roth (Oak Ridge National Laboratory), Xian-He Sun (Illinois Institute of Technology)

Existing HPC systems are designed primarily for workloads requiring high rates of computation. However, the widening performance gap between processors and storage, and trends toward higher data intensity in scientific and engineering applications, suggest there is a need to rethink HPC system architectures, programming models, runtime systems, and tools with a focus on data intensive computing. The 2nd International Workshop on Data Intensive Scalable Computing Systems (DISCS) builds on the momentum generated by its predecessor workshop, providing a forum for researchers and other interested people in the areas of data intensive computing and high performance parallel computing to exchange ideas and discuss approaches for addressing Big Data challenges. The workshop includes a keynote address and presentation of peer-reviewed research papers, with ample opportunity for informal discussion throughout the day. More information about the DISCS workshop can be found on the workshop website at data. *cs.ttu.edu/discs/.*

Friday, November 22

1st International Workshop on Software Engineering for HPC in Computational Science and Engineering (SE-HPCCSE)

8:30am-12pm
Room: 210/212

Jeffrey Carver (University of Alabama), Selim Ciraci (Pacific Northwest National Laboratory), Neil Chue Hong (University of Edinburgh)

Due to the availability of hardware and increasing computing requirements, there is a demand to utilize HPC, including GPGPUs and computing clusters, for computational science & engineering (CSE) applications. Unfortunately, developing HPC software is not an easy task. 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 various development tasks including: validation and verification, design, requirements management and maintenance. Development of HPC CSE software requires tailoring of SE tools and methods developed for more traditional software applications. 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 generate an agenda to improve tools and prac-

tices for developing HPC CSE software. For more information go to: *sehpccse13.cs.ua.edu.*

3rd International Workshop on HPC, Networking and Analytics for the Power Grid

8:30am-12pm
Room: 601/603

Daniel G. Chavarría, Bora Akyol, Zhenyu (Henry) Huang (Pacific Northwest National Laboratory)

This workshop promotes the use of high performance computing and networking for power grid applications. Technological and policy changes make this an urgent priority.

Sensor deployments on the grid are expected to increase geometrically in the immediate future, while the demand for clean energy generation is driving the use of non-dispatchable power sources such as solar and wind. New demands are being placed on the power infrastructure due to the introduction of plug-in vehicles. These trends reinforce the need for higher fidelity simulation of power grids, and higher frequency measurement of their state.

Traditional grid simulation and monitoring tools cannot handle the increased amounts of sensor data or computation imposed by these trends. The use of high performance computing and networking technologies is of paramount importance for the future power grid, particularly for its stable operation in the presence of intermittent generation and increased demands placed on its infrastructure. For more information go to: *gridoptics.pnnl.gov/sc13.*

8th Workshop on Virtualization in High-Performance Cloud Computing (VHPC '13)

8:30am-12pm
Room: 503/504

Michael Alexander (Vienna University of Technology), Gianluigi Zanetti (CRS4)

The 8th Workshop on Virtualization in High-Performance Cloud Computing (VHPC '13) brings together HPC operators and users from the scientific community with industrial technology providers to foster discussion, collaboration and mutual exchange of knowledge and practical experiences. With virtualization being currently deployed to scientific data centers, topics including management, performance and HPC I/O in virtualized environments will be looked at in talks and demonstrations followed by a panel discussion. For more information go to: *vhpc.org.*

Exascale MPI

8:30am-12pm
Room: 201/203

Stefano Markidis (KTH Royal Institute of Technology), Erwin Laure (KTH Royal Institute of Technology), Jesper Larsson Träff (Technical University of Vienna), Lorna Smith (Edinburgh Parallel Computing Center), Mark Parsons (Edinburgh Parallel Computing Center)

The MPI design and its main implementations have proved surprisingly scalable. Some issues that hampered MPI scalability have been addressed in the MPI 2.1-2.2 definition process, and continued into MPI 3.0. For this and many other reasons MPI is currently the de-facto standard for HPC systems and applications. However, there is a need for re-examination of the Message Passing (MP) model and for exploring new innovative and potentially disruptive concepts and algorithms, partially to explore other roads than those taken by the recently released MPI 3.0 standard. The aim of this workshop is to bring together researchers and developers to present and discuss innovative algorithms and concepts in the MP programming model, in particular related to MPI. Possible workshop topics include innovative algorithms for: collective operations, topology mapping, scheduling/routing to avoid network congestion, "fault-tolerant" communication, interoperability of MP and PGAS models, and integration of task-parallel models. For more information go to: *pdc.kth.se/exampi13.*

Obtaining Bitwise Reproducible Results: Perspectives and Latest Advances

8:30am-12pm
Room: 301/302/303

Noah Clemons (Intel Corporation)

It is a widely known HPC fact that many optimizations and scheduling techniques require a change in the order of operations, creating results that are not bitwise reproducible (BWR). This workshop will bring together members of the HPC community who are affected by this non-reproducibility phenomenon in various ways. A number of leading experts from numerical software tools, national laboratories, academia, the military, and both open-source and commercial software development, will showcase the latest developments concerning their favorite HPC tool or research method solving a BWR problem. Each will proctor a 30 minute hands-on lab module for the audience to work on.

VISTech Workshop: Visualization Infrastructure and Systems Technology

8:30am-12pm
Room: 205/207

Kelly Gaither, Brandt Westing (Texas Advanced Computing Center), Jason Leigh (University of Illinois at Chicago), Falko Kuester (University of California, San Diego), Eric Wernert (Indiana University), Aditi Majumder (University of California, Irvine)

Human perception is centered on the ability to process information contained in visible light, and our visual interface is a tremendously powerful data processor. Every day we are inundated with staggering amounts of digital data. For many types of computational research, 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 by studying 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 workshop will provide a space where 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. For more information go to: www.tacc.utexas.edu/web/sc13-workshop/vistech.

Workshop on Large-Scale Data Analytics (LSDA)

8:30am-12pm
Room: 501/502

Chaitan Baru, Natasha Balac (San Diego Supercomputer Center)

While advances in technology have led to an unprecedented ability to collect, compute, and store huge amounts of data, the ability to extract useful information from raw data is becoming progressively more challenging. Large-scale data analytics offers challenges along multiple dimensions. It requires the ability to deal with large volumes of data from diverse sources; field systems that perform well with diverse data and scale extremely well with increasing data volumes; design and build efficient and effective predictive models that can accommodate diverse data types; and develop benchmarks for systems and analytical processes to enable comparisons among alternative approaches at different data volumes. This workshop will provide a venue for discussing the challenges and solutions in this area, including keynote presentations, lightning talks, and a group discussion.

Birds of a Feather

Birds of a Feather/Meetings

Don't just observe, ENGAGE! The Birds-of-a-Feather (BOF) sessions are among the most interactive, popular, and well-attended sessions of the SC Conference Series. The BOF sessions provide a non-commercial, dynamic venue for conference attendees to openly discuss current topics of focused mutual interest within the HPC community with a strong emphasis on audience-driven discussion, professional networking and grassroots participation. SC13 will continue this tradition with a full schedule of exciting, informal, interactive sessions focused around a variety of special topics of mutual interest.

There were 113 BOF proposals submitted to SC13 for peer review. A committee of 22 members selected 74 BOF sessions to be held this year. These sessions involve a wide range of HPC-related topics such as accelerators, programming models, big data, and education. Sessions will be held Tuesday, November 19 and Wednesday, November 20-—at lunchtime and in the evening—and Thursday, November 21 at lunchtime. See the SC13 schedule for detailed information about which sessions will be held and when.

Tuesday, November 19

ACM SIGHPC Annual Members Meeting
12:15pm-1:15pm
Room: 605

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.

Birds of a Feather

Tuesday, November 19

Best Practices for Commissioning Liquid Cooling Infrastructure
12:15pm-1:15pm
Room: 404

Primary Session Leader: David Martinez (Sandia National Laboratories)

Secondary Session Leaders: Thomas Durbin (National Center for Supercomputing Applications), Michael Ellsworth (IBM)

The Energy Efficient HPC Working Group (http://eehpcwg.lbl.gov/) is developing a methodology for commissioning liquid-cooling infrastructure for HPC sites deploying liquid-cooled systems. Proper infrastructure operation is critical for liquid-cooled HPC equipment because safety margins are very small and cooling fluid flow cannot be disrupted without causing a system outage and/or damage to computing equipment. For anyone involved in liquid cooling facilities fit-up or retrofit, come to this BOF and hear lessons learned presentations from Steve Hammond, NREL, Anna Maria Bailey, LLNL and Detlef Lebrenz, LRZ. You will also be asked to provide your lessons learned and review the draft methodology.

Building on the European Exascale Approach
12:15pm-1:15pm
Room: 210/212

Primary Session Leader: Mark Parsons (EPCC, The University of Edinburgh)

Secondary Session Leaders: Alex Ramirez (Barcelona Supercomputing Center), Thomas Lippert (Juelich Supercomputing Centre)

The drive to develop, produce and exploit exascale platforms has led to distinct strategies across continents: differences in expertise and drivers defining the approaches taken. The focus of a few years ago on identifying the challenges has now moved on to early results and prototypes.

In Europe, combined funding of €25 million has been invested in three complementary research projects: CRESTA, DEEP and Mont-Blanc. These projects are now well established, making a comparison of Europe's exascale approaches with that of other countries and continents particularly timely. This BOF will debate these approaches and look to build cross continent communities.

Collaborative Opportunities with the Open Science Data Cloud
12:15pm-1:15pm
Room: 205/207

Primary Session Leader: Robert Grossman (University of Chicago)

Secondary Session Leader: Heidi Alvarez (Florida International University)

Scientists in a wide variety of disciplines are producing unprecedented volumes of data that is transforming science. Unfortunately, many scientists are struggling to manage, analyze, and share their medium to large size datasets.

The Open Science Data Cloud (OSDC) was developed to fill this gap. It is a cloud-based computing infrastructure that allows researchers to manage, analyze, integrate and share medium to large size scientific datasets. It is operated and managed by the not-for-profit Open Cloud Consortium.

Come to this Session to learn more about the OSDC and how you can use the OSDC for your big data research projects.

Getting Scientific Software Installed: Tools and Best Practices
12:15pm-1:15pm
Room: 708/710/712

Primary Session Leader: Kenneth Hoste (Ghent University)

Secondary Session Leader: Andy Georges (Ghent University)

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

EasyBuild (http://hpcugent.github.com/easybuild/), an open-source build and installation framework written in Python that allows to implement build procedures of (scientific) software in so-called easyblocks (Python modules that 'plug in' to the EasyBuild framework), will be briefly presented.

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

Hadoop 2 and HPC: Beyond MapReduce

12:15pm-1:15pm
Room: 507

Primary Session Leader: Martha Dumler (Cray Inc.)

Secondary Session Leaders: Howard Pritchard (Cray Inc.), Bill Sparks (Cray Inc.)

Apache Hadoop 2 greatly extends the capabilities of the Hadoop Ecosystem, and its potential applicability to large scale HPC Data Analytics and also expands the applications/usability of HPC systems.

Integration of NVM Technologies in HPC Architectures: Challenges and Opportunities

12:15pm-1:15pm
Room: 601/603

Primary Session Leader: Dirk Pleiter (Forschungzentrum Juelich / Juelich Supercomputing Centre)

Secondary Session Leaders: Robert Germain (IBM), Felix Schuermann (EPFL)

I/O performance of HPC systems is a major challenge on the path towards exascale. As of today the computing performance continues to improve at a faster path than the performance of the I/O subsystem. Architectural opportunities based on the use of existing non-volatile memory technologies and those arising from advances in NVM technologies can play a role in addressing these challenges. The Blue Gene Active Storage (BGAS) architecture is an example of this approach as it realizes an active storage concept within a top-class HPC architecture. A key question will be how applications can utilize these technologies and this architecture.

Maximize Data Center Power Efficiency through Energy Aware Computing

12:15pm-1:15pm
Room: 201/203

Session Leader: Luigi Brochard (IBM)

Tens of thousands of data centers now exist, most of them consuming vast amounts of energy. While energy efficiency varies widely from company to company, some centers are wasting massive amounts of the electricity they pull of the grid. In this session, we'll discuss all aspects of this problem: how to control and reduce power consumption and how to reduce the power consumed by the cooling. Through the use of green IT strategies such as power usage optimization and direct water cooling, administrators can maximize application workload throughput while significantly reducing data center costs.

Open MPI State of the Union

12:15pm-1:15pm
Room: 301/302/303

Primary Session Leader: Jeffrey Squyres (Cisco Systems)

Secondary Session Leader: George Bosilca (University of Tennessee)

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

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

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

Python for High Performance and Scientific Computing

12:15pm-1:15pm
Room: 401/402/403

Primary Session Leader: William Scullin (Argonne National Laboratory)

Secondary Session Leaders: Andreas Schreiber (German Aerospace Center (DLR)), Andrew Terrel (University of Texas at Austin)

This BOF is intended to provide current and potential Python users and tool providers in the high performance and scientific computing communities: a forum to talk about their current projects; ask questions of experts and leading vendors; explore methodologies; delve into issues with the language, modules, tools, and libraries; build community; discuss the concerns of the community with advocates and active developers present; and help continue the path forward.

Scalable Adaptive Graphics Environment (SAGE) for Global Collaboration

12:15pm-1:15pm
Room: 501/502

Primary Session Leader: Jason Leigh (University of Illinois at Chicago)

Secondary Session Leaders: Maxine Brown (University of Illinois at Chicago), Luc Renambot (University of Illinois at Chicago)

SAGE, the Scalable Adaptive Graphics Environment, receives major funding from the National Science Foundation to provide the scientific community with persistent visualization and collaboration services for global cyberinfrastructure. SAGE is

a widely used open-source platform and the de facto operating system, or framework, for managing "big data" content on scalable-resolution tiled display walls. The SC BOF provides an unparalleled opportunity for the global SAGE user community, and potential users, to meet, review current development efforts, share community-developed use cases and applications, and have dynamic discussions on user requirements and future roadmaps.

Seventh Graph500 List

12:15pm-1:15pm
Room: 705/707/709/711

Primary Session Leader: Richard Murphy (Micron Technology, Inc)

Secondary Session Leaders: David Bader (Georgia Institute of Technology), Andrew Lumsdaine (Indiana University)

Large-scale data analytics applications represent increasingly important workloads, but most of today's supercomputers are ill suited to them. Backed by a steering committee of over 30 international HPC experts from academia, industry, and national laboratories, Graph500 works to establish large-scale benchmarks that are representative of these workloads. This BOF will unveil the seventh Graph500 list, and an improved second benchmark. We will further explore the new energy metrics for the Green Graph500 and unveil the second Green Graph500 list.

The Lustre Community: At the Intersection of HPC and Big Data

12:15pm-1:15pm
Room: 405/406/407

Primary Session Leader: Galen Shipman (OpenSFS & Oak Ridge National Laboratory)

Secondary Session Leader: Hugo Falter (European Open File System & ParTec)

Lustre is the leading open source file system for HPC. Since 2011 Lustre has transitioned from a single vendor focus to a community developed file system with contributors from around the world. As a result, Lustre is now more widely used and in more mission-critical installations than ever. Many organizations are leveraging Lustre for both traditional HPC workloads as well as emerging Big Data workloads. At this year's Lustre Community BOF the worldwide community of Lustre developers, administrators, and solution providers will gather to discuss new challenges and corresponding opportunities emerging at the intersection of HPC and Big Data.

The Role of Software-Defined Networking in the Big Data Equation

12:15pm-1:15pm
Room: 703

Primary Session Leader: Bithika Khargharia (Extreme Networks)

Secondary Session Leader: Glenn Ricart (US Ignite)

While both HPC and its commanding processing power play a role in Big Data analytics, the real challenges are storage, fast migration, and fast access to all of that data - issues in which the underlying network architecture plays a critical role. As we are rapidly transitioning toward a fully-networked world, SDN has huge potential to address the evolving connectivity requirements of Big Data. SDN offers features that assist with management, integration, and analysis of Big Data. This session will examine the four fundamental questions users will need to answer to enable software-defined networks to support a fully mobile world.

TORQUE: Where Are We Now and Where Are We Going

12:15pm-1:15pm
Room: 503/504

Primary Session Leader: Kenneth Nielson (Adaptive Computing)

Secondary Session Leaders: David Beer (Adaptive Computing), Jill King (Adaptive Computing)

The purpose of this session is to present to the community: (1) what Adaptive has contributed to the community over the last year; (2) present Adaptive Computing's contribution plans to move the technology forward; and (3) get feedback from the TORQUE community concerning those plans and solicit suggestions for new development.

Asynchronous and Adaptive Parallel Programming with Charm++

5:30pm-7pm
Room: 702/704/706

Primary Session Leader: Laxmikant Kale (University of Illinois at Urbana-Champaign)

Secondary Session Leader: Eric Bohm (University of Illinois at Urbana-Champaign)

A BOF for the community interested in parallel programming using Charm++, Adaptive MPI, and the associated ecosystem (mini-languages, tools, etc.), along with parallel applications developed using them. Intended to engage a broader audience and drive adoption.

Charm++ is a parallel programming system with increasing usage. Next to MPI (and now, possibly OpenMP) it is one of the most used systems deployed on parallel supercomputers, using a significant fraction of CPU cycles. A unified programming model with multicore and accelerator support, its abilities include: dynamic load balancing, fault tolerance, latency hiding, interoperability with MPI, and overall support for adaptivity and modularity.

Campus Bridging with XSEDE and Globus Online

5:30pm-7pm
Room: 601/603

Primary Session Leader: Steve Tuecke (University of Chicago)

Secondary Session Leader: Rachana Ananthakrishnan (University of Chicago)

As science becomes more computation- and data-intensive, computing needs often exceed campus capacity. Thus campuses desire to scale from the local environment to other campuses, to national cyberinfrastructure providers such as XSEDE, and/or to cloud providers. But given the realities of limited resources, time, and expertise, campus bridging methods must be exceedingly easy to use. This BOF will explore Globus Online's transfer and sharing tools in the XSEDE setting, which address the important campus bridging use case of moving, sharing, and synchronizing data across institutional boundaries, achieving ease of use for researchers and ease of administration for campus IT staff.

Defining BigData: Industry Views on Real-Time Data, Analytics, and HPC Technologies to Bring Them Together

5:30pm-7pm
Room: 405/406/407

Primary Session Leader: Tony Allen (PayPal)

Secondary Session Leader: Barney Maccabe (Oak Ridge National Laboratory)

PayPal, Twitter, Google, Map-D, Bank of America and Oak Ridge National Laboratory will discuss opportunities for HPC in the emerging area of industry real-time decision making and analytics for activities like fraud detection, dynamic pricing, and web analytics. This BOF presents an opportunity for industry experts to explain some of the challenges and current solutions they are using to solve these real-world Big Data problems to the HPC community in the hope of sharing experiences, listing best practices, and inspiring future collaborations.

Early Experiences Developing and Debugging on the Intel® Xeon Phi™ Coprocessor

5:30pm-7pm
Room: 401/402/403

Primary Session Leader: Chris Gottbrath (Rogue Wave Software)

Secondary Session Leaders: Dave Hiatt (Citigroup Inc.)

This BOF will highlight user experiences with development and debugging on systems which incorporate the Intel® Xeon Phi™ Coprocessor. While getting a program that runs on Linux-x86 running on the Xeon Phi is relatively easy, deeper changes are often required to take full advantage of multi-core processors. Frequently, multi-threading needs to be added or reworked to allow higher levels of thread concurrency. Vectorization, memory alignment, and data locality also play important roles in optimization. This BOF will start with 5 brief presentations where speakers from labs and industry share their experiences and a moderated discussion of the topics raised.

Eclipse Parallel Tools Platform (PTP)

5:30pm-7pm
Room: 210/212

Primary Session Leader: Greg Watson (IBM)

Secondary Session Leader: Beth Tibbitts (N/A)

The Eclipse Parallel Tools Platform (PTP) is an open-source project providing a robust, extensible workbench for the development of parallel and scientific codes. PTP makes it easier to develop, build, run, optimize, and debug parallel codes on a variety of remote clusters using a single unified interface. PTP includes support for MPI, OpenMP, UPC, Fortran, and other languages as well.

The BOF will consist of brief demos and focused discussions about PTP, and an overview of new and upcoming features. Information obtained from participants will be an important input into the development of new functionality and enhancements in future releases.

Energy Efficient High Performance Computing

5:30pm-7pm
Room: 205/207

Session Leader: Kurt Keville (MIT)

Energy efficiency has become a principal driving factor in HPC. With processors developed for smartphones now capable of delivering higher performance per Watt than even the densest servers, an emerging trend is the adoption of embedded computing technologies for HPC. Many challenges remain before the community can fully exploit the embedded space.

Much of the software stack on which the community relies, such as Fortran compilers and optimized libraries, still do not exist for embedded processors. This BOF will feature talks from major proponents in this area, including developers with Open Source Hardware compliant products built for the embedded space.

G8 Extreme Scale and Big Data Program: Progress on Current Exascale Projects and Future Funding Opportunities

5:30pm-7pm
Room: 501/502

Primary Session Leader: William Tang (Princeton Plasma Physics Laboratory)

Secondary Session Leader: Franck Cappello (Argonne National Laboratory)

There are six G8-funded exascale research projects in progress (2011-14) that address key scientific challenges from the Climate, Fusion Energy, Seismology, and Bio-molecular domain applications. This BOF will provide new information on the G8 international HPC research program, including plans beyond 2014. Representatives from funding agencies supporting the G8 projects will present the funding landscape, and members from the current exascale projects will highlight progress and experiences. This BOF will enable the HPC community to get a snapshot and engage in discussions of the exciting work and research opportunities in the G8 extreme scale and large data HPC activities.

HPC and the Web

5:30pm-7pm
Room: 703

Primary Session Leader: Annette Greiner (Lawrence Berkeley National Laboratory)

Secondary Session Leaders: Rion Dooley (University of Texas at Austin), Shreyas Cholia (Lawrence Berkeley National Laboratory)

The HPC and the Web BOF will provide infrastructure builders, service providers, gateway developers, and other stakeholders a venue to exchange visions of how we can collectively leverage the web to enhance HPC in the next decade. Consisting entirely of interactive discussion between participants, it will offer the opportunity to engage in topics cutting across many different areas. These include, but are not limited to, development tools, APIs, data locality, federated identity, proprietary rights, job scheduling, and recommendation engines for scholarly activity.

Library of Mini-Applications for Exascale Component-Based Performance Modeling

5:30pm-7pm
Room: 708/710/712

Primary Session Leader: Marie-Christine Sawley (Intel SAS)

Secondary Session Leaders: Naoya Maruyama (RIKEN), John Shalf (Lawrence Berkeley National Laboratory)

This BOF is focused on community building around a coordinated effort between some of the most prominent groups in the US, in Japan with Riken for the Exascale Feasibility Project, and in Europe with EESI2 and groups such as the Intel Exascale Labs and their partners. All have experience in mini- applications and are developing roadmaps for forthcoming years. We propose to discuss how to join efforts around a library of new kernels, mini applications representative of new usage models, or built around new innovative algorithms, and use them to extend the capacities for exascale component-based performance modeling.

MPICH: A High-Performance Open-Source MPI Implementation

5:30pm-7pm
Room: 201/203

Primary Session Leader: Rajeev Thakur (Argonne National Laboratory)

Secondary Session Leaders: Pavan Balaji, Rusty Lusk (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 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 Goes Heterogeneous With OpenMP 4.0

5:30pm-7pm
Room: 301/302/303

Primary Session Leader: Michael Wong (IBM)

Secondary Session Leaders: Bronis de Supinski (Lawrence Livermore National Laboratory), Barbara Chapman (University of Houston)

During the past few years, the OpenMP ARB has worked hard to add significant new features to the OpenMP specification. These include features for large core counts and for execution on platforms with heterogeneous cores and accelerators. The extensive set of new features that have been defined in the recently released OpenMP 4.0 will be discussed by key members of the Language Committee. The intended audience consists of OpenMP users and implementers. The format will include discussions on the use of the new features, and a lively cross fire session with vendors on their implementation plans for OpenMP 4.0.

Science and Scientific Workflows: Putting Workflows to Work

5:30pm-7pm
Room: 404

Primary Session Leader: Marlon Pierce (Indiana University)

Secondary Session Leaders: Mats Rynge (University of Southern California), Suresh Marru (Indiana University)

The purpose of this session is for computational scientists, scientific workflow software researchers, cyberinfrastructure architects, and cyberinfrastructure operations leaders to share ideas on how to best use cyberinfrastructure such as the NSF's XSEDE to tackle the most challenging and scientifically valuable distributed computing problems. The emphasis will be on discussing open, unsolved scientific problems that will benefit from using multiple XSEDE, Open Science Grid, campus, and international cyberinfrastructure resources in coordinated fashion. The BOF will be mediated by the XSEDE Workflow Community Applications Team.

SUPReMM: Comprehensive Open Source Resource Management

5:30pm-7pm
Room: 507

Primary Session Leader: Thomas Furlani (SUNY at Buffalo)

Secondary Session Leaders: James Browne (University of Texas at Austin), Abani Patra (SUNY at Buffalo)

SUPReMM is the first comprehensive open-source tool chain to provide resource management capabilities to users and managers of HPC systems. SUPReMM was created by integrating data collected by the TACC_Stats monitoring system with the XDMoD (XSEDE Metrics on Demand, https://xdmod.ccr.buffalo.edu) analysis and reporting system. SUPReMM provides detailed reports on the performance (CPU and memory usage, swapping/paging activities, network bandwidth, filesystem usage, and interconnect fabric traffic) of all jobs running on open-source Linux based systems. This BOF will have an interactive, user participant demonstration and a discussion session focusing on future development of this NSF funded effort.

TOP500 Supercomputers

5:30pm-7pm
Room: Mile High

Session Leader: Erich Strohmaier (Lawrence Berkeley National Laboratory)

The TOP500 list of supercomputers serves as a "Who's Who" in the field of 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 42nd TOP500 list will be published in November 2013 just in time for SC13.

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.

Towards Exascale Runtime Systems: Challenges and Opportunities

5:30pm-7pm
Room: 503/504

Primary Session Leader: Hans-Christian Hoppe (Intel Corporation)

Secondary Session Leader: Robert Wisniewski (Intel Corporation)

Exascale runtime systems will need to perform careful orchestration of millions of hardware and software elements in a near-optimal way to reach the overall system performance and power targets. The unprecedented number of components, their dynamic and nonlinear behavior, and the increasing complexity of applications greatly raise the bar. The BOF should kick off a community of hardware and OS specialists, runtime system experts, programming model & tools providers that work towards a solution. Leaders in the fields will briefly discuss the problem and start a discussion with the audience on the way towards a truly exascale capable runtime system.

Trends in Small HPC Center Management

5:30pm-7pm
Room: 705/707/709/711

Primary Session Leader: David Stack (University of Wisconsin-Milwaukee)

Secondary Session Leader: Beth Anderson (Intel Corporation)

This session will explore the evolving nature of campus-scale, HPC capacity clusters. Attendees will provide their reactions to the findings of a pre-conference survey that investigated the compute, storage, submit/compile and scheduler environments that are common in systems of this size. Trend data from three consecutive years will be available for discussion. Attendees are also invited to share how scientific software is provided to the end users of their clusters, which operating systems run on the worker nodes and how the system configurations are deployed and maintained.

Wednesday, November 20

Application Migration and Performance Expectation for Manycore Programming

12:15pm-1:15pm
Room: 205/207

Primary Session Leader: Jim Cownie (Intel Corporation)

Session Leaders: Jim Cownie (Intel Corporation), John Michalakes (National Renewable Energy Laboratory), Shuo Li (Intel Corporation)

As we enter the era of manycore computing, hardware led innovations have offered us a plethora of parallel paradigms, programming models and languages. Application migrations across different programming modes has become a major challenge and created issues that involve language syntax, programmability, runtime support optimization, etc. In this BOF, we invite the manycore application developers of any programming models and the programming language implementers, compiler and library writers to join us in discussing the issues fundamental manycore programming language design, expectation of the performance portability and appreciate the challenges in runtime library support.

Applications of LLVM to HPC

12:15pm-1:15pm
Room: 201/203

Primary Session Leader: Michael McCool (Intel Corporation)

Secondary Session Leader: Hal Finkel (Argonne National Laboratory)

The LLVM system provides a state-of-the-art optimizing compiler infrastructure, supporting both static and dynamic compilation, that can be used for a variety of production and research purposes in HPC. Implementations of languages used in HPC production such as C, C++, and OpenCL use LLVM, and new experimental languages, or high performance versions of existing interpreted languages, can also be based on LLVM. This forum will bring together developers working on implementing languages for HPC based on the LLVM infrastructure with users of these languages.

At the Intersection of Big Data and Extreme Computing

12:15pm-1:15pm
Room: 503/504

Primary Session Leader: James Reaney (SGI)

Big Data and research computing go hand-in-hand, especially for research disciplines that operate at extreme scales. Current Big Data solutions can be improved by borrowing from HPC techniques. A panel of experts in bioinformatics, quantum chemistry, engineering mechanics / fluid dynamics, and weather modeling will review Big Data challenges in their respective disciplines and discuss where bottlenecks exist or where HPC techniques could be useful. Anyone interested in learning about Big Data challenges and applicability to specific research disciplines in particular should attend. A Q&A session with the audience will follow and audience participation in the BOF is strongly encouraged.

Chapel Lightning Talks 2013

12:15pm-1:15pm
Room: 601/603

Primary Session Leader: Sung-Eun Choi (Cray Inc.)

Secondary Session Leader: Bradford Chamberlain (Cray Inc.)

Are you a scientist considering a modern high-level language for your research? Are you a language enthusiast who wants to stay on top of new developments? Are you an educator considering using Chapel in your classroom? Are you already a Chapel fan and wondering what's in store for the future? Then this is the BOF for you!

In this BOF, we will hear "lightning talks" on community activities involving Chapel. We will begin with a talk on the state of the Chapel project, followed by a series of talks from the broad Chapel community, wrapping up with Q&A and discussion.

Community MOOC's for Computing Technologies and Applications

12:15pm-1:15pm
Room: 705/707/709/711

Primary Session Leader: Geoffrey Fox (Indiana University)

This BOF discusses the possible value of Massive Open Online Courses (MOOC) for training and education in computing technologies and applications (HPC, Clouds, Software, Algorithms, Applications, Infrastructure, Systems Administration, Computational Science). We will explore community interest in building and using such MOOC's in a new "X-MOOC" repository (which could be set up as a result of this BOF). Speakers will briefly describe using and building MOOC's including open source learning management software from Google and EdX. Comments from partners with commercial companies Udacity and Coursera will be presented. We will discuss how to support computing laboratories associated with MOOC's.

Heterogeneous Computing Platforms: Merging HPC and Embedded

12:15pm-1:15pm
Room: 708/710/712

Primary Session Leader: Markus Levy (Multicore Association)

Secondary Session Leader: Barbara Chapman (University of Houston)

This BOF brings together the multicore and manycore user/developer community to discuss the impact that multicore software will have given the fast emergence of the hardware chips containing homogeneous and heterogeneous cores. The focus is on programming models and the use of software standards as applied to multicore applications and the needs of the multicore community to be considered for future standard releases. The motivation behind this is to facilitate the merging of embedded and HPC applications.

High Precision Arithmetic Operations: Libraries and Applications

12:15pm-1:15pm
Room: 301/302/303

Primary Session Leader: Naohito Nakasato (University of Aizu)

Secondary Session Leaders: Fukuko Yuasa (High Energy Accelerator Research Organization)

The emergence of large-scale and high-speed parallel computing forces us to consider, on a new level, rounding errors in repeated arithmetic operations. Many scientific and engineering problems rely on the numerical stability of the algorithms used, even with conventional "double precision" arithmetic operations. The development of techniques that are truly high precision is important to solve very compute-intensive problems, while an extreme solution is to implement "high precision" arithmetic as hardware. In this BOF, researchers working on high precision arithmetic techniques and applications will meet and present recent progress. We will also discuss on the future developments relevant to HPC.

High-Performance Communications for High Performance Computing

12:15pm-1:15pm
Room: 404

Primary Session Leader: Jack Wells (Oak Ridge National Laboratory)

Secondary Session Leaders: Richard Coffey (Argonne National Laboratory), Michele De Lorenzi (Swiss National Supercomputing Centre)

This BOF will bring together leaders from supercomputing centers, universities, industry, and associated fields (like science journalists and HPC solution providers) to discuss the challenges in communicating the value of supercomputing for society at large. We will analyze the range of topics that need to be communicated, identify their respective audiences, and discuss possible strategies and practices for achieving the highest impact. 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.

HPC Job Scheduling Challenges and Successes from the PBS Community

12:15pm-1:15pm
Room: 702/704/706

Primary Session Leader: Bill Nitzberg (Altair)

Secondary Session Leader: Greg Matthews (CSC and NASA Ames)

Decades ago, the first HPC systems were single-user. As soon as the technology allowed sharing, job scheduling became a necessity (for the sanity of the users and the administrators). More than 20 years ago, NASA developed the PBS software to fill this need. Now, PBS is consistently among the top 3 most-reported technologies for HPC job scheduling (as reported by IDC). But scheduling is hard---every site has unique processes, unique goals, and unique requirements. Join fellow members of the community to share challenges and solutions and learn others' tips and tricks for scheduling in the modern world.

HPC Systems Engineering and Administration
12:15pm-1:15pm
Room: 501/502

Primary Session Leader: Adam Yates (Louisiana State University)

Secondary Session Leaders: William Scullin (Argonne National Laboratory), Adam Hough (Petroleum Geo-Services)

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 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.

INCITE and Leadership-Class Systems
12:15pm-1:15pm
Room: 703

Session Leader: Julia White (INCITE)

The INCITE program grants annual allocations of more than 5 billion core hours at the Argonne and Oak Ridge Leadership Computing Facility (LCF) centers. Come learn about INCITE and the LCF systems: a 27-petaflops Cray XK7 (Titan) and a 10-petaflops IBM Blue Gene/Q (Mira). We'll talk about successful proposals and discuss the work required to refactor and port applications to Titan and Mira. Join us for a presentation followed by an open forum and Q&A on topics ranging from LCF system architectures to INCITE award criteria.

Opportunities and Barriers for Computational Science Education
12:15pm-1:15pm
Room: 507

Primary Session Leader: Steven Gordon (Ohio Supercomputer Center)

It is critical that computational science concepts and techniques are introduced into the university curriculum. Yet this interdisciplinary field faces a number of obstacles to implementation. The XSEDE Education program offers some assistance to offset these obstacles. We will introduce model curricula, the nature of the implementation problems and discuss common problems and issues concerning computational science education with the audience.

PGAS: The Partitioned Global Address Space Programming Model
12:15pm-1:15pm
Room: 401/402/403

Primary Session Leader: Tarek El-Ghazawi (George Washington University)

Secondary Session Leaders: Lauren Smith (US Government)

The partitioned global address space (PGAS) programming model strikes a balance between the ease of programming due to its global address memory model and performance due to locality awareness. While developed for scalable systems, PGAS is gaining popularity due to the NUMA memory architectures on many-core chips. Some PGAS implementations include Co-Array Fortran, Chapel, UPC, X10, Phalanx, OpenShmem, Titanium and Habanero. PGAS concepts are influencing new architectural designs and are being incorporated into traditional HPC environments. This BOF will bring together developers, researchers and users for the exchange of ideas and information and to address common issues of concern.

The Open Community Runtime (OCR) Framework for Exascale Systems
12:15pm-1:15pm
Room: 405/406/407

Primary Session Leader: Vivek Sarkar (Rice University)

Secondary Session Leaders: Rob Knauerhase (Intel Corporation), Richard Lethin (Reservoir Labs, Inc.)

Extreme-scale and exascale systems impose new requirements on software to accommodate platforms with hundreds of homogeneous and heterogeneous cores, as well as energy, data movement and resiliency constraints within and across nodes. The goal of this BOF is to build on the momentum from the unveiling of the first OCR open-source release at SC12. We will discuss improvements to OCR from multiple contributors, examples of growing interest from the community, and a roadmap for future extensions. Through hands-on demonstrations of OCR code obtained from different programming models, we will demonstrate the efficacy of porting applications to the OCR execution model.

Total Power Usage Effectiveness: A New Take on PUE

12:15pm-1:15pm
Room: 210/212

Primary Session Leader: Michael Patterson (Intel Corporation)

Secondary Session Leaders: Chung-Hsing Hsu (Oak Ridge National Laboratory), Anna Maria Bailey (Lawrence Livermore National Laboratory)

The Energy Efficient HPC Working Group proposes two new metrics: ITUE (IT-power usage effectiveness), similar to PUE but "inside" the system and TUE (total-power usage effectiveness), which combine for a total efficiency picture. TUE provides a ratio of total energy, (internal and external support energy uses) and the specific energy used in the HPC. ORNL has alpha tested these new metrics and LLNL, NCAR, LBNL, NREL, VTech and LRZ have all expressed interest in beta testing. This BOF will provide a forum for reviewing the beta test results and also provide for further community feedback on the new metrics.

Big Data, Big Compute: Data-Intensive and Extreme-Scale Computing

5:30pm-7pm
Room: 301/302/303

Primary Session Leader: Lucy Nowell (DOE Office of Advanced Scientific Computing Research; NITRD High End Computing Interagency Working Group)

Secondary Session Leader: John West (DoD High Performance Computing Modernization Program; NITRD High End Computing Interagency Working Group)

This BOF will highlight the important union between the Big Data and HPC communities, and spark a discussion with attendees to explore ways in which hardware and software challenges faced by both communities align. The goals of the session are to examine synergies and challenges that have already been identified, gather inputs on additional ones, look for additional application areas that span both domains, and outline a few opportunities to fund research and development activities that benefit both areas. As common threads emerge, we will seek collaborations from the community to address the opportunities.

Codesign for the Department of Energy's Computational Science Community

5:30pm-7pm
Room: 702/704/706

Primary Session Leader: Richard Barrett (Sandia National Laboratories)

Secondary Session Leaders: Allen McPherson (Los Alamos National Laboratory), Charles Still (Lawrence Livermore National Laboratory)

The Department of Energy HPC community is actively engaged in codesign efforts as a means of ensuring that new architectures more effectively support mission critical workloads. The combinations of application codes, represented by proxy programs, algorithms, programming models, system software, and architecture provides a concrete and powerful means for interacting with vendors. In this session we will present and discuss concrete examples of the impact of these efforts. We invite members of the computational science community to participate in this session, as a means of learning about, forming collaborations, and influencing the direction of this work.

Cost-Benefit Quantification for HPC: An Inevitable Challenge

5:30pm-7pm
Room: 601/603

Primary Session Leader: Thomas Ludwig (DKRZ)

Secondary Session Leaders: Albert Reuther (MIT Lincoln Laboratory), Amy Apon (Clemson University)

HPC is an indispensable though expensive technology. While the discussion of cost has been part of the conversation for many years, only recently have we also seen the analyses of its benefits. Quantification is difficult and we still lack an economical model to evaluate the cost-benefit ratio, and effects to efficiency, in the industrial process and in scientific development.

This BOF will present approaches to this issue and foster a discussion between investors and users.

Better understanding of the methods to quantify and qualify the benefit and cost aspects will result in cost efficient computational science and engineering.

Creating a Training Roadmap for New Computational Consultants and Their Users

5:30pm-7pm
Room: 703

Primary Session Leader: Kimberley Dillman (Purdue University)

Secondary Session Leaders: Timothy Stitt (University of Notre Dame), Vincent Betro (University of Tennessee, Knoxville)

This session will primarily be interactive and focus on the development of a list of training topics and the appropriate level of training that will meet the needs of new "Computational Consultants" and the users they support. These training topics will be linked to "user models" that have been previously identified at a recent BOF session held during the XSEDE13 Conference. The concept of a computational consultant is not limited to a single program, such as the XSEDE Campus Champions, but is a much broader idea that can be embraced by both educational institutions and industry.

Critically Missing Pieces in Heterogeneous Accelerator Computing

5:30pm-7pm
Room: 401/402/403

Session Leader: Pavan Balaji (Argonne National Laboratory)

Heterogeneous architectures play a massive role in architecting the largest systems in the world. However, much of the interest in these architectures is an artifact of the hype associated with them. For such architectures to truly be successful, it is important that we look beyond this hype and learn what these architectures provide and what is critically missing. This continuing BOF series brings together researchers working on aspects of accelerator architectures---including data management, resilience, programming, tools, benchmarking, and autotuning---to identify critical gaps in the accelerator ecosystem.

Drilling Down: Understanding User-Level Activity on Today's Supercomputers

5:30pm-7pm
Room: 205/207

Primary Session Leader: Robert McLay (University of Texas at Austin)

Secondary Session Leader: Mark Fahey (University of Tennessee, Knoxville)

Let's talk real, no-kiddin' supercomputer analytics: drilling down to the level of individual batch submissions, users, and binaries. And we're not just targeting performance: we're 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, and will be the kickoff meeting for a new Special Interest Group.

Evolution of and Experiences with OpenACC

5:30pm-7pm
Room: 705/707/709/711

Primary Session Leader: Duncan Poole (NVIDIA)

Secondary Session Leaders: Michael Wolfe (NVIDIA), Barbara Chapman (University of Houston)

This BOF brings together the OpenACC user community to discuss the recent changes to the OpenACC standard. The focus this year is the major 2.0 update and gathering feedback from users on the new features and changes. We also plan to discuss the future road map, and in more detail, needs of the community to be considered for the next release. Additionally, discuss recent efforts in the development of an OpenACC benchmark. Lastly, we will discuss some of the research efforts towards building an open source compiler for the OpenACC standard.

Exascale IO Initiative: Progress Status

5:30pm-7pm
Room: 708/710/712

Primary Session Leader: Toni Cortes (Barcelona Supercomputing Center)

Secondary Session Leaders: Peter Braam (Xyratex), André Brinkmann (Johannes Gutenberg-Universität Mainz)

EIOW intends to architect and implement an open source, upper level I/O middleware system suitable for exascale storage. It intends to be primarily motivated by the requirements of the applications, management and system architectures, and to a lesser extent by the constraints and traditions of the storage industry. The resulting middleware system is targeting adoption in the HPC community by creation of or integration into various HPC software components, such as libraries. We also target adoption by storage vendors to layer this on either existing or new scalable high performance storage products.

OpenCL: Version 2.0 and Beyond

5:30pm-7pm
Room: 405/406/407

Primary Session Leader: Tim Mattson (Intel Corporation)

Secondary Session Leaders: Ben Bergen (Los Alamos National Laboratory), Simon McIntosh-Smith (University of Bristol)

OpenCL enables heterogeneous computing (e.g. a combination of CPUs, GPUs and coprocessors) without locking users into a single vendor's products. In the BOF, we will discuss the

latest developments in OpenCL including the new OpenCL 2.0 specification. We will then convene a panel representing divergent views on the future of OpenCL to engage in an interactive debate on where to take OpenCL. We will close the BOF with a survey of the audience to create "hard data" we can present to the OpenCL standards committee to help them address the unique needs of the HPC community.

OpenSHMEM: Further Developing a Standard for the PGAS & SHMEM Community

5:30pm-7pm
Room: 201/203

Primary Session Leader: Tony Curtis (University of Houston)

Secondary Session Leaders: Steve Poole (Oak Ridge National Laboratory)

The purpose of this BOF is to engage collaboration and input from users and developers of systems, libraries, and applications to further expand the open organization for OpenSHMEM.

The current API of OpenSHMEM is expected to develop with a richer feature set over time to accommodate advances in system design and scale. This BOF is an excellent face-to-face opportunity to provide your input into this ongoing process.

We invite attendees to present and discuss contributions to the ecosystem of OpenSHMEM tools and applications, and hardware solutions.

Reconfigurable Supercomputing

5:30pm-7pm
Room: 503/504

Primary Session Leader: Martin Herbordt (Boston University)

Secondary Session Leaders: Alan George (University of Florida), Herman Lam (University of Florida)

Reconfigurable supercomputing (RS) is characterized by hardware that adapts to match the needs of each application, offering unique advantages in speed per unit energy. With a proven capability of 2 PetaOPS at 12KW, RS has an important role to play in the future of high-end computing. Systems such as Novo-G in the NSF CHREC Center at Florida are rapidly moving towards production in multiple scientific and engineering domains. This BOF introduces concepts and 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.

Research Data Alliance (RDA) for HPC

5:30pm-7pm
Room: 501/502

Session Leader: Beth Plale (Indiana University)

The ubiquity of today's data is not just transforming what is, it is transforming what will be – laying the groundwork to drive new innovation. Today, research questions are addressed by complex models, by large data analysis tasks, and by sophisticated data visualization techniques, all requiring data. To address the growing global need for data infrastructure, the Research Data Alliance (RDA) was launched in FY13 as an international community-driven organization. We propose to bring together members of RDA with the HPC community to create a shared conversation around the utility of RDA for data-driven challenges in HPC.

Super-R: Supercomputing and R for Data-Intensive Analysis

5:30pm-7pm
Room: 404

Primary Session Leader: Weijia Xu (Texas Advanced Computing Center, University of Texas at Austin)

Secondary Session Leaders: George Ostrouchov (Oak Ridge National Laboratory and University of Tennessee), Hui Zhang (Pervasive Technology Institute, Indiana University)

R has become popular for data analysis in many fields, drawing power from its high-level expressiveness and numerous domain-specific packages. While R is clearly a "high productivity" language, it is not known as a "high performance" language. However, recent efforts have resulted in methods for effectively scaling R to the power of supercomputers. This BOF will consist of presentations at the intersection of supercomputing and R, followed by audience discussion to share experiences, needs, and questions. The ultimate goal is to help build a community of users and experts interested in applying R to solve data intensive problems on supercomputers.

Techniques and Strategies for Extreme Scale Debugging

5:30pm-7pm
Room: 210/212

Primary Session Leader: Chris Gottbrath (Rogue Wave Software)

Secondary Session Leader: Dong Ahn (Lawrence Livermore National Laboratory)

This BOF will bring together users interested in debugging parallel programs at large scale for the purposes of sharing experiences and disseminating best practices. There will be a number of brief 5-10 minute user-presentations focusing on their large scale debugging experiences, highlighting the strategies they used, as well as details on what worked and what didn't.

Seven members of the community representing sites such as LLNL, TU-Dresden, ANL, University of Wisconsin, and JSC have already expressed interest in speaking. The remainder of the time will be a moderated discussion among BOF participants.

The Green500 List and Its Evolution

5:30pm-7pm
Room: Mile High

Primary Session Leader: Wu Feng (Virginia Tech)

Secondary Session Leaders: Natalie Bates (Energy-Efficient HPC Working Group), Erich Strohmaier (Lawrence Berkeley National Laboratory)

The Green500, entering its seventh year, encourages sustainable supercomputing by raising awareness in the energy efficiency of such systems. This BOF will present (1) evolving metrics, methodologies, and workloads for energy-efficient HPC, (2) trends across the Green500, including the trajectory towards exascale, and (3) highlights from the latest Green500 List. In addition, we will discuss a collaborative effort between Green500, Top500, and EE HPC WG to improve power measurement while running a workload, such as HPL, and solicit feedback from the HPC community. We will close with an awards presentation, recognizing the most energy-efficient supercomputers in the world.

Women in HPC Around the World

5:30pm-7pm
Room: 507

Primary Session Leader: Rebecca Hartman-Baker (iVEC)

Secondary Session Leaders: Fernanda Foertter (Oak Ridge National Laboratory), Yashema Mack (University of Tennessee, Knoxville)

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 HPC. There are numerous systemic causes of women's low participation in STEM fields, which differ very little across the globe. In this BOF we bring together women from different continents to share their stories and experiences and determine ways we can boost the number of women around the globe enjoying the unique challenges and rewards of a career in HPC.

Thursday, November 21

Application Grand Challenges in the Heterogeneous Accelerator Era

12:15pm-1:15pm
Room: 401/402/403

Session Leader: Pavan Balaji (Argonne National Laboratory)

Accelerators have gained prominence as the next disruptive technology with a potential to provide a non-incremental jump in performance. However, the number of applications that have actually moved to accelerators is still limited because of many reasons, arguably the biggest of which is the gap in understanding between accelerator and application developers. This BOF is an application oriented session that aims to bring the two camps of application developers and accelerator developers head-to-head.

Computer Architecture Repositories for Open Simulation and Modeling Tools

12:15pm-1:15pm
Room: 404

Primary Session Leader: Daniel Mosse (University of Pittsburgh)

Secondary Session Leaders: Noel Wheeler (Laboratory for Physical Sciences), Bruce Childers (University of Pittsburgh)

The goal of this BOF is to engage academia, government and industry in a discussion about OCCAM (Open Curation for Computer Architecture Modeling), an open framework and repository for repeatable experimentation with simulation and tools for computer architecture. OCCAM is based on an interoperable, robust, community-supported simulation and emulation repository. By bringing together researchers, users, and implementers of computer architecture simulation, emulation, and modeling tools, through discussion, involvement, and feedback, OCCAM will be truly a community-supported.

Coprocessors, GPUs, MICs: Enticing Users to Jump onto Accelerators

12:15pm-1:15pm
Room: 405/406/407

Primary Session Leader: Laura Carriere (NASA Center for Climate Simulation)

Secondary Session Leader: Dan Duffy (NASA Center for Climate Simulation)

More and more, HPC sites are offering significant capacity increases via coprocessors, GPUs, MICs and other accelerators. But most user application codes were developed for traditional processors, and efforts for converting codes may be perceived as reducing resources for the primary mission. How can HPC sites entice users to make the jump to adapt codes for

accelerators? HPC sites will share their situations and approaches, compare notes, and determine interest in developing a community group to continue to share plans and strategies.

Discussing an I/O Framework to Accelerate Improvements in Application I/O Performance

12:15pm-1:15pm
Room: 601/603

Primary Session Leader: Paul Grun (Cray Inc.)

Secondary Session Leader: Sean Hefty (Intel Corporation)

In many cases, server I/O comprises a complex dance across a well-defined interface between an application and the underlying network. The common practice has been for coders to take the I/O interfaces as a given and to code to them at some cost in terms of application performance. This BOF brings application developers together with the open source networking community to discuss whether application I/O can be improved through a framework of interfaces allowing developers to achieve significant improvements in I/O efficiency with minimal effort. The outcome will provide direction to drive forward the development of open source I/O software.

Ethernet's Rate Roadmap

12:15pm-1:15pm
Room: 702/704/706

Primary Session Leader: John D'Ambrosia (Ethernet Alliance)

There is something comforting in predictability. With Ethernet predictability has been its rate progression - 10x increments from its initial 10Mb/s to 100Mb/s, to 1Gb/s to 10Gb/s, with little to no controversy. The simultaneous introduction of 40GbE and 100GbE effectively ended this legacy. It appears that Ethernet has abandoned its 10x increment legacy in favor of an x4 increment, leading to considerable speculation by some regarding the future trajectory of Ethernet. Will the next speed after 400Gb/s be 1.6Tb/s? This would seem to be the case if we were to simply multiple 400Gb/s by 4.

Harnessing Accelerator Technology for Next-Gen Sequencing Bioinformatics

12:15pm-1:15pm
Room: 503/504

Primary Session Leader: Shel Swenson (University of Southern California)

Secondary Session Leaders: Srinivas Aluru (Georgia Institute of Technology), David Bader (Georgia Institute of Technology)

Next generation sequencing (NGS) technologies enable rich and diverse applications but have found wide gaps between these technologies' potential and current computational reality. Research bridging these gaps delivers broad and immediate impact in both industry and academia. This BOF explores how biological sciences and HPC communities can combine efforts to solve pressing NGS problems using emerging and future computing platforms. Topics include identifying high-impact problems, discussing software-related challenges and building a community research agenda. This BOF part a broader NSF Software Institute Conceptualization project, "Software Institute for Accelerating Grand Challenge Science with Future Computing Platforms" (future-compute.usc.edu).

HPC Training Perspectives and Collaborations from PRACE (Europe), XSEDE (US) and RIKEN AICS (Japan)

12:15pm-1:15pm
Room: 501/502

Primary Session Leader: David Henty (Edinburgh Parallel Computing Centre at the University of Edinburgh)

Secondary Session Leaders: Scott Lathrop (University of Illinois at Urbana-Champaign), Mitsuhisa Sato (University of Tsukuba)

HPC training plays a vital role in fostering an ever-growing community of researchers and developers who conducts their work on supercomputers. In this BOF session, representatives from the Partnership for Advanced Computing in Europe (PRACE), the Extreme Science and Engineering Discovery Environment (XSEDE), and the RIKEN Advanced Institute for Computational Science (AICS) will each present their HPC training programs in Europe, the US, and Japan, respectively. The presentations will be followed by a discussion of the respective experiences and lessons learned, opportunities for sharing materials, processes for assessing and certifying knowledge acquisition, and a discussion of opportunities for collaboration.

New Developments in the APGAS Programming Model and X10

12:15pm-1:15pm
Room: 507

Primary Session Leader: Vijay Saraswat (IBM)

Secondary Session Leaders: David Grove (IBM), Shigeru Chiba (University of Tokyo)

This BOF provides a forum to discuss the APGAS (Asynchronous Partitioned Global Address Space) programming model and the X10 programming language. APGAS is one of the first models for concurrency and communication that has integrated constructs for asynchrony within a node and messaging between nodes. The APGAS programming model is realized natively in the X10 programming language. Via the new APGAS runtime, the APGAS model is now also directly available to C++ programmers who desire to write efficient, asynchronous, multi-node, programs.

The BOF will also introduce Resilient X10, an extension of APGAS with support for fault tolerance and place elasticity.

Petascale Systems with Intel Xeon Phi Co-Processors

12:15pm-1:15pm
Room: 301/302/303

Primary Session Leader: Ilene Carpenter (National Renewable Energy Laboratory)

Secondary Session Leaders: Gary Skouson (Pacific Northwest National Laboratory)

This BOF is a forum for the administrators, systems programmers and support staff behind large HPC systems that include the new Intel Xeon Phi co-processors to share solutions and approaches to deploying and managing these systems. Subjects for discussion include image deployment, software environments, filesystems and supporting different usage models.

Slurm User Group Meeting

12:15pm-1:15pm
Room: 205/207

Primary Session Leader: Morris Jette (SchedMD)

Secondary Session Leaders: Eric Monchalin (Bull), David Wallace (Cray Inc.)

Slurm is an open source workload manager used 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 13.12, describe the Slurm roadmap, and solicit user feedback. Everyone interested in Slurm use and/or development is encouraged to attend.

The 2013 HPC Challenge Awards

12:15pm-1:15pm
Room: 201/203

Primary Session Leader: Piotr Luszczek (University of Tennessee, Knoxville)

Secondary Session Leaders: Jeremy Kepner (MIT Lincoln Laboratory)

The 2013 HPC Challenge Awards BOF is the 9th edition of an award ceremony that seeks high performance results in broad categories taken from the HPC Challenge benchmark as well as elegance and efficiency of parallel programming and execution environments. The performance results come from the HPCC public database of submitted results that are unveiled at the time of BOF. The competition for the most productive (elegant and efficient) code takes place during the BOF and is judged on the spot with winners revealed at the very end of the BOF. Judging and competing activities are interleaved to save time.

The Message Passing Interface: Version 3.0 and What Comes Next?

12:15pm-1:15pm
Room: 705/707/709/711

Primary Session Leader: Martin Schulz (Lawrence Livermore National Laboratory)

Secondary Session Leader: Richard Graham (Mellanox Technologies)

Last year, the MPI forum released MPI-3 with large enhancements like nonblocking and neighborhood collectives, enhanced RMA, and a tools information interface. However, this milestone is not the end of the standardization process: the forum continues to discuss new additions to the standard, such as support for fault tolerance and improved support for hybrid models. We will use this BOF to discuss the changes in MPI-3 and to continue an active discussion with HPC community on priorities for MPI-4 in an effort to make the forum more approachable for a wider community.

The UDT Forum: A Community for UDT Developers and Users

12:15pm-1:15pm
Room: 703

Primary Session Leader: Allison Heath (University of Chicago)

Secondary Session Leader: Robert Grossman (University of Chicago)

UDT is an open source library supporting high performance data transport. It is used by a growing number of cyberinfrastructure projects and has been commercialized by over 12 companies. UDStar is an application that integrates UDT with common utilities, such as rsync and scp, to make it easier to transport data over high performance networks. In this session, we will provide updates on UDT and UDStar development, a roadmap for the future, and have a discussion with the SC13 attendees from the UDT community of ways that the UDT core developers can support the community of UDT developers and users.

Emerging Technologies

Emerging Technologies is a new element of the Technical Program at SC13. It provides a showcase on the show floor for novel projects at a national and international scale. Emerging Technologies is different from other aspects of the technical program, such as contributed presentations and posters, in that it provides a forum for discussing large-scale, long-term efforts in high performance computing, networking, storage, and analysis, rather than a recent research result that such a project might have achieved. Visit Booth 3547 to check out the latest in recent developments and capabilities.

Emerging Technologies

Emerging Technologies

Monday, November 18
6pm-9pm

Tuesday, November 19-Wednesday, November 20
10am-5:30pm

Thursday, November 121
10am-2pm
Room: Booth 3547

Power Efficient, Scalable Hybrid DSP+ARM+FPGA Platform for Computer Vision Tasks

Jason Bakos (University of South Carolina)

Our objective is to demonstrate the capabilities of a new Digital Signal Processor (DSP), the Texas Instruments Keystone-II, with respect to its ability to perform real-time computer vision tasks on high resolution data. Our platform is comprised of a Keystone-II evaluation module (EVM) that is connected to a Xilinx Zynq 7020 FPGA which acts as an interface to both a USB video camera and an HDMI monitor. We will characterize this platform in terms of its performance, power efficiency, scalability, and applicability to next-generation aerial surveillance tasks. When serving as a coprocessor this DSP architecture potentially offers an equal or better performance-power ratio than state-of-the-art GPUs. General purpose computing using the TI DSP technology is still relatively obscure and its development tools and programming models are still immature. We hope this demonstration will inspire insightful discussion concerning the future of this technology as it relates to high performance computing.

CLARISSE: Cross-Layer Abstractions and Run-time for I/O Software Stack of Extreme-Scale Systems

Florin Isaila (University Carlos III of Madrid)

The current uncoordinated development model of independently applying optimizations at each layer of the I/O software stack on High End Computing platforms will most likely not scale to the expected levels of concurrency, storage hierarchy, and capacity of future Exascale systems. The main goal of CLARISSE project is to investigate a radical new approach of reforming the I/O software stack in order to simplify the development of storage I/O optimizations and facilitate the development of Exascale systems. CLARISSE project will propose novel abstractions and mechanisms for cross-layer monitoring and control of the software I/O stack. More exactly, CLARISSE will explore: the cross-layer dissemination of run-time information throughout the I/O software stack, the cross-layer dynamic shipping of I/O functionality for load balancing, the adaptive control of the I/O data path, and the cross-layer dynamic exploitation of the data locality.

Monte Carlo as an Emerging Technology for High Performance Computing

Michael Mascagni (Florida State University)

Recent developments in Monte Carlo methods (MCMs) for the solution of partial differential equations (PDEs) provide new algorithms that are computationally competitive with traditional deterministic methods for many problems. In fact, there are a variety of new Monte Carlo method alternatives to a variety of numerical computations traditionally solved with deterministic methods. Besides the fact that in certain circumstances these new MCMs are faster than deterministic methods, MCMs have some generic properties that make them especially interesting from the high-performance computing point-of-view.

Thus, MCMs offer a very different computational paradigm for the computation of many numerical problems. They offer many desirable numerical properties, and they are so compute intensive, that a distributed Monte Carlo computation can be undertaken quite efficiently on a very loosely coupled collection of computing resources. MCMs can take advantage of relatively cheap computational resources and can be expected to achieve high levels of parallel efficiency.

The Mont-Blanc Project: Building HPC out of Mobile Phone Processors

Alex Ramirez (Barcelona Supercomputing Center)

High Performance Computing was initially built on vector processors and data level parallelism. During the 1990's, processors targeted at personal computers and workstations integrated the floating point unit, which made them capable of running HPC workloads. These commodity processors were 10 times slower, but they were 50 times cheaper. That fact, coupled with the transition from DLP and shared memory programming to message passing programming models, changed HPC replacing vector systems for more cost-efficient distributed systems built on commodity processors. Smartphone processors now have integrated floating point units. They are still 10 times slower than today's HPC commodity processors, but they are also 50 times cheaper, and they do not require a fundamental change in the parallel programming model to use. We may be about to see another step in HPC evolution, when a new class of super-commodity processor replaces the current established technology for a more cost-efficient alternative.

Automata Processor: A New Accelerator for High Performance Computing

Paul Dlugsoch (Micron Technology, Inc.), Srinivas Aluru (Georgia Institute of Technology), Kevin Skadron (University of Virginia), Michela Becchi (University of Missouri)

For the past seven years, Micron 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. The purpose of this submission is to present and fully describe this technology for the first time, and showcase its applications through talks and demonstrations. Automata processing is a novel approach to parallel processing enabling a new programming paradigm and fresh approach to many standard algorithms. The reconfigurable processor can be programmed with NFAs consisting of thousands of paths which are evaluated simultaneously, or thousands of independent, smaller automata. A new parallel hardware configuration language and other software tools for automata processing will be exhibited. On-going work in the development of automata processor-based algorithms and applications in bioinformatics and other areas will be demonstrated.

Evolving the Power Architecture into Microserver DenseCore Nodes

Kevin Moran (System Fabric Works), Roger Ronald (System Fabric Works), Bill Boas (System Fabric Works)

Our challenges are to provide sufficient compute power for an interesting set of applications and to enable the networking infrastructure for Remote Direct Memory Access (RDMA) protocols over standard interfaces (i.e., 10/40/100Gb Ethernet and/or InfiniBand). A high level of compute power and RDMA support are both critical because neither capability by itself is sufficient. To implement these goals, our project plans to build HPC micro servers using System on a Chip (SoC) network processors. Network processor SoCs have been driven by the need for extremely high I/O rates, fast processing, and low power consumption. Current network processors are capable of switching and routing I/O data rates as fast as 200Gb/sec.

PowerPC is the market leader here. While we believe that ARM (or even Intel) may eventually be a contender in this market, we have selected the Freescale QorIQ family's T4240 as the current strongest candidate for low power exascale exploration in HPC.

HPCMP CREATE-AV Project Exhibit

Robert Meakin (DoD HPCMP), Nathan Hariharan (Chenega Federal Systems, LLC), Scott Morton (University of Dayton Research Institute), Roger Strawn (US Army Aeroflightdynamics Directorate), Joseph Laiosa (NAVAIR), Robert Nichols (USAF AEDC), Greg Roth (USAF LCMC/XR)

The project proposed for exhibit at SC13 is a software development project known as HPCMP CREATE-AV. The project is one of three primary elements of the HPCMP CREATE Program, established in 2008 by the U.S. DoD to enable major improvements in defense acquisition engineering workflows associated with the design and analysis of Ships, Air Vehicles and Radio Frequency Antennas. The HPCMP CREATE-AV Project is tasked to develop, deploy, and support a set of multi-disciplinary, physics-based simulation software products for the engineering workforces supporting air vehicle defense acquisition programs.

The products are designed to leverage the capacity of next generation HPC systems and exploit the use of multi-disciplinary, physics-based testing to drive design iterations prior to the manufacture of physical prototypes; to subsequently aid in design of physical tests and provide engineering data in advance of key design and acquisition decisions prior to production; and to support sustainment processes after deployment.

Application-Aware Traffic Engineering for Wide Area Networks using OpenFlow

Michael Bredel, Artur Barczyk, Harvey Newman (California Institute of Technology)

Scientific collaborations on global scale rely on the presence of high performance networks. The efficiency in data movement translates directly into the capability to reach scientific goals. In order to achieve these goals, we need to increase the efficiency of these networks. We believe that Software Defined Networking and OpenFlow are key factors to success. In this demo we address the challenge of traffic optimization for large flows in wide area layer-2 OpenFlow networks. We present an OpenFlow controller implementation that enables an application aware networking. In particular, we present a simple API that can be used by data movement applications to provide additional flow information to the network controller. Using this information, we use multipath forwarding to leverage network capacity according to application requirements. We show how our network load balancing can improve the network utilization, reduce the transfer times of big data, increase the users quality of experience.

Multi-Scale Reactive Modeling of Insensitive Munitions

William D. Mattson (US Army Research Laboratory)

The Institute for Multi-Scale Reactive Modeling of Insensitive Munitions (MSRM) has the mission to develop a science-based capability to simulate damage to military munitions by capturing the effects that microstructure impose on macroscopic events. This effort will transform the modeling and simulation (M&S) in the DoD for materials design by incorporating essential but currently lacking micro- and meso-level modeling capabilities needed to capture key physiochemical properties in continuum level codes. This project is introducing an agile and robust set of M&S tools applicable for a wide range of material design that are amenable to future growth and expansion. This multi-scale approach provides a true predictive capability of system-level response resulting in 1) faster design and implementation of material solutions 2) reduced risk with integrating new materials. This HPC-intensive multi-scale design and analysis tool suite will be transitioned and disseminated to DoD, industry and university researchers, developers and the acquisition community.

Fault Tolerance Interface

Leonardo A. Bautista-Gomez (Argonne National Laboratory), Franck Cappello (INRIA)

HPC is changing the way scientists make discoveries. Large scale simulations performed in supercomputers are allowing researches of all domains to better understand and study complex natural phenomena. Extreme scale computing promises great opportunities to the entire scientific community, motivating the design and development of always larger supercomputers. However, at extreme scale, component failures start compromising the usability of these systems. Multiple techniques can be used to guarantee the successful completion of the simulation, checkpoint/restart being the most popular of them. This method has been useful for several decades, but it is starting to show some limitations: while the computational power of supercomputers has been increasing exponentially, the I/O system has been growing linearly. This causes a bottleneck while writing large amounts of data to the File System. Fault Tolerance Interface (FTI) is a library that aims to provide researchers easy and scalable multilevel checkpointing.

HPCMP CREATE (TM) Cloud Deploy Project Exhibit

Christopher A. Atwood (HPCMP CREATE)

The project proposed for exhibit at SC13 is a software development project known as the Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP) Portal. An objective of this new software-as-a-service (SaaS) application environment is to enable acquisition engineers and researchers secure access to capabilities without desktop software installation. The web user interface coupled with the HPCMP SaaS environment hosted at HPCMP DoD Supercomputing Resource Centers (DSRC) provide significant capability to DoD engineers and scientists without the need to install or configure software.

The HPCMP CREATE (TM) program is designed to improve the DoD acquisition process (AP) by developing and deploying three sets of advanced computational engineering design tools for use in acquisition programs for the design of military air-vehicles, ships, and RF-antennas.

Preparing for the Next Pandemic by Harnessing the Power of Pervasive Supercomputing

Christopher L. Barrett, Keith Bisset, Madhav Marathe (Virginia Tech)

We propose to showcase a 15 year program on Computational Epidemiology: development and use of computer models for understanding the spatio-temporal diffusion of disease through populations. Controlled experiments used to understand scientific phenomenon are much harder and often impossible to do when studying epidemiology, due to ethical, and often practical reasons. As a result, computational models play an important role in elucidating the space-time dynamics of epidemics. They also serve an important role in evaluating various intervention strategies, including pharmaceutical and non-pharmaceutical interventions. Recent advances in high performance computing, software services and proliferation of the web and handheld devices and sensors have created entirely new opportunities to understand epidemics and devise new ways to control them. By combining interactive games, videos, talks by experts and presentation, we will showcase our program highlighting the use of pervasive high performance computing technologies to support important federal case studies.

Accelerating Improvements in HPC Application I/O Performance and Efficiency

Jim Ryan (Intel Corporation), Paul Grun (Cray Inc.), Bill Lee (Mellanox Technologies)

In most cases, parallel computing I/O comprises a complex dance between the application codes, the compute node based software including MPI and the underlying interconnects. Until now, the common practice in application programming has been to take the I/O interfaces as a given and to code to them, at some cost to the application in terms of performance. The proposed poster will provide application developers/programmers and the providers of open source middleware architectures and interfaces a basis for discussing whether there are classes of applications that can become higher performance, more efficient and use less power by implementing a new framework of I/O APIs and, possibly, protocols. The prospect is API's focused on a specific class of applications may achieve significant improvements in I/O

efficiency with minimal effort. The outcome of the discussions at SC13 and thereafter will provide direction to drive the work of the OpenFabrics Alliance going forward.

Chapel: An Emerging Parallel Programming Language

Bradford L. Chamberlain, Sung-Eun Choi, Martha Dumler, Greg Titus (Cray Inc.)

Chapel is an emerging parallel programming language whose design and development are being led by Cray Inc. in collaboration with partners from academia and computing centers, both domestic and international. Chapel has been designed to improve programmer productivity on large-scale machines—by simplifying the creation of parallel programs; by teasing details related to architectural mapping away from the overall expression of parallel algorithms; and by permitting end-users to specify their own implementations of key parallel features such as array distributions and layouts, parallel iteration strategies, and architectural models. While Cray has spearheaded the Chapel project, the language and its implementation have been designed from the outset with portability in mind, supporting execution on multicore laptops and desktops, commodity clusters, and large-scale HPC systems developed by Cray and other vendors. For the Emerging Technologies track and booth, we intend to provide posters, presentations, and staffing that present an overview of the project.

Scalable Tools for Debugging, Performance Analysis and Performance Visualization

Martin Schulz, Abhinav Bhatele, Peer-Timo Bremer, Todd Gamblin (Lawrence Livermore National Laboratory)

Future HPC systems will exhibit a new level of complexity, in terms of their underlying architectures and their system software. At the same time, the complexity of applications will rise sharply, both to simulate new science and to exploit the new hardware features. Users will expect a new generation of tools that help address these challenges, work seamlessly with current and new programming models, scale to the full size of the machine, and provide automatic analysis capabilities.

We are addressing these requirements through a wide range of tool research and development efforts targeting performance analysis, performance visualization, debugging, and correctness tools, as well as automatic tuning and optimization capabilities. These efforts are supported by activities on tool infrastructures that enable us to work with more modular tool designs and support rapid tool prototyping. This “emerging technologies” booth will showcase our approaches to users, system designers and other tool developers.

Mystic: a Framework for Optimal Performance Modeling

Michael McKerns, Houman Owhadi, Clint Scovel (California Institute of Technology), Tim Sullivan (University of Warwick)

We have developed a new rigorous mathematical framework for quantifying risk designed to utilize all available information and predict the impact of Black Swan events. These rigorous predictors are global optimizations over all possible valid scenarios, and do not rely on common approximations in computing priors that lead to the exclusion of high-impact rare events. Such optimizations, however, are high-dimensional, highly-constrained, and non-convex, and generally impossible to solve with current optimization technology. We will overview of our mathematical framework, and the powerful optimization software we have developed to rigorously solve real-world problems in predictive science, finance, and technology. The software runs in standard python on a simple laptop -- however, it trivially scales up to potentially petascale and larger calculations running on some of the largest computers on the planet. We will also discuss the application of this technology to exascale system performance and design.

Carbon Nanotube Digital Circuits

Max Shulaker, Subhasish Mitra (Stanford University)

Although advances with silicon-based electronics continue to be made, alternative technologies are being explored. Carbon nanotube (CNT)-based digital circuits have the potential to outperform silicon by improving the energy-delay product, a metric of energy efficiency, by more than an order of magnitude, making CNTs an exciting complement to existing semiconductor technologies. However, substantial inherent CNT imperfections have limited past demonstrations to very basic circuit blocks. We show that these imperfections can be overcome, and demonstrate the first computer built entirely using CNT-based transistors. The CNT computer is capable of multitasking: as a demonstration, we perform counting and integer-sorting concurrently. In addition, we implement 20 different instructions from the commercial MIPS instruction set to demonstrate the generality of our CNT computer. It is a considerable advance because CNTs are prominent among a variety of emerging technologies that are being considered for the next generation of highly energy-efficient electronic systems.

Exhibitor Forum/ HPC Impact Showcase

Exhibitor Forum/ HPC Impact Showcase

The Exhibitor Forum offers an opportunity to learn about the latest advances in the supercomputing marketplace. Manufacturers, vendors and HPC service providers will present new products, services and roadmaps. Industry leaders will share their insights into customer needs, market drivers, product strategies and technology trends. Case studies will illustrate how their current and future solutions can be used effectively.

We're particularly excited this year to restructure the sessions to include topics focused on forefront issues such as data, HPC services, usability and exascale. The Exhibitor Forum topic areas this year are:

- Hardware and Architecture
- Moving, Managing and Storing Data
- Software for HPC
- Effective Application of HPC
- HPC Futures and Exascale

Please join us for a few Exhibitor Forum sessions this year. We have a great lineup with 42 talks across 15 sessions.

HPC Impact Showcase

New for SC13, the HPC Impact Showcase reveals real-world HPC applications via presentations in a theater setting on the exhibit floor. The Showcase is designed to introduce attendees to the many ways that HPC is shaping our world through testimonials from companies, large and small, not directly affiliated with an SC13 exhibitor. Their stories relate real-world experience of what it took to embrace HPC to better compete and succeed in their line of endeavor.

Whether you are new to HPC or a long-time professional, you are sure to see something new and exciting in the HPC Impact Showcase. Presentations will be framed for a non-expert audience interested in technology and will discuss how the use of HPC has resulted in design, engineering, or manufacturing innovations (and will not include marketing or sales elements!).

Exhibitor Forum

Tuesday, November 19

Hardware & Architecture I

10:30am-12pm
Room: 501/502

Micron Technology for System Differentiation

Todd Farrell (Micron Technology)

Designing systems that are differentiated from the competition and still meet stringent requirements for performance, energy consumption, and form factor is a daunting challenge. Memory solutions - taken for granted as a commodity element in systems for decades – are now offering game-changing solutions to solve new system challenges. This presentation will introduce three of Micron's newest technology solutions and what they have to offer.

Hybrid Memory Cube: HMC topples the memory wall with high bandwidth, low energy per bit, RAS (dependability) features, and small form factor. Today's commodity DRAM technologies don't compare.

PCIe SSD: Low-latency, high IOPs, reliable, power efficient… all attributes of our new PCIe SSDs. We'll show you how these can dramatically impact your system design.

NVDIMM: Merges DRAM, NAND and an energy efficient power source into one highly reliable system. Perform workloads at DRAM speeds and in the event of a power failure…data is securely preserved.

Implementation Details for Creating a PCI Express-Based Fabric

Larry *Chisvin* (PLX Technology)

There has been significant interest over the past year about the potential for a high-speed fabric created with PCI Express. The clear advantages of lower cost, power, and latency – coupled with high performance and the ability to use the vast, existing infrastructure - make this a compelling solution. The primary features of such a fabric are high speed clustering capability, and the ability to share common storage and communication subsystems among a potentially large number of processing elements.

A key requirement for such a solution is to offer these features, but to do so with existing hardware devices and their software and firmware drivers. This presentation will provide the details of how this will be accomplished through DMA, RDMA, and Shared I/O, and integrate into existing HPC data centers. The performance of the resulting fabric will be quantified.

PTX: The Next Generation Core for HPC NRENs

Suresh C. Roy (Juniper Networks)

PTX Series Packet Transport Routers are high-performance Converged Supercore platforms for HPC NREN (High Performance Computing - National Research and Education/Engineering Network) providers. These routers deliver powerful capabilities based on the Junos Express chipset and forwarding architectures optimized for MPLS and Ethernet, with integrated, coherent 100GbE technology. HPC NREN Providers need innovative network architectures that are dense, cost optimized, highly available, and simplified. The PTX Series dynamic Converged Supercore model with 100 GbE coherent optical technology provides that, and improves total cost of ownership (TCO), optimizes power use, and drastically reduces space utilization.

This session discusses the architecture of the PTX and how it can be leveraged by the next generation of HPC NRENs. It will also touch on recent trials in which the PTX with 100 Gbe coherent optical technology was used to demonstrate the practicality of running Alien Waves over the layer one infrastructure of production HPC NRENs.

Moving, Managing & Storing Data I

10:30am-12pm
Room: 503/504

Implement End-to-End Reliability and Resilience

Henry Newman (Instrumental)

We discuss how to implement end-to-end reliability and resilience into very large storage systems. We start with the basics, including how applications interface to the operating system and local and global namespace semantics and implementation. Then we cover low-level hardware issues, including storage interfaces like SAS and SATA, RAID and parity checking, checksums and ECC, comparing and contrasting different approaches for achieving resilience in current systems. Then we show specific examples from today's large storage systems where low-level bits can change in a way that remains undetected in the system

FhGFS - A Flexible Parallel File System for Performance Critical Applications

Christian Mohrbacher (Fraunhofer ITWM)

FhGFS (http://www.fhgfs.com) is the parallel file system from Fraunhofer's Competence Center for High Performance Computing. Its distributed metadata architecture has been designed to provide the scalability and flexibility that is required to run today's most demanding applications. The file system has been adopted by a wide range of users and to reliably power the storage for systems of all kinds and sizes, ranging from a handful of nodes to current Top500 clusters. The presentation will introduce the filesystem's architecture and give an idea of the flexibility, which makes FhGFS suitable for various scenarios from classic dedicated HPC storage to cloud-style installations with combined storage and compute resources. Besides that, the talk will feature several benchmark results and it will give an overview of the latest developments.

Optimizing Data for Performance, Capacity and TCO for all Storage Tiers - Active or at Rest

Janis Landry-Lane (IBM Corporation)

Worldwide, we now generate the equivalent of all the data that existed in the world up to 2003 every two days. In order to leverage the vast quantities of data being produced daily, and maximize the ability to effectively use the data, it must be optimized for performance, capacity and total cost of ownership (TCO) for all storage tiers – active or at rest.

To set the stage to examine this topic, Janis Landry-Lane, IBM WW Program Director for Higher Education and Life Sciences, will speak about best practices for data access, management, archive and reuse, which are applicable to all research computing disciplines. There are a variety of approaches to the topic. At IBM, our focus is on empowering the users of the data to take ownership of the data without compromising data integrity and accessibility.

Hardware & Architecture II

3:30pm-5pm
Room: 501/502

NEC's Brand-New Vector Supercomputer and HPC Roadmap

Shintaro Momose (NEC Corporation)

NEC presents a brand-new model of vector supercomputer and an HPC roadmap for the future.

NEC has just launched its new generation SX vector supercomputer as the successor model of the SX-9 by aiming at much higher sustained performance particularly for memory-intensive applications. This system is based on the big core strategy targeting higher sustained performance. It provides both the world's highest single core performance of 64GFlops and the world's highest memory bandwidth per core of 64GBytes/s. Four cores, memory controllers, and network controllers are integrated into a single CPU LSI, enabling the CPU performance of 256GFlops and the memory bandwidth of 256GBytes/s. Moreover, in order to boost the performance of memory-intensive applications, each CPU core is designed to exploit the whole memory bandwidth of a CPU (256GBytes/s). This brand-new system is quite beneficial to real scientific and engineering applications. NEC continues to pursue higher effective performance on real applications.

Integrating Intel® Xeon Phi™ Co-Processors Into Existing Cluster Solutions

Michael Hebenstreit (Intel Corporation)

Overview: With the advancement of Intel Xeon Phi Coprocessors integration into existing cluster installations becomes an important aspect of High Performance Computing decisions. This talk highlights various options and demonstrates in examples the advantages of Intel technology.

Topics include: • Standard Cluster models • Review of Xeon Phi programming models • Cluster integration of Xeon Phi as coprocessor • Cluster integration of Xeon Phi as pseudo cluster node • Comparison of user environments between Xeon and Xeon Phi • How to achieve a single system image • Integration of cluster file systems (Lustre, Panasas) • Porting of typical utilities • Integration into job scheduling systems • Integration into cluster management solutions.

Experience level: Intermediate Moving, Managing & Storing Data II

3:30pm-5pm
Room: 503/504

Increasing Storage Server Efficiency and Performance: Is an Edge-Core Storage Architecture Right for You?

Michael Kazar (Avere Systems)

Traditional file-based storage servers are responsible for performing several functions: efficient storage (whether measured in dollars, or watts per gigabyte), quick data delivery and data protection. Diverging from this traditional view of servers performing multiple functions is the contention that a more effective file storage architecture can be built by dividing these functions into different components: primary storage to hold the primary copy of the data, and edge filers holding the most active data in high performance hardware. • Examine pros and cons of both architectures in multi-vendor environments. Topics include migrations between different vendor storage systems and implementation of data management features such as data mirroring. • Compare cost and performance

benefits of the traditional storage server and the new Avere storage architecture. Typical use cases, including storage and users connected via WAN links, where latency concerns impact the selection of a storage architecture and migrating to cloud-based storage.

Building High-IOPS Flash Array with Innodisk FlexiRemap Technology

Charles Tsai (Innodisk Corporation)

With the rapid advancement of processor technologies, disk access has been identified as the next performance bottleneck in many cloud computing applications. In recent years, storage appliances based on flash memory technologies have been deemed as practical solutions to resolving the performance bottleneck. High-end flash appliances are mostly built with proprietary hardware designs, aiming at particular scenarios in larger-scale data centers, and are barely affordable by enterprise and industry customers that also deploy private clouds. Innodisk FlexiRemap technology deals with the challenges of performance, data endurance, and affordability through innovations in software and firmware, creating a new category of flash-collaborating (in contrast to flash-aware or flash-optimized) storage appliances that deliver sustained high IOPS, even for random write operations. The speaker will elaborate on how the Innodisk FlexiArray storage appliance, a high-IOPS flash array built upon Innodisk FlexiRemap technology, provides a cost-effective alternative to customers with demands for high-speed data access.

Analysis of Use Cases of SSDs with Parallel File Systems

James Coomer (DataDirect Networks)

We look at multiple methods of implementing SSDs to benefit aspects of parallel file system performance. SSD's can be deployed with the aim of accelerating metadata operations and small file I/O. However, bottlenecks can exist elsewhere in a parallel file system that prevent the IOPs benefits translating to improved user experience. We test SSD benefits when implemented as simple LUNs for metadata and data and also look at more advanced techniques for deploying Flash acceleration for Exascale file systems.

Wednesday, November 20

Effective Application of HPC I

10:30am-12pm
Room: 503/504

AweSim Platform Offers Web-based Supercomputing Applications

Alan Chalker (Ohio Supercomputer Center), Robert Graybill (Nimbis Services)

The Ohio Supercomputer Center (OSC) and Nimbis Services along with P&G, Intel, TotalSim, AltaSim, and Kinetic Vision have developed the AweSim platform based on an extensive track record of providing web-based modeling and simulation resources to small and medium sized enterprises (SMEs). This effort offers an innovative commercial marketplace consisting of web-based app store, supporting infrastructure and app development tools.

The Nimbis commercial web-based marketplace will be personalized and rebranded for AweSim. OSC brings their experience of creating apps and development tools that will be tightly integrated. Other partners will develop their own new web-based apps.

The presentation will consist of a review of the challenges and lessons learned in working with SMEs that have impacted the architecture of the AweSim Platform. The AweSim marketplace, components and technology will be described. Examples of apps that integrate manufacturing domain expertise, simulation software and cloud-based resources will be provided.

Web-Based Remote Interaction with, and Visualization of, 3D Applications

Marie Granier (SysFera), Olivier Rouchon (National Computer Center for Higher Education)

CINES, one of France's leading large-scale computing centers, is committed to driving the adoption of HPC and visualization technologies. However, in spite of quality equipment and generous access conditions to a large community, usage of HPC and visualization resources is not as high as it should. The main reasons for this are that hardware resources are generally underexposed and hard to use.

To address these issues, SysFera has developed SysFera-DS WebBoard, a web-based portal that can be used to offer remote interaction with and visualization of 3D applications, collaborative sessions, in-advance resource booking, and regular HPC job-submission (including on the visualization resources). These features are made available through most modern web browsers, which greatly reduces administration costs.

The presentation will describe the process of deploying the WebBoard at CINES, and the impact it had over the first 6 months of usage in production. Future requirements and enhancements will be discussed.

HPC Futures & Exascale I

10:30am-12pm
Room: 501/502

Statistic Multiplexed HPC – A Blueprint for Exascale Computing

Justin Y. Shi (Temple University)

The basic challenge for exascale HPC is an application architecture that is capable of delivering incrementally better performance and reliability at the same time as we build larger applications.

We address this challenge by applying the statistic multiplexing principle (like packet switching in data communications) to HPC applications by treating each HPC application as an application-dependent data network where the data objects are statistic multiplexed. We show that SMC-enabled applications have zero checkpoint and rollback overheads. We also show that SMC-enabled application can deliver counter-intuitively better performance than "bare-metal" MPI programs via optimal granularity probing. The project is a collaboration between Temple University and NCAR. It is currently in beta-phase. The source code is available on github.com.

AMD's Dense Computing Solutions for HPC

Michael Schulte (Advanced Micro Devices, Inc.)

During this session, AMD will talk about its move towards dense ARM and Accelerated Processing Unit (APU) based solutions for HPC and share plans for these solutions. We will also discuss the latest developments on the software ecosystem for these solutions including progress being made by the open source ARM community through Linaro and on the Heterogeneous Software Architecture (HSA) through the HSA foundation.

Effective Application of HPC II

1:30pm-3pm
Room: 503/504

GreenButton(tm) Cloud Fabric and the Consumerization of HPC

Robert M. Demb (GreenButton)

The GreenButton™ logo on an HPC application represents the ability for users to have the power of a supercomputer available to them. Users no longer have to care about where their compute resources come from and how they are managed. The Cloud is effectively sedimented into a GreenButton™-enabled application. GreenButton™ Cloud Fabric is a Cloud-Agnostic compute PaaS that implements strong abstractions for general computing and management of compute jobs in the Cloud.

GreenButton™ Job Prediction learning software allows users to pick, for a workload, a time vs. cost point for completing that job. In addition, GreenButton™ software technologies manage all aspects for completing that job.

GreenButton™ Cloud Fabric manages all aspects of billing and governance of those jobs, and thus eliminates "shadow IT," and does this in a seamless, automatic, and transparent way so that users never need be aware of and/or in need of IT staff, administrators, and Dev-Ops.

Architecting Your HPC Strategy with Gompute

Devarajan Subramanian (Gompute Inc)

Learn how large enterprises, ISVs and SMEs can architect their HPC strategy with Gompute. Learn how data center consolidation, HPC cloud offload and remote visualization can be implemented using the Gompute HPC Platform.

HPC and Visualization Tools: Remote Access to Your Simulation

Andrea Rodolico (NICE)

The evolution to remote visualization is the virtualization of the workstation and the creation of special purpose visualization servers that sit alongside the computational server to run CAD, CAE and other 3D applications in your Technical Cloud. A number of burning issues like network overload, data security issues and administration costs are causing many organizations to strive towards this new model. Such a deployment requires support for a mixed Linux and Windows user community on a single server, the efficient sharing of resources, and the flexibility for users to be able to run applications on physical or virtual machines running side by side on the same hardware. DCV and Enginframe, NICE's offering for remote 3D visualization and HPC, tick all these boxes and help customers move from the old model to the new one. By using them, technical staff can really focus on solving difficult engineering challenges rather than IT management issues.

HPC Futures & Exascale II

1:30pm-3pm
Room: 501/502

HPC Efficient Design to Top Green 500

Giovanbattista Mattiussi (Eurotech)

Eurotech would like to talk about their approach in building the HPC efficient systems that topped 1st and 2nd position of Green 500. Obtaining high energy efficiency from standard components is a matter of architectural decisions and design capable of optimizing the contribution of accelerators, liquid cooling and power conversion to maximize efficiency. Eurotech will briefly describe the record systems and their scalability to petascale size.

Energy efficient and green IT are not only a matter of Green 500, but they rely on holistic approaches that aim to minimize the energy and space consumption of a data centers as a whole, while keeping high reliability and low complexity in design, installation and maintenance of HPC systems.

Eurotech will also give an overview of what means being green in project results over and beyond Green 500.

A Practical Approach to Exascale – a New RSC's Ultra High Density Solution with Intel® Xeon Phi™

Alexander Moskovsky (RSC Group)

RSC Group (www.rscgroup.ru/en), the leading developer of direct liquid cooling technologies for HPC and supercomputing systems, will announce a new breakthrough product at SC13: with its new ultra-high density, energy efficient solution with Intel Xeon Phi coprocessors - the exascale systems seem tractable now, due to RSC's platform compactness, scalability and reliability features. The platform design is focused to reach exascale levels of performance can be also successfully utilized for Big Data applications. RSC Tornado architecture, first presented and deployed in 2009, currently delivers over 200 TFlops per small rack (2.6ft x 2.6ft x 42U or 80cm x 80 cm x 42U) is de-facto the densest solutions based on electronic COTS components on the market. RSC has proven its technology in a track record of projects, including Europe's largest Intel Xeon Phi based supercomputer at Russian Academy of Sciences.

Map to Discovery: Data-Intensive Course

William Blake (Cray Inc.)

Research estimates put the volume of data at 40,000 exabytes by 2020. This explosion of data from sensors, data feeds and instruments means science and business will be increasingly data driven. While supercomputing has predominantly been utilized to solve problems based around mathematical models, Cray is seeing a strong trend toward combining data intensive applications with traditional HPC. This presentation will look at the Big Data space and how Cray is building on its success in delivering highly scalable systems to solve the Big Data challenge with integrated computing, storage and large-scale data analytics solutions.

HPC Futures & Exascale III

3:30pm-5pm
Room: 501/502

The World's First Containerized Total Liquid Submersion Cooling Data Centers

Christiaan Best (Green Revolution Cooling)

Green Revolution Cooling introduced a containerized data center solution in mid-2013, a scalable system of dielectric liquid submersion cooling racks installed in a standard 40' ISO shipping container. The first of its kind in the world, the CarnotJet Container, as it is called, represents a new kind of HPC data center – one that provides high performance and high density in an extremely compact and scalable design.

Join Christiaan Best, CEO & Founder of Green Revolution Cooling, as he discusses the broad implications that containerized liquid submersion cooling may have on the future of HPC. In particular, Best will discuss what happens when an HPC data center no longer requires layers of air-cooling infrastructure and how ultimately an HPC center can support world-class capacity while reducing upfront costs and operating costs. Best will also discuss early results from the initial product deployments, one at a Top-10 Supercomputing site.

Tools for Data Analysis and Visualization at Exascale

Utkarsh Ayachit (Kitware Inc.)

Visualization and analysis are critical to understanding complex data and phenomena. However, current approaches will need a paradigm shift in order to scale to the next generation of computers. This session will highlight the Visualization Toolkit (VTK) and ParaView, the leading open-source tools for scientific visualization and analysis. It will also cover the collaborative efforts on the Data Analysis at Extreme (DAX) toolkit and Portable Data-Parallel Visualization and Analysis Library (PISTON), and how to leverage these next-generation paradigms with VTK and ParaView to move your work to the exascale. Further, the session will detail the business benefits of basing a custom solution on open-source tools, such as avoiding vendor lock-in and leveraging a community's software maintenance skills.

Fujitsu HPC Roadmap Beyond the Petascale Computing

Toshiyuki Shimizu (Fujitsu)

Fujitsu offers a range of HPC solutions with the high-end PRIMEHPC FX10 supercomputer, the PRIMERGY x86 clusters and associated ecosystems. For petascale computing, we viewed technologies include a SIMD extension for HPC applications (HPC-ACE), automatic parallelization support for multi-threaded execution (VISIMPACT), a direct network based on a 6D mesh/torus interconnect (Tofu), and hardware support for collective communication (via the Tofu & its dedicated algorithms). The evaluations confirm these features contribute to scalability, efficiency, and usability of application programs in a massively parallel execution environment. In addition, initial indications show that future platforms, which will have more processor cores with wider SIMD capability, will exhibit similar characteristics.

Thursday, November 21

Moving, Managing & Storing Data III

10:30am-12pm
Room: 503/504

Ultra-High Speed, Long Reach 400G Transmission Made Real

Rodney Wilson (Ciena Corporation)

Recently a number of computer communications research breakthroughs and clever techniques have made 400Gb/s a reality. This presentation will discuss these breakthroughs and provide a look at exactly how modifications to modulation techniques and forward error correction result in easily achieving 200G, and how efficient use of optical fiber spectrum allows 400G connectivity. While 400G can be achieved for very short spans, this talk will also explain methods to achieve 400G for wide area spans that meet needs of high-performance computing (HPC) beyond short hop or Metropolitan Network distances.

Attendees of this session will gain a greater understanding about how these cutting-edge transmissions systems work, and how they will greatly assist HPC network engineers in determining how to select and deploy these latest technologies. The session will also provide insights in how to determine which vendors have the depth and portfolio to deliver high-performance networks for the HPC community.

Hybrid Packet-Optical Circuit Switching Helps HPC Data Centers Handle Large Data Flows

Daniel Tardent (CALIENT Technologies)

Large persistent data flows in high-performance computing data centers can drag down performance of the entire data center fabric by keeping other servers from accessing networking resources. This can happen in several circumstances including virtual machine and data migrations, and in the MapReduce function within a Hadoop parallel processing environment.

A new option for offloading these flows is the hybrid packet-circuit switched network, which uses layer-one pure photonic Optical Circuit Switches alongside a packet-switched aggregation network to siphon off these large flows onto high-speed, low-latency fiber optic links. This offloading works with planned data flows and real-time data flows and can be triggered by switch management or by a software-defined networking controller.

This presentation will introduce the audience to hybrid packet-optical network and the role of photonic Optical Circuit Switches in the HPC data center.

Accelerating High Performance and Super Computing with Ethernet

Nadeem Zahid (Extreme Networks)

Ethernet is poised to erode the InfiniBand market in the near future. High Performance Computing cluster architectures have started adopting low-latency, high-speed Ethernet due to its highly favorable performance, efficiency, and scalability. This is happening in the interconnect-fabric while InfinBand still has its hold with the cluster. Trends influencing this shift include: •Ethernet bandwidth's increased demand. The current 10 and 40G demand is making way for 100G in clusters and the need for 400G is anticipated in the future •Shrinking gap between latency and pricing of Ethernet and InfiniBand •Unmatched distance and reach flexibility of Ethernet •Developments in big-fast Ethernet switches that are now aligned with compute and storage trends to create improved synergies Nadeem Zahid, at Extreme Networks, will share insights into how clients in oil & gas, seismology, bio-medical, meteorology and energy sectors can benefit from the enhanced deployment of high-speed and -performance Ethernet networks in HPC environments.

Software for HPC I

10:30am-12pm
Room: 501/502

Announcing OpenMP API Version 4.0

Michael Wong (OpenMP)

The OpenMP® Application Program Interface (API) supports shared-memory parallel programming in C, C++, and Fortran on a wide range of computer systems and operating systems. The OpenMP API provides a portable, scalable programming model that gives shared-memory parallel programmers a

simple and flexible interface for developing parallel applications for a wide variety of platforms. It is jointly defined by a group of major computer hardware vendors, software vendors, and software developers.

The new OpenMP API version 4.0 includes new features such as: Accelerator support for a heterogeneous hardware environment; Enhanced tasking model with groupings, dependencies, and scheduling constraints; Thread affinity to support binding and to improve performance on non-uniform memory architectures; Clean termination of parallel execution through cancellation constructs; User defined reductions (UDRs); Fortran 2003 base language support; Support for array sectioning and shaping in C/C++; Sequentially consistent atomic operations with atomic swaps; and Support for SIMD instruction-level parallelism.

Arista Networks EOS and SDN Solutions

Kulin Shah (Arista Networks)

In the modern high performance data center SDN applications enable greater visibility, allow for agile workload mobility and API based flexible programming of the network devices. Arista EOS supports a broad set of SDN solutions based around OpenFlow, eAPI and VXLAN together with integration into VMWare and Openstack environments. This technical session will present an overview of Arista's SDN applications, the use cases and benefits.

OpenCL 2.0: Unlocking the Power of Your Heterogeneous Platform

Tim Mattson (Intel Corporation)

OpenCL is an industry standard for programming heterogeneous computers built from CPUs, GPUs and other processors. It includes a framework to define the platform as a host (e.g. a CPU) and one or more compute devices (e.g. a GPU) plus a C-based programming language for writing programs that execute on the compute devices. Using OpenCL, a programmer can write parallel programs that use all the resources of a heterogeneous platform. In this talk we will introduce the new OpenCL 2.0 specification. This major update to OpenCL adds shared virtual memory, dynamic parallelism, and much more.

Moving, Managing & Storing Data IV

1:30pm-3pm
Room: 503/504

EXTOLL: A Scalable Interconnect Targeted to FPGAs and ASICs

Ulrich Bruening (University of Heidelberg)

EXTOLL is a new, scalable and high-performance interconnection network targeted for FPGAs and ASICs. The EXTOLL RTL code can be mapped to ASICs and FPGAs from the same code base and thus opens many new application areas where high speed and high flexibility must be combined. It implements a 3-D torus topology designed for low latency and reliable transmission of messages. Cut-through switching allows for low latency beyond nearest neighbor communication. The feature rich network EXTOLL was selected by the EU-Project DEEP to implement a separate network allowing for the virtualization and aggregation of accelerator resources into a so called Booster. EXTOLL's unique feature of a PCIe root-complex allows for autonomous boot and operation of accelerators within the Booster sub-system.

Arista-Designed Networks for HPC Clusters in Oil and Gas Clusters

Wyatt Sullivan (Arista Networks)

Oil and Gas seismic processing is a growing technology that requires massive amounts of compute and storage. As compute and storage scales at Moore's Law, the networking must keep pace and allow for future growth. Arista-designed networks can scale from a single rack lab to 64,000+ 10G nodes at no oversubscription, and anywhere in between. Working with major Oil and Gas compute clusters, Arista has identified and solved many of the major pain points seen in todays large scale datacenters. This presentation will cover the current state of large IP based clusters, current scaling limitations and requirements, best practices using Arista's open networking architecture, designing for future growth.

OrangeFS Drop-In

Walter Ligon (Clemson University), Boyd Wilson (Omnibond), Becky Ligon (Omnibond)

Join our team for a casual drop-in to catch up on another year of exciting progress for OrangeFS, the scalable storage solution for HPC and Big Data. Meet developers, support leaders and a growing user community in this open discussion on current and upcoming releases. As momentum continues to build for this versatile, open source, next-generation parallel file system, the value of building relationships and gaining your insights and individual needs continues to grow. And since it's a drop-in, please feel free to stop by any time during the session.

Software for HPC II

1:30pm-3pm
Room: 501/502

Developing Heterogeneous Applications Confidently with Rogue Wave's TotalView Debugger

Chris Gottbrath (Rogue Wave Software)

Heterogeneous architectures featuring accelerators based on NVIDIA's Tesla, Fermi and Kepler processors or coprocessors based on Intel's Xeon Phi seem to be increasingly common. While programming paradigms like OpenMP, OpenACC and LEO exist to make incremental development a possibility, porting and optimizing existing applications for these heterogeneous systems remains daunting, especially if you are looking to make significant performance improvements. Tools that provide seamless control and visibility of the entire heterogeneous application allow you to troubleshoot problems regardless of what the root cause ends up being.

This talk will provide an update on Rogue Wave's TotalView parallel debugger within the context of heterogeneous architectures. It will also briefly review the benefits of other Rogue Wave technologies, such as the embeddable analytics available in our IMSL Numerical Libraries.

Bright Cluster Manager and Hadoop: From Bare Metal to Big Data in an Hour

Martijn de Vries (Bright Computing, Inc.)

Bright Cluster Manager is uniquely positioned to deliver a robust and comprehensive solution for the ongoing convergence of HPC with Big Data on three fronts. (1) Bright significantly eases the management of Hadoop clusters. Starting from bare metal, and in a manner that is completely compatible with the most-popular Hadoop distributions, Bright provisions, monitors and manages Hadoop clusters. (2) Bright seamlessly integrates the management of Hadoop clusters alongside HPC clusters in a single, unified management context, while preserving distinguishing characteristics of resources (e.g., a Hadoop name or data node) through Hadoop-specific roles. (3) Bright allows IT resources available from the cloud to be seamlessly incorporated for use in HPC, Big Data or hybrid utilization scenarios. In summary, Bright Cluster Manager delivers a threefold advantage to those who seek to make efficient and effective use of their IT resources for HPC and Big Data.

OpenACC Performance Portability

Michael Wolfe (Portland Group)

The promise of high-level directive-based programming APIs, such as OpenMP and OpenACC, is high productivity, incremental migration and portability across a range of platforms. OpenMP, for instance, has been successful in the multiprocessor and multicore arena by allowing high productivity with very good portability across a range of system architectures. OpenACC aims to deliver high productivity, high performance and performance portability across heterogeneous accelerator-enabled platforms, but until now those have only been promises. This presentation will discuss the latest PGI Accelerator compilers and tools and how they deliver performance portability across NVIDIA and AMD accelerators. We will present small and application examples, show delivered performance, introduce the aspects of the compiler and runtime that deliver performance portability, and discuss PGI's plans for other targets in the future.

Moving, Managing & Storing Data V

3:30pm-5pm
Room: 503/504

Integrating HPC Across Global Sites - Scheduling, Data and More

Nick Werstiuk (IBM Corporation)

This session will address how organizations can share and integrate computational and data storage resources on a global basis to improve time to results, maximize resource utilization and enable improved collaboration. IBM Research has developed a distributed caching technology called Active File Management (AFM) which expands the IBM General Parallel File System (GPFS) global namespace across geographical distances – empowering globally distributed teams by providing low-latency access to the same synchronized data, no matter where they are in the world. In addition, IBM Platform Computing workload management tools take maximum advantage of geographically distributed compute infrastructure, ensuring optimal resource utilization and the shortest-possible time-to-results. Together these solutions enable organizations to ensure that infrastructure resources are used to their fullest with no wasted compute cycles and no wasted time waiting to access data, while promoting the sharing of resources and results between teams.

Expanding Mid-Range Lustre Deployment with the ClusterStor 1500

John Fragalla (Xyratex)

Explosive growth among technical computing users in commercial, industrial, scientific and academic communities has created intense demand for parallel data processing and storage combining superior I/O performance with high availability. The third-generation Lustre scale-out storage solution from Xyratex, ClusterStor 1500, delivers a new standard in affordable file system performance, scalability and efficiency with unmatched integration of high availability scale-out storage building blocks. The solution also concentrates storage processing power in a compact footprint suitable for use in customer-provided 19-inch industry standard racks. This unique combination of efficiency and availability characteristics far exceeds limitations associated with legacy entry and mid-range solutions relying solely upon serial data access from a Network File System (NFS). Designed to expand mid-range Lustre file system deployment, ClusterStor 1500 enables commercial and industrial technical computing users to attain faster job completion, higher job iteration rates and more precise application insights with higher resolution accuracy.

HAMMR Enterprise, HPC Principles for Big Data

Brian Heilig (ET International), Nathan Bishop (ET International)

Current Big Data solutions require organizations to store data within a single file system such as HDFS. However, organizations typically store data different systems, each with its own unique storage model. HAMMR Enterprise enables organizations to tie directly into multiple data sources, and analyze data without unnecessary movement or the creation of a unique data set. From its years of HPC experience, ETI is applying its learning's to Big Data. By applying HPC data flow principles HAMMR Enterprise pushes the bar of analytics and allows organizations to broaden their analytics capabilities.

Software for HPC III

3:30pm-5pm
Room: 501/502

Pick Your Battles: Getting Results Faster with Xeon Phi and CUDA

Mark O'Connor (Allinea Software)

New hardware and accelerated computation models can dramatically reduce the compute time needed to reach new results, but do they reduce the overall time, including the research and development needed? Can we hedge training and infrastructure investments in new hardware models with cross-platform performance profiling and debugging tools?

While the hardware manufacturers have been working on reducing the computation time, Allinea have been working with teams across the world to reduce their development, research and training time. In this talk we discuss effective ways to hedge against technological uncertainty and present the productized results of collaboration with research labs across the world.

Creating a Better Infrastructure to Manage Big Data

Trev Harmon (Adaptive Computing)

While traditional HPC is at the core of supercomputing, recent advancements are allowing HPC centers to offer new and exciting compute services to their users and other customers. Adaptive Computing will discuss two such advancements---topology-based scheduling and HPC cloud---and how together they provide an infrastructure that allows HPC centers to provide Big Data and services along with traditional offerings.

Virtualization for Flexible HPC Environments

Shai Fultheim (ScaleMP)

High-end virtualization for HPC environments is an extremely valuable technology which helps to reduce costs and create a more flexible HPC environment. By aggregating industry-standard systems into powerful, large scale SMPs, organizations can quickly react to various workloads. Applications can get full access to terabytes of memory and hundreds of compute cores in a simple and easy to use system. The latest technology is available to users and in addition they can get access to massive quantities of compute cores and memory. In addition to fully supported versions, a free version is available so that demanding users can experience how their applications can scale, without limiting the apps to small amounts of memory.

HPC Impact Showcase

Monday, November 18

Co-Chairs: David Halstead (National Radio Astronomy Observatory), Ray Bair (Argonne National Laboratory)
7pm-9pm
Room: Booth 3947

Introduction to the HPC Impact Showcase

Raymond Bair (Argonne National Laboratory), David Halstead (National Radio Astronomy Observatory)

Overview and highlights from of the exciting presentations to be given by industry leaders throughout the week.

The HPC Initiative at the Council on Competitiveness

Cynthia R. McIntyre, Walter L. Kirchner (Council on Competitiveness)

The 21st century poses new challenges to American competitiveness. Globalization, high-speed communications, enterprise resilience, and energy sustainability issues are forcing organizations at all levels to rethink and redefine how U.S. companies will remain competitive. The Council on Competitiveness is a nonpartisan and nongovernmental organization composed of CEOs, university presidents, and labor leaders working to ensure U.S. prosperity (Compete.org). The council's High Performance Computing Initiative stimulates and facilitates wider usage of HPC across the private sector, with the goal of propelling productivity, innovation, and competitiveness. The initiative includes efforts to foster public-private sector partnerships (e.g., the National Digital Engineering and Manufacturing Consortium, NDEMC.org) in order to better leverage resources and expertise and help overcome barriers to more widespread private sector usage.

This presentation will highlight the council's efforts to stimulate use of HPC in manufacturing. The presentation will include problems the council members faced and the approaches they developed in partnership with companies, government-funded facilities, and programs across the country. The result has been a number of success stories, where these organizations advanced their R&D, accelerated innovation, created important new knowledge, and shortened the time to market for new products—all essential to business success in the face of global competition.

Tuesday, November 19

Co-Chairs: David Halstead (National Radio Astronomy Observatory), Ray Bair (Argonne National Laboratory)
10:20am-2:30pm
Room: Booth 3947

HPC and Racecar Design

Donour Sizemore (Michael Waltrip Racing)

In professional motorsports, performance is measured by small margins – fractions of a second and thousandths of an inch. Close competition leads to tight regulation, which in turn escalates costs. Significant resources are needed to gain small advantages.

Aerodynamics is crucial to vehicle performance. Unfortunately, field testing aerodynamic systems is extraordinarily expensive. Most automotive engineers now use computational fluid dynamics to develop car bodies and cooling systems. In this talk, I discuss how Michael Waltrip racing uses high-performance computing to accelerate our design/test cycle. Topics include application development, resource acquisition, and results validation.

Michael Waltrip Racing is a professional stock car racing organization in Cornelius, NC. The company designs, builds, tests, and fields three teams in the NASCAR Sprint Cup Series.

Advanced Modeling of the Human Skin Barrier Using High-Performance Computing

Russell DeVane (Proctor & Gamble)

The stratum corneum (SC) forms the outermost layer of mammalian skin, providing a defense against external assault while also preventing dehydration. The 15 µm-thick structure is composed of protein rich corneocytes, embedded in a multilamellar lipid matrix. Although extensive experimental work has produced a somewhat detailed picture of the SC lipid matrix, molecular-level details are still lacking. Two key pieces missing are (1) the lack of knowledge regarding the molecular arrangement of the SC lipids and (2) a detailed mechanistic understanding of skin barrier disruption. These missing pieces, as well as others, restrict efforts to quantitatively model skin especially with respect to skin penetration. Through a Department of Energy INCITE award, we are using the Oak Ridge Leadership Computing Facility system Titan to run large-scale molecular dynamics simulations in order to shed light on these questions regarding the SC lipid matrix. Guided by the most reliable experimental data available, we are investigating ensembles of coupled subsystems of SC lipids; the SC lipid matrix is a complex superstructure that cannot be modeled by independent simulations of a simplified model. Our massively parallel approach exploits leadership computing resources, while using fully open source software. At this point we not only have gained insight into the mechanisms of skin lipid disruption but also have acquired a better understanding of the roles of specific lipid components in the SC lipid matrix.

HPC – A Great Enabler in Advancement of Chemical Technology and Science

Dee Dickerson (Dow Chemical Company)

Early in the 1990s, The Dow Chemical Company (TDCC) ventured into the arena of computational fluid dynamics simulations using a supercomputer at NCSA. The purpose was to understand fluid migration and absorption by the superabsorbent (Drytech® particles) in baby diapers. Since then Dow has been performing computational simulations not only to solve problems but also to advance discoveries in chemistry.

Today Dow uses over 4,000 CPU cores for a multitude of research projects, manufacturing process design, problem solving, and process optimization.

HPC has also enabled better materials design and models for plant troubleshooting. Recently Dow Agro Sciences (DAS) used HPC to identify and validate millions of direct numerical simulation (DNS) variants from thousands of plants. Such discovery is an important piece of the DAS pipeline for new variety development.

HPC enabled Dow researchers to meet stringent timelines and deliver implementable solutions to businesses. HPC provides the platform where large, complex, reacting flow models can be analyzed in parallel to significantly shorten the delivery time. Dow's science and engineering community advances technologies and develops cutting-edge computational capabilities involving complex multiphase reactive processes in dispersion applications, possible only with the advancement of high-speed, high-capacity, and large-memory HPC systems.

This presentation will demonstrate the wide application of HPC across disciplines and businesses in Dow. Many invention disclosures, priority applications, patents, external papers, and internal research reports have been directly based on the results of computational simulations. HPC is definitely an integral part of the scientific process at TDCC.

HPC, A Journey You Must Take

Mohamad S. El-Zein (John Deere)

Computing is a journey where the speed of the vehicle has increased dramatically within a short period of time. The faster the vehicle, the shorter the product cycle time provided the training, the software, the personnel, and the financial resources are available. There are many surprising destinations and false advertising that need to be sorted out; plus many set legacy systems that companies are dependent on which are operating at low cost structure, speeds.

For many enterprises, the journey starts by the pressures of the market place and the need to harvest the fruit of new technology and innovation. Along the way, the hurdles become overwhelming whether it is the infrastructure, the software cost or the trained resources. Surprisingly, the main drawback to HPC is not the compute power or the hardware cost but the pre-historic cost models of the "engineering software".

This presentation will outline the computing journey over the years and where HPC is/will be used.

Accelerating Innovation While Reducing Risk – High Impact HPC at a Formula One Team and a Racing Bike Wheels Manufacturer

Steve M. Legensky (Intelligent Light)

For performance-oriented engineering organizations, high-performance computing (HPC) offers the promise of accelerating performance while reducing risk. The HPC-enabled success of two such organizations will be compared: one with large in-house HPC resources and dedicated computational fluid dynamics (CFD) engineering staff and the other a small engineering and manufacturing company producing consumer products. Common needs and different HPC strategies are discussed. These organizations depend on working with commercial software vendors to develop their capabilities and direct their investments to core skills and knowledge.

Red Bull Racing dominates Formula 1 racing with three consecutive Constructors and Drivers Championships, while Zipp Speed Weaponry has developed leading-edge products that resulted in a doubling of category revenue and 120 new manufacturing jobs for a market leader in a mature segment. The audience will see how different investment and HPC use models are producing extraordinary results and are illustrative of the path forward for innovators in many industries of any size. Through the use of HPC for CFD, these companies rapidly develop and deliver accurate, actionable information to engineers and managers through engineered CFD workflows. Each company has developed the ability to rapidly explore many ideas, put more mature designs into testing and production, and reduce the risks of exploring novel design approaches.

Next Generation Sequencing: HPC Challenges and Cloud Based Solutions

Jason Tetrault (AstraZeneca)

Pharma companies are faced with unprecedented amounts of genetics data coming out of the public domain that they need to combine with their own sequencing datasets in order to gain better understanding of disease, its translation across species, and corresponding stratification of a given specie population. Most of today's public datasets are stored in the cloud, and it is almost impossible to bring all the data

internally for analysis. Thus, new hybrid HPC solutions need to be implemented. This new cloud reality brings an interesting challenge to the risk-averse regulated pharmaceutical industry. The question is not "if" but rather "how" we run computing in the cloud and still protect patients' and donors' privacy (e.g., HIPAA, EU vs US legislations). This presentation will discuss "how" we got there, describing the challenges we faced, decisions we took, and solutions we built along the way to support our next-generation sequencing efforts.

Wednesday, November 20

Co-Chairs: David Halstead (National Radio Astronomy Observatory), Ray Bair (Argonne National Laboratory)
10:20am-2:30pm
Room: Booth 3947

High Performance Computing in Pharmaceutical R&D - Transforming Data Into Therapies

Guna Rajagopal (Janssen Research & Development, LLC)

New experimental modalities incorporating technological advances in miniaturization, parallelization, and especially robotics have opened up new ways to study biological systems

in their normal and diseased states and has catalyzed transformational changes in how we study disease and search for cures. The knowledge gained will enable the development of predictive models that dramatically increase the specificity and timeliness of diagnosis for a wide range of common diseases. These quantitative, data-based models could also guide the search for new therapies by delivering accurate predictions about the efficacy of treatment options, thus realizing the era of personalized medicine. Because of the enormous scale of the data in this new era of ubiquitous sequencing/imaging and concomitant developments of other data- intensive 'omics and imaging technologies, the computational challenge of making sense of massive, heterogeneous data sets to guide our R&D programs throughout Janssen Pharma is formidable indeed. Our goal at Janssen is to develop and deploy state-of-the-art Informatics capabilities and computational expertise in high-performance computing in order to extract information and knowledge from these "big data" sets. This focus is guiding our ongoing efforts, illustrated by a few examples, to support researchers throughout our R&D enterprise in helping them bring products to market – as the patients are waiting!

Bicycle Aerodynamics: Fast Paced Research and Innovation Using Cloud HPC

Mio Suzuki (Trek Bicycles)

In the ever-accelerating pace of the bicycle industry product development cycle, use of cloud-based high-performance computing (HPC) and macro-based automation has been playing a critical role in pushing the boundaries of Trek's engineering R&D. This presentation will highlight the benefits that HPC-based computational fluid dynamics (CFD) has brought to Trek. Examples will be used to showcase the industry-leading design process.

Trek has been successfully implementing an HPC-based CFD program, where the production-related simulation solver elapsed time was reduced as much as 90%. In addition to the speed, HPC has enabled the engineers to change the course of design method. More intelligent model choice is now possible with fast DOE and optimization. Additional computation power has also brought high-fidelity simulation using detached eddy simulation (DES) to further enhance the understanding of bicycle aerodynamics. URANS and overset simulations have also been used to uncover aerodynamics insights that may be unique to bicycle testing.

The successful integration of the on-demand cloud HPC and local process automation leads to significant increase in productivity, specifically in analysis turnaround time reduction and rapid design iterations. This cost-effective, high-output approach will continue to drive the future of innovation culture at Trek.

HPC: Where the Rubber Meets the 'Simulated' Road

Christopher M. Baker (Goodyear Tire & Rubber Company)

Historically, new tires were developed by build-and-test cycles. Tires were beyond the capability of simulation technology due to the extreme loads, high speeds, large deformations, and soft/stiff materials. Goodyear has invested over 25 years in High Performance Computing They have collaborated with national laboratories developing highly paralleled modeling codes to simulate structural, fluid, noise and thermal analyses under all types of loads and road conditions. Goodyear has moved from punch cards and single CPU mainframe computers, sharing cycles, to expensive 'big iron' super computers, and on to thousands of low cost Linux compute nodes, multi-core processors, high speed interconnects, parallel storage and multi-site clustering that spans the globe. The switch to commodity hardware, Linux operation systems and open source software has propelled Goodyear's tire development and shortened its time-to-market allowing Goodyear to grow even through economic downturns. HPC has allowed Goodyear to replace most of its physical build-and-testing with proven simulations that are a fraction of the cost and yield data not possible from driving tires around a race track for months. HPC enables Goodyear tire designers to model multiple iterations of a tire, varying materials, tread patterns, weight, thicknesses and sizes until the customer's needs are met.

High Performance Computing at the Bosch Research and Technology Center for Industrial Applications

Aleksandar Kojic (Robert Bosch Research and Technology Center)

Researchers at the Bosch Research and Technology Center use high-performance computing for numerous projects with direct industrial relevance. This talk will provide examples of projects in energy technology from among the following: ab initio simulation of battery and thermoelectric materials; multiscale simulation of battery performance and safety; and the application of large-eddy simulations to turbulent combustion, including injection, combustion, and cavitation. Additional examples from the data-mining team will also be presented. This team researches and develops machine learning algorithms that process, iteratively, data sizes that are substantially larger than typical memory sizes for application to areas such as health informatics, manufacturing quality, and predictive maintenance.

Jeco Plastic Products – HPC Pallet Design

Paul M. Thompson (Jeco Plastic Products, LLC.)

Jeco Plastic Products designs and manufactures durable pallets and containers for the worldwide printing and automotive industries, as well as other industries handling heavy product or delicate items that must be protected from damage. A recent collaboration with Purdue University and the National Digital Engineering and Manufacturing Consortium provided us with access to powerful computers and software, which were used to perform stress calculations and demonstrate to customers that plastic pallets can perform as effectively as metal pallets in many situations. With HPC and finite element analysis, Jeco was able to significantly strengthen an existing pallet design by making structural changes as a result of detailed 3D models. Jeco plastic pallets are stronger and last significantly longer than wooden pallets. They are much lighter in weight than metal pallets and can be designed to hold much larger loads than would seem possible. The calculations we now can do enable us to develop plastic structures for a wide variety of applications.

Extreme Data for Business Situational Awareness Demands High Performance Computing (HPC)

Earl J. Dodd (LGT Advanced Technology Limited)

Juxtaposing several trends—growth of big data, the fact that most data centers use as much "overhead" energy as they do to power their computing resources, inadequate data center infrastructure, and the move to the smart grid—a new agile, scalable, and sustainable energy technology market has formed.

LGT Advanced Technology seeks to address this new market with its sustainable energy system (SES). The SES represents the first renewable energy system that allows all forms of energy to cogenerate while storing energy for on-demand use. At its core is a fully decoupled wind turbine and tower complex using hydraulic transmission. Its energy storage system feeds ground-based generators that are uniquely clustered, delivering high-voltage direct current to specialized energy consumers in nearly all weather conditions.

Our global optimization processes will couple an HPC system and command operation center, dynamically optimizing an entire energy farm based on real-time operating and financial metrics.

LGT use of HPC for breakthrough innovation of tough engineering problems: Modeling and simulations of systems-of-systems; Computational fluid dynamics modeling to predict turbine performance; Continuous monitoring, reporting, and autonomic functions HPC-supported nontraditional aspects of the SES; Acoustic modeling for bat avoidance; Reductions in infrastructure size and weight.

Just as the Internet transformed information and knowledge, clean and sustainable power producing technology—efficient turbines, energy storage, and agile cogeneration—will have similar global impacts. Renewable energy is universally recognized as a necessary path on the route to global environmental and economic recovery. This recovery is impossible without HPC.

Thursday, November 21

Co-Chairs: David Halstead (National Radio Astronomy Observatory), Ray Bair (Argonne National Laboratory)
10:20am-1pm
Room: Booth 3947

When "Big Iron" Isn't Big Enough

Matthew H. Evans (Lockheed Martin Corporation)

Lockheed Martin relies on high-performance computing (HPC) to deliver advanced technology solutions, from undersea vehicles to aircraft to satellites. High-fidelity, physics-based models and increasingly powerful machines provide critical input to design and test efforts. Designing, building, and sustaining integrated systems require more than just powerful, stand-alone modeling, however. Understanding and optimizing the performance of interconnected system-of-systems depend on linking physics-based models across computing systems. For example, simulating a regional, integrated air and missile defense system requires simultaneously running interdependent models in several HPC environments. From basic components to full-scale systems, high-performance computing at Lockheed Martin relies on powerful machines and powerful networks.

High-Performance Computing at BP

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.

Our team will complete construction of a new computing facility and deliver a 2.2-petaflop system this year. 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.

Presentation in Preparation for Assessing the impact of HPC on Wall Street

Bradford Spiers (Bank of America Corporation)

Abstract: To Be Announced

HPC for High Patient Throughput Magnetic Resonance Imaging

Steven Cauley (A.A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital)

The increased importance of health care cost containment has led to the development of techniques for achieving high patient throughput in magnetic resonance imaging (MRI). These clinically relevant techniques, often a product of collaborative industrial and academic research efforts, reduce cost by minimizing patient imaging time. Sophisticated image reconstruction algorithms are typically required to ensure clinically relevant image quality. High-performance computing (HPC) resources are then utilized to facilitate real-time imaging for radiologists and neuroscience researchers.

Through partnership with Siemens Medical Solutions we illustrate successful collaboration for the development, integration, and distribution of the Simultaneous Multi-Slice (SMS) imaging technique. SMS allows for information from multiple regions of the human body to be encoded concurrently. The expertise of our industry collaborator was leveraged for the formulation of appropriate image-quality metrics. Based on these metrics, penalty-driven optimization models were formulated. Utilizing continuous feedback from Siemens, we successfully embedded our SMS reconstruction algorithm into several MRI platforms. Software packages were jointly released for testing at major hospital and research centers around the world. Work is on-going to develop SMS into a real-time tool through utilization of the multi-CPU multi-GPU HPC environment available on the newest Siemens scanners.

At this time we are involved in a new joint collaboration with Siemens for the development of real-time compressed sensing (CS) algorithms. CS is a general technique for producing images using rapidly acquired data. Here, we will focus on the balance between algorithm simplicity, algorithm efficiency, and the cost-benefit of HPC for MRI applications.

HPC Interconnections

HPC brings people together, ideas together, communities together. And it is through the power of connections so evident at the Supercomputing Conference that some of HPC's greatest contributions to life today got their first exposure and traction. Building on the power of community, SC13 will focus on the increasing integration of HPC into modern life through HPC Interconnections – an enhanced version of the long-standing community elements of SC. HPC Interconnections will help everyone get more out of the conference, providing programs for everyone interested in building a stronger HPC community, including students, educators, researchers, international attendees and underrepresented groups.

HPC Interconnections provides a welcoming entry to SC13 for attendees who may be new to the community or the conference. A range of programs will help attendees get the most out of their time at the conference and bring them into the HPC community. HPC Interconnections comprises Broader Engagement, HPC Educators, Mentors & Protégés, a program for international attendees, and programs for students, including the Doctoral Showcase, the Student Job Fair, the Student Cluster Competition and Student Volunteers. The goal is to ensure that all attendees make meaningful connections to help them during SC13 and after. Please join us.

HPC Interconnections

Monday, November 18

SC First-Timers Introduction

Chair: Bruce Loftis (University of Tennessee, Knoxville)
6:00pm-6:30pm
Room: 705/707/709/711

Participants who are attending the SC conference for the first time are invited to this 30-minute overview. Presenters will briefly describe the various parts of the conference, make some suggestions about what events you should attend and answer your questions. The session will end before the Opening Gala Exhibition at 7 p.m. Participants may suggest questions at sc-first-time@info.supercomputing.org. These will be answered at the Monday session.

Wednesday, November 20

Student Job/Opportunity Fair

Chair: Bernd Mohr (Juelich Supercomputing Center)
10:00am-3:00pm
Room: 605/607

This face-to-face event will be open to all students and postdocs participating in the SC13 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. New this year is the opening of the fair to postdocs, in addition to students.

Thanks to the following supporters of the HPC Interconnections Program:

- **Bank of America**
- **Proctor & Gamble**
- **Disney Research**
- **United Technologies Research**
- **National Security Agency**
- **National Science Foundation/Beyond Broader Engagement**
- **Lawrence Livermore National Laboratory**

Also thanks to those who contributed to the Students Scholarship Fund during registration.

Broader Engagement/HPC Educators

Saturday, November 16

LittleFe Setup

Chair: Charles Peck (Earlham College)
12pm-4:30pm
Room: 702/704/706

Charlie Peck (Earlham College), Aaron Weeden (Shodor Education Foundation), Mobeen Ludin (Shodor Education Foundation), Skylar Thompson (University of Washington), Kristin Muterspaw (Earlham College), Elena Sergienko (Earlham College), Tom Murphy (Contra Costa College)

LittleFe is a portable mini-cluster small enough to fit in a shipping case, light enough to easily move between classrooms that travels to conferences and workshops as standard checked baggage. The 2013 buildout will feature v4a LittleFe units. This design utilizes dual-core Intel® Atom™ processors with NVIDIA ION2 chipsets which support CUDA/OpenCL programming. The buildout will consist of participants assembling their LittleFe unit from a kit; installing the Bootable Cluster CD (BCCD) software on it; learning about the curriculum modules available for teaching parallel programming, HPC and CDESE; and learning how to develop new curriculum modules for the LittleFe/BCCD platform.

Broader Engagement Arrival and Registration

Chair: Damian Rouson (Stanford University)
1:30pm-5pm
Room: 705/707/709/711

Broader Engagement participants should plan to arrive at the Colorado Convention Center before 5pm Saturday, November 16, and check in at the SC13 conference registration area to pick up their badges and other conference materials. Conference registration is open on Saturday from 1pm-6pm. After you have gotten your badge, please attend the Broader Engagement orientation at 5pm, the Joint Broader Engagement/HPC Educators orientation at 6pm, and the Broader Engagement/HPC Educators networking event from 7pm-9pm.

Broader Engagement Orientation

Chair: Mary Ann Leung (e_SciTek)
5pm-6pm
Room: 705/707/709/711

Damian Rouson (Stanford University), Mary Ann Leung (e-SciTek), Raquell Holmes (improvscience)

This program begins with a welcome from Broader Engagement Chair Damian Rouson and Broader Engagement Deputy Chair Mary Ann Leung. An overview of the program activities will be presented and important topics guiding the week's activities will be discussed. The session will also include an

interactive session led by Raquell Holmes from improvscience focused on improving your ability to collaborate, communicate and work with each other to fully take advantage of the conference and each other's expertise.

Broader Engagement/HPC Educators Orientation

Chair: Laura McGinnis (Pittsburgh Supercomputing Center)
6pm-7pm
Room: 705/707/709/711

William Gropp (University of Illinois at Urbana-Champaign), Elizabeth Jessup (University of Colorado Boulder), Damian Rouson (Sourcery, Inc), Almadena Chtchelkanova (National Science Foundation)

This program will begin the SC13 Broader Engagement and HPC Educators programs with a welcome from SC13 General Chair William Gropp; HPC Interconnections Chair Elizabeth Jessup; Broader Engagement Chair Damian Rouson; and HPC Educators Chair Almadena Chtchelkanova. The session will also include an overview of the SC13 conference programs and areas.

Broader Engagement/HPC Educators Networking Event

Chair: Elizabeth Bautista (Lawrence Berkeley National Laboratory)
7pm-9pm
Room: B Lobby

This event facilitates closer interaction and exchange of ideas around posted content on high performance computing. Broader Engagement and HPC Educators program participants are invited to present technical posted resources that can serve as a backdrop for discussion and networking among the participants and other attendees.

Sunday, November 17

Broader Engagement/HPC Educators Plenary I

Chair: Damian Rouson (Stanford University)
8:30am-10am
Room: 705/707/709/711

Making Parallelism Easy: A 25-Year Odyssey

Kunle Olukotun (Stanford University)

In this talk, I will trace my involvement with high-performance computing and parallel programming over the last 25 years. I will relate my experience with the challenges of developing computer architectures and system software to make parallelism easy to use for application software developers. I will touch on message passing machines, the genesis of chip-multiprocessors (CMP), speculative parallelism, commercial throughput processor design, and my most recent work in domain specific programming languages for parallelism.

Fortran 2008 Coarrays and Performance Analysis

Chair: Fernanda Foertter (Oak Ridge National Laboratory)
10:30am-12pm
Room: 705/707/709/711

Platform-Agnostic Multiphase Flow Simulation via Fortran 2008 Coarrays

Karla Morris (Sandia National Laboratories)

Computational modeling of turbulent combustion is vital for our energy infrastructure and offers the means to develop, test and optimize fuels and engine configurations. In the case of internal combustion engines, fuel injection simulations provide insight into phenomena that determine engine efficiency. In modeling the dilute spray regime, away from the injection site and downstream of the atomization processes, considerable doubts persist regarding how to best parameterize and predict the various couplings between the spray and the surrounding fluid turbulence. This talk will focus on the computational issues addressed by Fortran 2008 coarrays. Coarrays provide a partitioned global address space (PGAS), high-level view of the hardware without reference to the underlying communication layer, which compiler teams are free to implement via any of several open-source or proprietary communication protocols. This talk will demonstrate the use of coarrays in a spray modeling package.

Performance Evaluation Using the TAU Performance System

Sameer Shende (University of Oregon)

To meet the needs of computational scientists to evaluate and improve the performance of their parallel, scientific applications, we present the TAU Performance System. This talk will describe a versatile tool for performance profiling, tracing, and debugging I/O and memory issues. This talk will describe some interesting problems in performance evaluation of HPC applications. The talk will introduce profiling and debugging support in TAU. TAU now includes support for tracking call-stacks at the point of program failure to isolate runtime faults. The talk will briefly cover performance evaluation of parallel programs written in Python, Fortran, C++, C, and UPC, using MPI, and other runtime layers such as CUDA, C++OpenCL, SHMEM, and OpenMP using the TAU Performance System. We will demonstrate the different techniques for program instrumentation. The talk will include support for memory and I/O evaluation.

Panel on Fortran 2008 Coarrays and Performance Analysis

Fernanda Foeretter (Oak Ridge National Laboratory), Karla Morris, Sameer Shende (Sandia National Laboratories)

The panel discussion will focus on technical as well as career and professional development topics.

Supercomputing in Plain English

Chair: Steven Brandt (Louisiana State University)
10:30am-12pm
Room: 708/710/712

Henry J. Neeman (University of Oklahoma), Charles Peck (Earlham College)

This session provides a broad overview of HPC. Topics include what is supercomputing?; the fundamental issues of supercomputing (storage hierarchy, parallelism); hardware primer; introduction to the storage hierarchy; introduction to parallelism via an analogy (multiple people working on a jigsaw puzzle); Moore's Law; the motivation for using HPC. Prerequisite: basic computer literacy.

LittleFe Buildout

Chair: Charles Peck (Earlham College)
10:30am-5:00pm
Room: 702/704/706

Charlie Peck (Earlham College), Aaron Weeden (Shodor Education Foundation), Mobeen Ludin (Shodor Education Foundation), Skylar Thompson (University of Washington), Kristin Muterspaw (Earlham College), Elena Sergienko (Earlham College), Tom Murphy (Contra Costa College)

LittleFe is an educational appliance designed to make teaching parallel and distributed computing easy to do and inexpensive. The SC13 buildout will feature LittleFe units with dual-core Intel® Atom processors and NVIDIA ION2 chipsets which support CUDA/OpenCL programming. On Sunday participants will assemble and test the hardware for their LittleFe units.

Curriculum Workshop: Mapping CS2013 and NSF/TCPP Parallel and Distributed Computing Recommendations and Resources to Courses

Chair: Thomas Hacker (Purdue University)
1:30pm-5pm
Room: 708/710/712

Joel C. Adams (Calvin College), Richard A. Brown (St. Olaf College), Elizabeth Shoop (Macalester College)

The ACM/IEEE CS Curriculum 2013 report is expected to be available by the time of the SC13 conference. The IEEE/TCPP Curriculum Initiative on Parallel and Distributed Computing (PDC) is also available. In both of these documents, PDC moves from being an elective into the core CS curriculum, meaning all CS students now need to learn about it. In this workshop, we will (i) provide an overview of the PDC topics that should now be in every CS department's core curriculum, (ii) explore different curricular approaches to covering these topics, (iii) present existing pedagogical and computational resources to aid CS educators in adding coverage of PDC to the core CS curriculum, and (iv) provide time for participants to brainstorm, discuss, and plan how they will incorporate these PDC topics into the core CS curriculum at their home institutions.

Hands-On Exercise: Performance Analysis with TAU

Chair: Damian Rouson (Stanford University)
1:30pm-3pm
Room: 705/707/709/711

Karla Morris (Sandia National Laboratories)

The hands-on portion of the tutorial will guide the developers through the steps involved in using the TAU Performance System (http://tau.uoregon.edu) to study a sample Fortran 2008 coarray parallel program. Participants will use a HPC Linux LiveDVD that will allow them to boot their laptops to a Linux distribution with the requisite tools installed. Participants are encouraged to bring a laptop with them with the VirtualBox virtualization software and related OVA files installed from http://www.hpclinux.org/. Participants will also receive temporary guest accounts on Oak Ridge National Laboratory's Titan supercomputer.

Graphics and Visualization Technologies

Chair: Ritu Arora (Texas Advanced Computing Center)
3:30pm-5pm
Room: 705/707/709/711

Visualizing in the Information Age — Gaining Insight Against Insurmountable Odds

Kelly Gaither (University of Texas at Austin)

We are living in an age in which digital information makes up 90% of all information produced. In this digital age, we are struggling to manage the deluge of data, much less derive meaning and insight from it. Visualization has become an essential tool in scientific and information research, spanning many disciplines. As complexity and data size have scaled dramatically, so have the requirements for processing, managing, and storing the enormous amounts of data needed for visualization and analysis. In this talk, I will discuss why visualization provides such a powerful medium for analyzing, understanding and communicating the vast amounts of data being generated every day. I will show examples of visualization over the years and will discuss what the future will hold with so much of our everyday lives being permeated by digital data/information.

Using Supercomputers to Create Magic: HPC in Animation

Rasmus Tamstorf (Walt Disney Animation Studios and Disney Research)

Creating the magic in an animated feature film from Disney today requires many millions of CPU hours. Historically the clusters used to perform these computations have ranked highly on Top500, and these systems share many characteristics with traditional HPC systems. Yet, animation provides unique challenges due to the diversity of workloads. In this talk I will first give an overview of the many compute challenges in an animation production system. Then I'll look in more detail at a physically based simulation system for computing soft tissue deformation. By using a number of techniques from HPC, this system is able to provide near realtime feedback to artists. Providing such fast feedback can be transformative for the creative process but requires strong scaling of the underlying algorithm and system. Many other problems can potentially benefit from such a transformation which makes animation and 3D graphics an exciting domain for developing HPC algorithms.

Opportunities in HPC and Graphics: Perspectives of a Current Student

Michael Driscoll (University of California, Berkeley)

3D computer graphics abounds with problems for high-performance computing research. Many graphics applications strive to provide high-fidelity experiences to artists and audiences, and HPC techniques are one means of achieving that goal. In this talk, I'll focus on two graphics kernels that can be implemented efficiently by reducing them to common patterns found in scientific computing. The first kernel computes the location of surfaces for complex shapes commonly found in animated feature films. I'll show how it can be implemented with sparse matrix-vector multiplication to yield twice the performance of existing approaches. The second kernel computes the location of a character's hair as he or she moves through time, i.e. hair simulation. I'll describe a new, communication-avoiding algorithm for particle systems that promises good performance for hair simulations on distributed-memory machines. I'll conclude with a few more problems in graphics that can benefit greatly from high-performance implementations.

Monday, November 18

Broader Engagement/HPC Educators Plenary Talk II

Chair: Almadena Chtchelkanova (National Science Foundation)
8:30am-10am
Room: 705/707/709/711

Perspectives on Broadening Engagement and Educations in the context of Advanced Computing

Irene Qualters (NSF/CISE Advanced Computing Infrastructure Division)

As the role of computing and computational science has become integral to the advance of science and engineering, the associated cyberinfrastructure to support HPC within this framework has moved from a more isolated activity to one that is increasingly integrated to the larger cyberinfrastructure. In parallel, an increasing need for multidisciplinary research presents challenges both educational and workforce development. These trends offer unique opportunities for broadening engagement within the HPC community.

HPC Applications

Chair: Sadaf R. Alam (Swiss National Supercomputing Centre)
10:30am-12pm
Room: 705/707/709/711

Exploring the Structure of the Nucleus Using GPUs

M. A. Clark (NVIDIA Corporation)

The exponential growth of floating-point power in GPUs, combined with high memory bandwidth, has given rise to an attractive platform upon which to deploy HPC applications. When it comes to legacy applications there is a danger that entire codebases have to be rewritten to fully embrace this computational power. In this session we discuss how to efficiently port legacy lattice quantum-chromodynamics (LQCD) applications, e.g., MILC and Chroma, onto GPUs avoiding this rewriting overhead. The approach taken is a community-wide library (QUDA) which provides high-performance implementations for the time-critical LQCD algorithms thereby providing instant GPU acceleration. We discuss some of the bleeding-edge strategies taken by QUDA to maximize performance, including the use of communication reducing algorithms, mixed-precision methods and an aggressive auto-tuning methodology. We discuss various compile-and-run strategies to circumvent Amdahl's law, including the use OpenACC directives or retargeting the underlying DSL to generate GPU code directly from the original source.

Opening Frontiers with Extreme Capability Computing

Fred Streitz (Lawrence Livermore National Laboratory)

Lawrence Livermore has a long history of fielding some of the world's largest computers, fueled by our nearly insatiable need for both capacity and capability computing. Each new generation in computing brings with it the ability to perform simulations that were impossible with the earlier computers. I will discuss the development of three applications designed to run on the largest computer on earth: the solidification of a molten metal, the development of a fluid instability, and the electrophysiology of a beating human heart. Although the enormous scale of these simulations uncovered a number of unexpected technical hurdles, resolution of these challenges allows us to overcome limitations of both time and length scale–opening windows to discovery and spurring innovation.

This work is performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-¬-AC52-¬-07NA27344.

Panel on HPC Applications and Extreme Capability Computing

Sadaf Alam (Swiss National Supercomputing Centre), M. A. Clark (NVIDIA Corporation), Fred Streitz (Lawrence Livermore National Laboratory)

This panel discussion will focus on technical as well as career and professional development topics.

Workshop on Parallel, Distributed, and High Performance Computing in Undergraduate Curricula (EduPDHPC) (Part 1)

Chair: Steven Brandt (Louisiana State University)
10:30am-5pm
Room: 708/710/712

Prasad Sushil (Georgia State University)

Parallel and Distributed Computing (PDC) now permeates most computing activities. It is no longer sufficient for even novice programmers to acquire only traditional sequential programming skills. This workshop on the state of art in Parallel, Distributed, and High Performance Computing Education will include contributed as well as invited papers from academia, industry, and other educational and research institutes on topics pertaining to the teaching of PDC and HPC topics in the Computer Science, Computer Engineering, Computational Science, and Domain Science and Engineering (and related) curricula. The emphasis of the workshop will be on undergraduate education, and the target audience will include attendees from among the SC13 Educators and related programs, academia, and industry. This effort is in coordination with the NSF/TCPP curriculum initiative for CS/CE undergraduates (www.cs.gsu.edu/~tcpp/curriculum/index.php), and with the SC13 HPC Educators Program.

LittleFe Setup

Chair: Charles Peck (Earlham College)
10:30am-5:00pm
Room: 702/704/706

Charlie Peck (Earlham College), Aaron Weeden (Shodor Education Foundation), Mobeen Ludin (Shodor Education Foundation), Skylar Thompson (University of Washington), Kristin Muterspaw (Earlham College), Elena Sergienko (Earlham College), Tom Murphy (Contra Costa College)

LittleFe is an educational appliance designed to make teaching parallel and distributed computing easy to do and inexpensive. The SC13 buildout will feature LittleFe units with dual-core Intel® Atom processors and NVIDIA ION2 chipsets which support CUDA/OpenCL programming. On Sunday participants will assemble and test the hardware for their LittleFe units.

Broader Engagement Session V: Mentor/Protégé Session & Mixer

Chair: Gerardine F. Lamble (Santa Clara University)
1:30pm-3pm
Room: 705/707/709/711

Geri Lamble (Santa Clara University), Natasha Floersch (University of Colorado), Richard Coffey (Argonne National Laboratory), Tony Drummond (Lawrence Berkeley National Laboratory)

The Mentor-Protégé event is organized by the Broader Engagement committee to support the development of underrepresented professionals in HPC by creating opportunities for attendees to establish developmental professional relationships.

By pairing new SC participants with more experienced attendees, the Mentor-Protégé program provides introduction, engagement and support. Beginning with a kick-off exercise of digital storytelling to facilitate introductions and orientations, Mentor-Protégé pairs then move on to discuss both networking and technical development pathways to help BE newcomers to the conference integrate with the larger SC community.

A pre-conference technical development activity supported by the Mentor Protégé program is poster help mentoring made available to SC BE participants wherein applicants submitting a poster to the SC conference can receive pre-submission mentoring help.

Digital Storytelling: Developing Your Stories

Storytelling is a fundamental part of being human. Stories build understanding and improve collaboration between groups of different experiences. Personal stories provide us reminders of our past and allow others a glimpse of our worldview. Professional stories communicate to managers or collaborators your unique value.

Natasha Floersch from the University of Colorado,Denver, and Richard Coffey from the Argonne Leadership Computing Facility will lead a 50-minute session on developing our stories through Broader Engagement at SC13. This storytelling exercise will teach you the basics of developing a digital story as well as help you to continue to develop your professional network at Broader Engagement and SC13.

SV Professional Development I: Navigating a Computing Career

Chair: Jeanine Cook (New Mexico State University)
3:30pm-5pm
Room: 705/707/709/711

Patty Lopez (Intel Corporation)

The Professional Development Session on Navigating a Computing Career offers communication tips and strategies that are essential in cultivating professional relationships and building a trusted network of allies and advocates who can guide your career. Topics to be covered include delivering your elevator speech, honing your presentation skills, working with others, establishing credibility, and finding strong mentors. Attendees will have the opportunity to practice their elevator speech, learn why presentation skills are essential in sharing knowledge and expertise, demonstrating leadership, and creating visibility for their contributions, discover how their preferred work style meshes with those of their colleagues, and the secret of approaching a potential mentor.

Tuesday, November 19

Special Session on Education and Workforce Development

Chair: Almadena Chtchelkanova
(National Science Foundation)
10:30am-12pm
Room: 708/710/712

Scott Lathrop (University of Illinois at Urbana-Champaign), John E. West (Department of Defense), Wilf Pinfold (Intel Corporation), Barbara Horner-Milller (University of Alaska Fairbanks)

There are numerous national and international reports highlighting the critical important of preparing current and future generations of computational scientists, technologists, engineers, mathematicians, and scholars in all fields of study. Further, there is an essential need to engage more under-represented minorities and women across all fields of study. There are many challenges and countless opportunities for addressing the need to prepare a larger and more diverse community able to advance scientific discovery. The goal of this session is to gather community input to identify strategic activities that the community might pursue to support learning among the HPC community. The session leaders will describe work that has already been accomplished to advance HPC learning, and they will invite the community to share other challenges, accomplishments, and opportunities.

SV Professional Development II: Building Your Technical Resume

Chair: Jeanine Cook (New Mexico State University)
10:30am-12pm
Room: 601/603

Geri Lamble (Santa Clara University)

We will provide students with detailed recommendations for preparing a high-quality resume. We will also encourage students to participate in activities that lead to useful resume content.

LittleFe Buildout

Chair: Charles Peck (Earlham College)
10:30am-5:00pm
Room: 702/704/706

Charlie Peck (Earlham College), Aaron Weeden (Shodor Education Foundation), Mobeen Ludin (Shodor Education Foundation), Skylar Thompson (University of Washington), Kristin Muterspaw (Earlham College), Elena Sergienko (Earlham College), Tom Murphy (Contra Costa College)

LittleFe is an educational appliance designed to make teaching parallel and distributed computing easy to do and inexpensive. The SC13 buildout will feature LittleFe units with dual-core Intel® Atom processors and NVIDIA ION2 chipsets which support CUDA/OpenCL programming. On Sunday participants will assemble and test the hardware for their LittleFe units.

Serious Computational Examples for Science Classes Featuring Python, Mathematica, an eTextBook and More (Part 1)

Chair: Thomas Hacker (Purdue University)
1:30pm-5pm
Room: 705/707/709/711

Rubin Landau (Oregon State University), Richard Gass (University of Cincinnati)

Part I provides examples of research-level, high performance computing that can be used in courses throughout the undergraduate curriculum. At present such examples may be found in specialty courses in computational physics, although those courses too often focus on programming and numerical methods. In contrast, disciplinary classes tend to use computing as pedagogic tools without understanding the computation. The examples contain a balance of modern computational methods, programming, and interesting physics and science. The Python examples derive from an eTextBook including video-based lectures, Python programs, applets, visualizations and animations. The Mathematica examples focus on nonlinear dynamic, quantum mechanics, program optimization and visualizations. Whether using Mathematica or Python, the session looks inside the computation black box to understand the algorithms and to see how to scale them to research-level HPC. Part II covers materials not covered in Part I including suggested HPC tools and subjects for undergraduates.

Strategies for Introducing Parallelism with Python

Chair: Steven Brandt (Louisiana State University)
1:30pm-5pm
Room: 708/710/712

Steven Bogaerts (DePauw University), Joshua V. Stough (Washington and Lee University)

The Python programming language has seen rapid growth in popularity both in academia and industry. As a lightweight high-level language that supports both functional and object-oriented programming, it provides many tools to allow programmers to easily express their ideas. This expressiveness extends to programming using parallelism and concurrency, allowing the early introduction of these increasingly critical concepts in the computer science core curriculum. In this half-day session we describe and demonstrate an educational module on parallelism, including materials on distributed systems and parallel image processing. We cover such key concepts as speedup, divide & conquer, communication, and concurrency. We consider how these concepts may be taught in the context of CS1 and CS2, and we provide extensive hands-on demonstrations of parallelized search, sort, image processing, and distributed computing.

SV Professional Development III: HPC Careers and Networking Panel

Chair: Jeanine Cook (New Mexico State University)
3:30pm-5pm
Room: 601/603

A panel of HPC leaders will briefly present their personal paths in HPC with recommendations for students exploring opportunities for their future. There will be time for questions and informal networking with the panelists. Panelists include: Les Button, Corning, Inc.; Dona Crawford, Lawrence Livermore National Laboratory; Dean Klein, Micron Technology, Inc.; Wilf Pinfold, Intel; Amanda Randles, Lawrence Livermore National Laboratory

Wednesday, November 20th

An Educator's Toolbox for CUDA (Part 1)

Chair: Steven Brandt (Louisiana State University)
10:30am-12pm
Room: 708/710/712

Karen L. Karavanic (Portland State University), Jens Mache (Lewis and Clark College), David Bunde (Knox College)

GPUs (graphical processing units) with large numbers of cores are radically altering how high performance computing is conducted. Programming frameworks such as CUDA and OpenCL can be used to program GPUs for computational tasks and achieve orders of magnitude improvement in performance over using the CPU alone. The importance of this approach, combined with the easy and inexpensive availability of hardware, combine to make this an excellent classroom topic. The purpose of this workshop is to provide CS educators with the fundamental knowledge and hands-on skills to teach CUDA materials. Three session leaders with a combined total of over three decades of teaching experience will present short lectures, exercises, course materials, and panel discussions.

Exploring Parallelization Strategies at the Undergraduate Level

Chair: Almadena Chtchelkanova (National Science Foundation)
10:30am-12pm
Room: 702/704/706

Eduard Ayguade (Universitat Politècnica de Catalunya), Rosa Maria Badia (Barcelona Supercomputing Center), Vladimir Subotic (Barcelona Supercomputing Center)

We are living the "real" parallel computing revolution. Something that was the concern of a "few" forefront scientists has become mainstream and of concern to every single programmer. This HPC Educator Session presents a Linux environment to be used at the undergraduate level to analyze different task decomposition strategies and their potential benefit. The core of the environment is Tareador, a tool that provides a very intuitive approach to visualize these strategies and understand their implications. The programmer needs to use simple code annotations to identify tasks and Tareador will dynamically build the computation task graph, identifying all data-dependencies among the annotated tasks. Tareador also feeds Dimemas, a simulator to predict the potential of the proposed strategy and visualize an execution timeline (Paraver). The environment also includes a decomposition explorer that implements a top-down approach to explore a rich set of potential task decomposition strategies (task decomposition and granularity).

An Educator's Toolbox for CUDA (Part 2)

Chair: Steven Brandt (Louisiana State University)
1:30pm-3pm
Room: 708/710/712

Karen L. Karavanic (Portland State University), Jens Mache (Lewis and Clark College), David Bunde (Knox College)

This is a continuation of Part 1.

CSinParallel: Using MapReduce to Teach Data-intensive Scalable Computing across the CS Curriculum

Chair: Thomas Hacker (Purdue University)
1:30pm-5pm
Room: 702/704/706

Richard A. Brown (St. Olaf College), Elizabeth Shoop (Macalester College), Joel C. Adams (Calvin College)

MapReduce, the cornerstone computational framework for cloud computing applications, has star appeal to draw students to the study of parallelism. Participants will carry out hands-on exercises designed for students at CS1/intermediate/advanced levels that introduce data-intensive scalable computing concepts, using WebMapReduce (WMR), a simplified open-source interface to the widely used Hadoop MapReduce programming environment, and using Hadoop itself. These hands-on exercises enable students to perform data-intensive scalable computations carried out on the most widely deployed MapReduce framework used by Facebook, Microsoft, Yahoo, and other companies. WMR supports programming in a choice of languages (including Java, Python, C++, C#, Scheme); participants will be able to try exercises with languages of their choice. Workshop includes brief introduction to direct Hadoop programming and information about access to cluster resources supporting WMR. Workshop materials will reside on csinparallel.org, along with WMR software. Intended audience: CS instructors. Laptop required (Windows, Mac, or Linux).

Broader Engagement Session VI: Journal Publishing 101

Chair: Mary Ann Leung (e_SciTek)
3:30pm-5pm
Room: 703

Jeff Hollingsworth (University of Maryland), Rebecca Capone (Elsevier), Victor Prasanna (University of Southern California)

You've completed your experiments and want to publish your work. Are you wondering where to start? Are you confused about the roles of workshop, conference, and journal papers? Do you know how to pick an appropriate journal for your specific work? Do you want to learn how to structure your paper or what to include in a cover letter? Is the submission and peer review process a mystery to you? Perhaps you've submitted your work before and are in need of some guidance interpreting and responding to reviewer comments. This session will provide new researchers, and those new to HPC, a chance to hear about the journal publishing process from the people involved: editors-in-chief and publishers of leading HPC journals.

Thursday, November 21

Going Parallel with C++11 (Parts 1&2)

Chair: Thomas Hacker (Purdue University)
10:30am-Noon/1:30pm-3:00pm
Room: 702/704/706

Joe Hummel (University of Illinois at Chicago), Jens Mache (Lewis and Clark University)

As hardware designers turn to multi-core CPUs and GPUs, software developers must embrace parallel programming to increase performance. No single approach has yet established itself as the "right way" to develop parallel software. However, C++ has long been used for performance-oriented work, and it's a safe bet that any viable approach involves C++. This position has been strengthened by ratification of the new C++0x standard, officially referred to as "C++11". This interactive session will introduce the new features of C++11 related to parallel programming, including type inference, lambda expressions, closures, and multithreading. The workshop will close with a brief discussion of other technologies, in particular higher-level abstractions such as Intel Cilk Plus and Microsoft PPL.

Doctoral Showcase

Thursday, November 21

Doctoral Showcase - Dissertation Research

Chair: Alistair Rendell (Australian National University)
10:30am-12pm
Room: 601/603

Towards Effective High Performance Computing in the Cloud

Abhishek Gupta (University of Illinois at Urbana - Champaign)

The advantages of the pay-as-you-go model, elasticity, and the flexibility and customization offered by virtualization make cloud computing an attractive option for meeting the needs of some HPC users. There is a mismatch, however, between cloud environments and HPC requirements. The poor interconnect and I/O performance in cloud systems, HPC-agnostic cloud schedulers, and the inherent heterogeneity and multi-tenancy in cloud are some bottlenecks for HPC in cloud.

This thesis goes beyond the research question: "What is the performance of HPC in cloud?" and explores "How can we perform effective and efficient HPC in cloud?" To this end, we adopt the complementary approach of making clouds HPC-aware, and HPC runtime system cloud-aware. Through intelligent application-to-platform mapping, HPC-aware VM placement, interference-aware VM consolidation, and cloud-aware HPC load balancing, we demonstrate significant benefits for both: users and cloud providers in terms of cost (up to 60%), performance (up to 45%), and throughput (up to 32%).

Algorithm/Architecture Codesign of Low Power and High Performance Linear Algebra Compute Fabrics

Ardavan Pedram (University of Texas at Austin)

We show the design of specialized compute fabrics that maintain the efficiency of full custom hardware while providing flexibility to execute a whole class of coarse-grain linear algebra and FFT operations. We have designed a specialized linear algebra processor (LAP) that can perform level-3 BLAS, more complex LAPACK level operations like Cholesky, LU (with partial pivoting), and QR factorizations, and even FFT operations. We present a power performance model that compares state of the art CPUs and GPUs with our design. Our power model reveals sources of inefficiencies in CPUs and GPUs, and our LAP design demonstrates how to overcome them. When compared to other conventional architectures for linear algebra applications, LAP is over orders of magnitude more power efficient. Based on our estimations up to 25 GFLOPS/W and double-precision efficiency is achievable on a single chip in standard 45nm technology.

Towards Semi-Automated Dense Linear Algebra Library Generation

Bryan A. Marker (University of Texas at Austin)

Design by Transformation (DxT) is an approach to software engineering where knowledge of software is encoded as transformations. Currently, experts implicitly apply this knowledge to manually create high-performing code. The long-term goal of my work is to encode expert knowledge to automatically apply it for more sustainable, efficient, and formal engineering. Without being explicitly stored, experts' critical knowledge can be lost when an expert retires or leaves. Further, when this knowledge is applied manually, it can be incorrectly applied, resulting in coding errors, or forgotten to be applied, resulting in inefficiencies. This is the case for domains such as dense linear algebra (DLA); experts inefficiently develop DLA code libraries manually. Instead, we should automatically generate such libraries. With DxT we encode knowledge and automatically generate the same or better code (as seen in my experiments with dense linear algebra). Code generated by my work is shipping with the Elemental library.

Addressing Shared Resource Contention in Datacenter Servers

Sergey Blagodurov (Simon Fraser University)

Servers are major energy consumers in modern datacenters. Much of that energy is wasted because applications compete for shared resources and suffer severe performance penalties due to resource contention. My thesis addresses the contention problem.

Enabling Efficient Data Movement in MPI and PGAS Models on Heterogeneous Clusters with High Performance Interconnects

Sreeram Potluri (Ohio State University)

Modern clusters with accelerators (GPUs) and coprocessors (Xeon Phis), while providing high compute density and high performance per watt, introduce heterogeneity and require developers to use a hierarchy of programming models and paradigms to design applications. This has increased the complexity and cost of data movement on these clusters. We propose MVAPICH2-GPU, a model to use MPI for communication from CPU and GPU memories, in a unified manner. We also extend the OpenSHMEM PGAS model to support such unified communication. These models considerably simplify data movement in MPI and OpenSHMEM applications running on GPU clusters. We also propose designs in MPI and OpenSHMEM runtimes to transparently optimize data movement on GPU clusters, using state-of-the-art GPU technologies such as CUDA IPC and GPUDirect RDMA. Further, we present an optimized MPI runtime for communication on clusters with Intel Xeon Phi coprocessors. We evaluate our designs using micro-benchmarks, application kernels and end-applications.

Empirical Search to Optimize Matrix Computation

Thomas Nelson (University of Colorado, Boulder)

High performance linear algebra is an important component to many numerical applications. Efficient matrix computation requires careful optimization based on application and hardware, typically a time consuming process. Today most scientific applications use BLAS (Basic Linear Algebra Subprograms) to provide high-performance linear algebra computation. This research presents an alternative approach, based on an empirical search strategy for linear algebra optimization. It uses Build-to-Order BLAS, a domain-specific language for linear algebra that produces highly-optimized multicore C code as output. This approach maintains a high-level understanding of the operations, which allows more aggressive code transformation. The search is based on a highly-tuned genetic algorithm, which performs global search over a transformation space while maintaining correctness and using application-specific heuristics. Our results show significant speedups compared to traditional library-based approaches for many matrix kernels.

Doctoral Showcase - Dissertation Research

1:30pm-3pm
Room: 601/603

Multi-Receptor High-Throughput Virtual Docking on Supercomputers with VinaMPI

Sally Ellingson (University of Tennessee, Knoxville)

VinaMPI has been developed to enable massively large virtual drug screens on leadership-class computing resources, using a large number of cores to decrease the time-to-completion of the screen. VinaMPI is a massively parallel Message Passing Interface (MPI) program based on the multithreaded virtual docking program AutodockVina, and is used to distribute tasks while multithreading is used to speed-up individual docking tasks. VinaMPI uses a distribution scheme in which tasks are evenly distributed to the workers based on the complexity of each task, as defined by the number of rotatable bonds in each chemical compound investigated. VinaMPI efficiently handles multiple proteins in a ligand screen, allowing for high-throughput inverse docking that presents new opportunities for improving the efficiency of the drug discovery pipeline. VinaMPI successfully ran on 84,672 cores with a continual decrease in job completion time with increasing core count.

Automated Floating-point Program Analysis

Michael O. Lam (University of Maryland, College Park)

As scientific computation continues to scale, it is crucial to use floating-point arithmetic as efficiently as possible. Lower precision allows streaming architectures to perform more operations per second and can reduce memory bandwidth pressure. However, using a precision that is too low will result in inaccurate results. Thus, developers must balance speed and accuracy when choosing the floating-point precision of their subroutines and data structures.

I have developed program analysis techniques to help developers learn about the runtime floating-point behavior of their programs, with the goal of helping them make decisions concerning the choice of precision in implementation. I have developed automated methods for cancellation detection, range tracking, mixed-precision configuration, mixed-precision search, and bit-level sensitivity analysis. This body of work improves the state of the field of program analysis and scientific computing by providing insights that were not previously available using automated analysis techniques.

High Performance Runtime for Next Generation Parallel Programming Languages

Vivek Kumar (Australian National University)

Processor design is moving toward large-scale parallelism, heterogeneous cores and accelerators to achieve performance and energy efficiency. Work-stealing is proven to be a promising approach for productively exploiting the software parallelism on these modern hardware. Despite the benefits, the work-stealing approach incurs some form of overheads as a necessary side effect of its implementation.

This thesis identifies key sources of overheads in the work-stealing scheduler, namely sequential and dynamic overhead. It proposes several novel techniques to reduce these overheads and evaluates the design using a range of benchmarks, and a variety of work-stealing implementations. Our new design can reduce the sequential overheads to just 15% and almost halves the dynamic overhead, leading to substantial performance improvements. These results and our insight into the sources of overhead for work-stealing implementations give further hope to an already promising technique for exploiting increasingly available hardware parallelism.

High Performance Computing for Irregular Algorithms and Applications

Oded Green (Georgia Institute of Technology)

Irregular algorithms such as graph algorithms, sorting, and sparse matrix multiplication present numerous programming challenges that include scalability, load balancing, and efficient memory utilization. The advent of Big Data increases the challenges of analyzing massive data sets such as Facebook and Twitter which change constantly. Thus, the continuous updating of the analytics of such data sets proves to be a big challenge. One such solution has been through the usage of special HPC systems like the Cray XMT. My current research focuses upon creating new algorithms for irregular algorithms with a particular emphasis on graph analytics which allow

monitoring of different types of networks. For example, Betweenness Centrality (BC) is an important analytic used to find key players in social networks. However, computing BC usually requires more time than the characteristic time in which the graph is updated thus making the entire process virtually useless.

Autonomic Failure Identification and Diagnosis for Building Dependable Computing Systems

Qiang Guan (University of North Texas)

In modern cloud computing systems, hundreds and even thousands of cloud servers are interconnected by multi-layer networks. In such large-scale and complex systems, failures are common. Autonomic failure detection and diagnosis are crucial technologies for understanding emergent, cloud-wide phenomena and self-managing resource burdens for cloud availability and productivity enhancement. They are determining the quality of the services that cloud could provide to the customers. In my PhD dissertation, I am focusing on developing an autonomic failure management mechanism with the assurance of dependability to the cloud services. My work first sets out by studying the health-related performance metric. With metric selection and metric extraction technologies, failures related metrics are selected. Then learning from the characteristics of metrics from time domain and frequency domain separately, failure detection and diagnosis are achieved dynamically.

An Auto-Parallelizing Library for Computer Simulations Which Scales From Smartphones to Supercomputers

Andreas Schaefer (Universität Erlangen-Nürnberg)

During my Ph.D .I have developed LibGeoDecomp (www.libgeodecomp.org), an auto-parallelizing library that allows researchers to make use of supercomputing resources they would otherwise not be able to tackle. The library is based on C++ templates and can harness all current supercomputer architectures (x86, x86 + accelerators (e.g. Tsubame 2.0), embedded CPUs (JUQUEEN, an IBM Blue Gene/Q)). One of its major contributions is the callback interface between user code and library: the interface is crafted in a way that allows legacy codes to be ported to the library while still enabling the library to apply powerful optimizations (e.g., cache blocking, GPU offloading, or letting a struct-of-arrays storage appear like an array-of-structs).

So far I have scaled the library up to 450k cores, 1080 GPUs, and 1.8M MPI ranks (4x overcommit on BG/Q). The zero-overhead interface and the library's portability and scalability set it apart from previous work.

Thursday, November 21

Doctoral Showcase - Early Research Showcase

Chair: Wojtek James Goscinski (Monash University)
3:30pm-5pm
Room: 601/603

Fast Multipole Method as a Preconditioner

Huda Ibeid (King Abdullah University of Science & Technology)

We employ FMMs as a preconditioner for sparse matrix solvers. We test our FMM-based preconditioner for the conjugate gradient method for elliptic problems. Unlike most previous FMM preconditioner/solver practice, where problems with free-field boundary conditions were investigated, we apply our preconditioner to Dirichlet/Neumann boundary conditions by coupling FMM with a boundary element method. The convergence rate is compared with Jacobi, ILU, geometric multigrid, and algebraic multigrid methods. The rate of convergence of our FMM-based preconditioner is similar to that of algebraic multigrid.

Taxonomy Cube: A Multi-Dimension Application-to-Architecture Mapping

Karan Sapra (Clemson University)

High application performance greatly depends on the choice of accelerator for the given application. In this paper, we propose an Application-to-Algorithm (A2A) Taxonomy Cube that maps an application from the algorithm space to an appropriate accelerator in the architecture space for optimal performance. The Taxonomy Cube is based on four major application performance factors, labeled performance dimensions, including computation-to-communication ratio, FLOPs (floating-point operations) to Non-FLOPs, device memory accesses, and Non-Uniform Memory Accesses (NUMA). We aim to evaluate the accuracy of A2A mapping by executing diverse representative applications on target architectures and formulation of performance models that will allow developers to further fine-tune their applications, once an appropriate A2A mapping is identified. The final research goal is to provide a set of guidelines for optimal A2A in the form of a Taxonomy Cube.

A Unified Sparse Matrix Format for Heterogeneous Systems

Moritz Kreutzer (Erlangen Regional Computing Center)

Sparse matrix-vector multiplication (spMVM) is the most time-consuming kernel in many numerical algorithms and has been studied extensively on all modern processor and accelerator architectures. However, the optimal sparse matrix storage format is highly hardware-specific, which could become an obstacle when using heterogeneous systems. Also, the efficient utilization of the wide SIMD units in current multi- and many-core processors in the context of unstructured spMVM has not

undergone thorough research. In this work, SELL-C-sigma, a variant of Sliced ELLPACK, is suggested as a SIMD-friendly data format. The advantages of SELL-C-sigma compared to established formats like CRS and ELLPACK and its suitability on a variety of hardware platforms (Intel Sandy Bridge, Intel Xeon Phi and Nvidia Tesla K20) for a wide range of test matrices from different application areas are shown. This leads to a hardware-independent sparse matrix format, which achieves very high efficiency for all test matrices across all architectures.

Local Memory Store - A Multi Purpose Memory Replacing On-Chip Cache

Nafiul A. Siddique (New Mexico State University)

Cache performance is poor for graph based applications due to high miss rate and excess data transaction to on-off chip memory. To improve the performance it is required to store the working set in an on-chip memory as scratchpad memory (SPM). SPM is a fast on-chip memory mapped into a predefined address space. But predicting the address range and implementing appropriate allocation technique is hard.

We propose a hardware controlled memory (Local Memory Store (LMStr)) that stores different sizes of blocks and can simultaneously be partitioned into and used as scratchpad and/or cache. We introduce a weighted based static searching strategy that predicts the block references to store in scratchpad space. To investigate the performance benefits of LMStr, we have created a functional simulator, and a block generator. Our results show that the miss rate reduces, but the number of words fetched to/from main memory, and overhead increases in LMStr.

Design and Analysis of Parallel Multiscale Algorithms for Modeling Platelets

Na Zhang (Stony Brook University)

We design and analyze a family of multidisciplinary integrated multi-scale parallel schemes for modeling flowing platelets in viscous blood plasma at disparate spatiotemporal scales. To effectively carry out our model, we customize LAMMPS software package and enable a simulation system size consisting of 100 million particles. To efficiently conduct numerical experiments, we propose a hybrid CGMD-DPD (coarse-grained molecular dynamics-dissipative particle dynamics) method at overlaying regions of platelets and viscous blood plasma for facilitating rheological information sharing over the spatial temporal scales. Through supercomputing and numerical variation separation technologies, we parameterize these algorithms and systematically analyze the impacts of model parameters on the accuracy of representing the physics observables and biomedical phenomena for optimal computing efforts.

Finally, we apply our optimal algorithms to realistic multi-scales studies of platelet adhesion and scale our model to thousands of processors on the Sunway Blue Light Massive Parallel Processing (MPP) system.

End-to-End Data Transfer Systems for 100 Gbps Networks and Beyond

Yufei Ren (Stony Brook University)

Data-intensive applications place stringent requirements on the performance of both back-end storage systems and front-end network interfaces in modern multi-core servers. However, for transferring data at ultra-high-speed, for example, at 100 Gbps and higher, the effects of multiple bottlenecks along a full end-to-end path, have not been resolved efficiently. My research leverages remote direct memory access (RDMA) and non-uniform memory access (NUMA) technologies to fully utilize bare-metal performance of state-of-the-art hardware. Particularly, my research designs network protocols for data-intensive applications based on RDMA one-sided operations and asynchronous programming interfaces across both local and wide area networks in front-end systems, and a NUMA-aware cache mechanism to align cache memory with the local NUMA node and schedule I/O requests to those threads that are local to the data to be accessed in back-end storage systems. These designs and implementation can significantly improve the performance of end-to-end data transfer systems.

Performance Analysis Towards Adaptive Thread Scheduling Policy

Patricia Grubel (New Mexico State University)

We conduct performance studies of the High Performance ParalleX (HPX) thread scheduling policies using a task scheduling micro benchmark, the Unbalanced Tree Search (UTS) benchmark, and an adaptive mesh refinement (AMR) application, on a variety of platforms to determine metrics that will aid in improving scheduling policies. Our objective is to determine if an adaptive management of scheduling policies can be used dynamically for task management among hardware threads on a node to improve performance of scaling impaired parallel applications. Once the metrics have been determined, we will apply results to dynamically adjust factors of the thread scheduling policy for improved performance and efficiency. Future work includes using similar performance studies to assess the possibility of using adaptive techniques on distributed processes in the Active Global Address Space (AGAS) module of HPX.

Ultra Fast Simulations for Ultra Cold Atoms

Christopher J. Watkins (Monash University)

The popularization of the Graphics Processing Unit (GPU), found in most modern desktop computers, has brought the age of supercomputing into the hands of average scientists. Using the highly parallel architecture the GPU we have been developing code that will allow us to probe the interesting quantum behavior of ultra cold atomic gases, generally less than a few hundred micro Kelvin. The project involves two separate codes, one that utilizes the Direct Simulation Monte Carlo technique, to simulate the quantum losses experienced as an atomic gas is evaporatively cooled in a magnetic trap. The other is a Finite Element code, taking advantage of the Discrete Variable Representation (DVR) to perform full three dimensional simulations of the evolution the gas after it has undergone a quantum phase transition. We are able to deploy these codes on the local GPU supercluster, MASSIVE, to gain excellent performance.

P-SLURM: Power-Aware Job Scheduling on a Hardware-Overprovisioned HPC Cluster

Tapasya Patki (University of Arizona)

While traditional supercomputers are worst-case provisioned and guarantee peak power to every node, next-generation supercomputers will require hardware-overprovisioning and intelligent power scheduling. Our results demonstrate that choosing the optimal configuration (number of nodes, number of cores per node, power per node) based on individual application characteristics on an overprovisioned cluster can improve performance under a power bound by up to 62% when compared to worst-case provisioning.

In this work, we propose P-SLURM, a power-aware resource manager for an overprovisioned cluster. P-SLURM takes job-level power bounds into account to implement policies that choose job configurations based on metrics such as average turnaround time and utilization. In addition, it facilitates reallocation of power at runtime as jobs enter and leave the system. We extend SLURM's easy backfilling algorithm with seven power-aware policies, and present some early results obtained by using job traces gathered from Lawrence Livermore National Laboratory.

Making GMRES Resilient to Single Bit Flips

James J. Elliott (North Carolina State University)

The collective surface area and increasing density of components has led to an increase in the number of observed bit flips. This effort works towards rigorously quantifying the impact of bit flips on floating point arithmetic. We exploit norm bounds to ensure the upper Hessenberg entries constructed in the GMRES algorithm do not exceed theoretical limits. We also combine analytical modeling to show how scaling the inputs to these algorithms mitigate silent faults should they occur. We then combine our upper Hessenberg check, analytical model of bit flip perturbations, and the concept of sandbox reliability to show how scaling data can enhance existing fault tolerant algorithms, namely Fault Tolerant GMRES.

Robust Processing of Health Stream Data

Kathleen Ericson (Colorado State University)

Personal health sensors offer the potential to reduce economic costs and improve outcomes. If data from from personal health sensors can be processed in a timely, robust, and failure-resilient manner, patients would be able to leave the hospital sooner after major surgeries and those who require full-time care facilities may be able to stay at home, freeing limited health care resources for emergencies.

Efficient scheduling is needed to meet processing guarantees. Optimal placements are an instance of the resource-constrained scheduling problem, either NP-hard or NP-complete depending on problem characterization. Our focus is to use statistical and ML algorithms to guide placements to satisfy constraints in real time.

We have implemented fault-tolerance through replication and leveraged replicas to enable migrations for load-balancing purposes, allowing quick migrations. We also developed monitoring tools, allowing us to track resource usage on a per-computation basis which can help drive load balancing and computation placement strategies.

Optimizing MPI Implementation on Massively Parallel Many-Core Architectures

Min Si (University of Tokyo)

Many-core architectures, such as the Intel Xeon Phi, provide user applications dozens of cores and hundreds of hardware threads. To efficiently utilize such architectures, application programmers are increasingly looking at hybrid programming models, where multiple threads often work with MPI, which is responsible for data communication. My doctoral research focuses on exploiting the capabilities of such massively threaded architectures in MPI. It can be broadly divided into three related pieces: internal multithreading in MPI, fine-grained consistency management in MPI, and MPI tasklets. We have so far focused on the first topic, enabling MPI to share idle threads from user applications. We parallelized the MPI functionality used in stencil computation routines as a case study. Preliminary results show that the parallel version provides up to a 40-fold speedup in MPI performance when parallelized across the entire Intel Xeon Phi board.

A Framework for Input and Architecture Aware Code Variant Autotuning

Saurav Muralidharan (University of Utah)

With computer systems becoming increasingly parallel and heterogeneous, autotuning is turning out to be essential in achieving performance portability for parallel applications.

My current research focuses on input-aware autotuning, since autotuners that take characteristics of the input data along with other execution context into account can optimize more aggressively. I describe a programmer-directed autotuning framework that facilitates autotuning of code variants, or alternative implementations of the same computation, and selects code variants at execution based on features of the input data set.

In real-world irregular benchmarks from sparse numerical methods and graph computations, variants selected using our system achieve up to 99.85% of the performance achieved by variants selected through exhaustive search.

I also discuss my future research focus: how autotuning systems can be made more accessible to non-expert users by integrating them with higher-level programming models that decouple computation and implementation.

Multiscale Modeling of Fine-Grained Platelet Suspension in Coarse-Grained Blood Flow Using Molecular Dynamics and Dissipative Particle Dynamics

Chao Gao (Stony Brook University)

The blood coagulation cascade may be initiated in prosthetic cardiovascular devices and cardiovascular pathologies by flow-induced platelet activation, leading to thrombus formation. Platelets undergo complex biochemical and morphological changes upon activation, aggregate, and adhere to blood vessels or device surfaces to form thrombi. Blood coagulation is inherently a multiscale problem: numerical simulations of this complex process traditionally treat blood as a continuum, typically by solving the Navier-Stokes equations to resolve the process to μm scales at the most. However, molecular mechanisms governing blood coagulation are on the nm scales. To address this challenge, we developed a multiscale approach based on Dissipative Particle Dynamics (DPD) and Coarse Grained Molecular Dynamics (CGMD) methods to bridge the gap between macroscopic flow and cellular scales. This model provides greater resolution for simulating platelets dynamics and kinematics and will be employed to simulate flow induced platelet shape change upon activation.

Optimization Methods for Energy Efficient HPC Architectures

Olga Datskova (University of Houston)

Our current research focuses on developing optimization methods for high-impact algorithms on low-power devices. Over the years, growing energy efficiency concerns within HPC (High Performance Computing) have brought to light a strong need for efficient software to take full advantage of power-optimized hardware. Exploitation of locality, data allocation and scheduling of operations in multi-core SoC (System on Chip) designs are critical and difficult for performance and energy efficiency in increasingly complex memory systems. In the present work, we focus on optimizing FFT (Fast Fourier Transform) as a case study for low-level data management and instruction scheduling in order to achieve execution and energy efficiencies close to the peak achievable by the architecture, the 40 nm TI (Texas Instruments) TMS320C6678 8-core DSP (Digital Signal Processor) with 64-bit floating point units and cores having a potential peak 64-bit performance of 7.5 GF at about 1W.

Techniques for Improving Observed Network Performance

Nikhil Jain (University of Illinois at Urbana-Champaign)

Research on understanding application communication performance on supercomputer networks is important because the scalability of many applications is adversely affected by communication overheads when running on large systems. My preliminary research has been focused on studying various aspects, in the network and in the software, that impact the observed network behavior. So far, I have worked on the following projects: 1) techniques for correlating performance with factors related to message communication such as link bandwidth, routing scheme etc. (I will present a paper based on it at SC13), 2) topology aware mapping, 3) use of simulation and modeling for predicting performance, and 4) topology-aware collective algorithms (papers presented at ICS, Euro-MPI). As I make progress in my doctorate program, the research will focus on end-to-end interaction of networks and applications. This will include classification of impacts of various factors on network performance, fast methods for performance prediction and design/study of future interconnects.

Inter-Application Coordination for Reducing I/O Interference

Sagar Thapaliya (University of Alabama at Birmingham)

Scientific applications running on HPC systems share common I/O resources, including parallel file systems. This sharing often results in inter-application I/O interference, leading to degraded application I/O performance. In this work, we study the benefits of coordinating between applications to reduce I/O interference and develop methods to perform this coordination.

Transactional Memory as Synchronization Solution for Parallel Graph Exploration

Mohammad Qayum (New Mexico State University)

Synchronization of parallel applications like large graphs is very complex. Two major synchronization policies are used in most the architectures are coarse grain, which sequentializes program execution, and fine grain, which is notoriously complex. Transactional memory (TM) attempts to mitigate complexity in programming and also provides concurrency. Parallel graph algorithms are mostly based on irregular data structures. In a large graph it is very unlikely that multiple threads will try to access the same vertex at the same time. TM benefits when there is low number of conflicts among the transactions. And, if enough numbers of transactions exist, they can run concurrently too. Some research results with graph benchmarks like SSCA-2 using Software based TM shows some improvements. Currently, we are working with graph benchmarks to use hardware TM support implemented in Intel Haswell architecture and believe that HTM will perform better than STM used in previous research.

Student Cluster Competition

Monday, November 18

Student Cluster Competition Kickoff

Chair: William Douglas Gropp (University of Illinois at Urbana-Champaign)
7pm-9pm
Room: Booth 3146

Join SC13 Conference Chair Bill Gropp to cheer on 12 twelve student teams as they embark on a real-time, non-stop, 48-hour challenge to assemble a computational cluster on the exhibit floor and race to demonstrate the greatest sustained performance across a series of applications.

Tuesday, November 19

10am-6pm
Wednesday, November 20
10am-4:30pm

Student Cluster Competition

Chair: Dustin Leverman (Oak Ridge National Laboratory)
Room: Booth 3146

Meet the 12 twelve student teams competing in this real-time, non-stop, 48-hour challenge to assemble a computational cluster on the exhibit floor and demonstrate the greatest sustained performance across a series of applications. Teams selected for this year's competition come from universities (and one high school) in the United States, Germany, China and Australia. The competition is divided into two tracks---the Standard Track, using big iron hardware with a 26 amp power limit, and the Commodity Track, using off-the-shelf (or off-the-wall!) hardware with a $2500 price limit and a 15 amp power limit.

Wednesday, November 20

Get Involved in the Student Cluster Competition

Chair: Daniel Kamalic (Boston University)
1:30pm-3pm
Room: 703

Come share questions and feedback with teams and organizers of the Student Cluster Competition and learn how your university, high school, or company can be involved next year in the "Premier Event in Computer Sports".

Student Cluster Competition Grand Finale

Chair: Dustin Leverman (Oak Ridge National Laboratory)
4:30pm-5:30pm
Room: Booth 3146

Come cheer on twelve student teams as they turn in their final results in this challenge to assemble a computational cluster on the exhibit floor and demonstrate the greatest sustained performance across a series of applications.

Student Cluster Competition Celebrity Pro-Am Cluster Challenge Kickoff

Chair: Dustin Leverman (Oak Ridge National Laboratory)
5:30pm-6pm
Room: Booth 3146

Come to booth 3146 as we introduce the first ever Celebrity Pro-Am Cluster Challenge, in which student teams from the Student Cluster Competition (the "amateurs") will race head-to-head with professionals in the industry (the "pros") to solve a computational problem set.

Thursday, November 21

Celebrity Pro-Am Cluster Challenge

Chair: Dustin Leverman (Oak Ridge National Laboratory)
10am-11:30am
Room: Booth 3146

Come watch the first ever Celebrity Pro-Am Cluster Challenge, which pits student teams from the Student Cluster Competition (the "amateurs") against industry professionals (the "pros") to solve a computational problem set.

Experiencing HPC for Undergraduates

Monday, November 18

Experiencing HPC For Undergraduates Orientation

Chair: Jeffrey K. Hollingsworth (University of Maryland)
3pm-5pm
Room: 703

Jeffrey K. Hollingsworth, Alan Sussman (University of Maryland)

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 19

Introduction to HPC Research Topics

Chair: Alan Sussman (University of Maryland)
10:30am-12pm
Room: 703

Chris Johnson (University of Utah), Andrew Chien (University of Chicago), John Mellor-Crummey (Rice University), Laura Carrington (San Diego Supercomputer Center)

No abstract

Wednesday, November 20

Graduate Student Perspective

Chair: Jeffrey K. Hollingsworth (University of Maryland)
10:30am-12pm
Room: 703

This session will be held as a panel. Current graduate students, all of whom are candidates for best student paper award in the technical papers program at SC13, 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 21

Experiencing HPC For Undergraduates Breakfast

Chair: Jeffrey K. Hollingsworth (University of Maryland)
7:30am-8:30am
Room: 703

Careers in HPC

Chair: Alan Sussman (University of Maryland)
10:30am-12pm
Room: 703

Boyana Norris (University of Oregon), Rob Schreiber (Hewlett-Packard), David Bernholdt (Oak Ridge National Laboratory), Steve Parker (NVIDIA Corporation), Kathy Yelick (University of California, Berkeley)

SCinet

For the duration of the SC13 conference, Denver will host 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 Colorado 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.

Building on the past successes of the SCinet Research Sandbox, XNet and the Bandwidth Challenge, this program is a new cross-cutting conference program designed to highlight breakthrough network research developments through live demonstrations on the SCinet infrastructure. For SC13, network researchers will highlight network research demon that also tie to technical papers, panels or posters, as well as planned activities in the general SC exhibit area.

SCinet

Thursday, November 21

Network Research Exhibition

10:30am-12pm
Room: 404

Future Network Testbed Enabling Research Projects for International Network Research Testbed Federation Prototype

Jim Chen, Joe Mambretti (Northwestern University), Mon-Yen Luo (National Kaohsiung University of Applied Sciences), Te-Lung Liu (National Applied Research Laboratories), Jian Yang (University of Science and Technology of China), Ronald van der Pol (SURFnet)

In the last few years, large scale network research testbeds have been recognized as necessary instruments for many network and computer science and engineering research projects to allow experimentation at scale. Consequently many countries and research communities have built or are in the process of building large scale network research testbeds to address this need. Even as this trend continues, a related initiative is creating a global capability to allow individual network research testbeds to interconnect through federations across the world. In the last few years, a consortium of network research institutions established an initiative to build large scale distributed testbeds and interconnect them to form a large scale heterogeneous network research testbed – an International Network Research Testbed Federation Prototype (INRTFP), to provide high degree of comprehensive functionality and resource diversity which cannot be achieved by any individual network testbed. Essentially, this testbed environment is a platform on which a wide spectrum of techniques for science research can be conducted. This SC13 exhibition will demonstrate the current status and capabilities of this testbed and it will show selected 5 research projects that contribute to making this testbed highly diverse, resource rich, multifunctional, and easy to use.

Demonstrations of Computational Genomic Analytics at 100 Gbps across a National Data-Intensive Computational Science Testbed

Joe Mambretti (Northwestern University), Bob Grossman (University of Chicago), Don Pruess (National Institutes of Health)

These demonstrations were designed to show novel approaches to large scale genomic analysis utilizing WAN scale 100 Gbps paths with extensions to SCinet. These demonstrations were designed, in part, as a response to an initiative proposed by the White House Office of Science and Technology Policy to showcase Big Data Science Networking at 100 Gbps. This special showcase initiative, which was motivated to demonstrate the importance of high capacity networking for advanced science research, has been promoted by the Large Scale Networking Committee, which coordinates major development activities among federal agency networks.

100 Gbps Networks for Next Generation Petascale Science Research and Discovery

Joe Mambretti (Northwestern University), Bill Fink (NASA), Paul Lange (NASA), Jim Chen (Northwestern University)

These demonstrations were designed to show innovative approaches to transporting large scale data for petascale scientific research using 100 Gbps WAN and LAN paths. These demonstrations were designed, in part, as a response to an initiative proposed by the White House Office of Science and Technology Policy to showcase Big Data Science Networking at 100 Gbps. This special showcase initiative, which was motivated to demonstrate the importance of high capacity networking for advanced science research, has been promoted by the Large Scale Networking Committee, which coordinates major development activities among federal agency networks.

Demonstrations of 100 Gbps Services and Technologies Supporting HPC Clouds

Joe Mambretti (Northwestern University), Robert Grossman (University of Chicago), Rod Wilson (Ciena), Marc Lyonnais (Ciena), Jim Chen (Northwestern University)

These demonstrations were designed to showcase advanced 100 Gbps optical networking services and technologies supporting HPC science clouds based on leading edge dense wave division multiplexing (DWDM), high-performance packet optical platforms, reconfigurable switching systems, and coherent optics based on an international/national testbed established on optical fiber lightpaths, with extensions to the SCinet conference in Denver and to booths on the showfloor. These demonstrations are being designed to underscore the importance of high capacity, flexible, low latency networking for advanced science research.

Demonstrations of Data-Scope at 100 Gbps across a National Data-Intensive Computational Science Testbed at SC13

Joe Mambretti (Northwestern University), Alex Szalay, Jan Vandenberg (Johns Hopkins University), Jack Wells, Jason Hill, Ryan Adamson, Daniel Pelfrey, Scott Kock (Oak Ridge National Laboratory), Eli Dart (Energy Sciences Network), Tom Lehman (Mid-Atlantic Crossroads), Robert Grossman (University of Chicago)

These SC13 Network Research Exhibition demonstrations were designed to show an innovative approach to moving large scale data for scientific analysis, by using Data-Scope. Data-Scope is a novel instrument designed to observe extremely large data sets by utilizing WAN scale 100 Gbps paths, provisioned at multiple sites across a national data-intensive computational science testbed provisioned on ESnet, with extensions to SCinet. These SC13 demonstrations were designed, in part, as a response to an initiative proposed by the White House Office of Science and Technology Policy to showcase Big Data Science Networking at 100 Gbps. This special showcase initiative, which was motivated to demonstrate the importance of high capacity networking for advanced science research, has been promoted by the Large Scale Networking Committee, which coordinates major development activities among federal agency networks.

Provider Backbone Bridging Based Network Virtualization

Ronald van der Pol (SURFnet)

An important use of Software Defined Networking (SDN) is network virtualization or slicing. This allows multiple users on the same physical infrastructure each having their own virtual network or slice. FlowVisor is one of the options to achieve this. It is a software module that sits between the OpenFlow switch and the OpenFlow controllers and gives each controller a part (slice) of the flowspace. The disadvantage of this is that controllers do not have access to the full OpenFlow tuple space and thus the experience is different from having direct access to a physical OpenFlow switch. In this demo we introduce Provider Backbone Bridging (PBB) as defined in IEEE 802.1ah as encapsulation technology in the OpenFlow data plane. In this way user traffic is separated and identified by the I-SID in the encapsulation header. The data part is the user's original Ethernet frame and users can create OpenFlow rules that match on any of fields that OpenFlow 1.3 supports, except for the PBB I-SID because this is used to map packets to users. We think this is a simple virtualisation method that gives users access to virtual OpenFlow switches that have the same OpenFlow capabilities as physical OpenFlow switches would have.

1:30pm-3pm
Room: 404

Advanced Network Analytics at 100 Gbps

Jordi Ros-Giralt (Reservoir Labs)

Reservoir Labs will demonstrate R-Scope, a high-performance cyber security appliance enabling a new generation of network analytics for the detection of sophisticated network threats at ingest bandwidth rates of 100+ Gbps.

R-Scope extends and enhances the functionality of Bro, a highly configurable and programmable network analysis framework, through a scale-out architecture leveraging the processing capabilities of advanced manycore network processors. Architecturally, R-Scope is comprised of a collection of R-Scope cells, each of which is composed of fundamental network security elements mapped to one or more processing cores. These R-Scope cells are individually capable of providing the full functionality of R-Scope and the underlying Bro network analysis. The R-Scope architecture narrows the traditional separation between load balancer and clustered analyzers seen in typical cyber security installations incorporating network analysis directly into the front-end load balancer. A uniform analytic expression language is used at all stages of processing, migrating intelligence closer to the front end to achieve more accurate and efficient discrimination of incoming packet flows.

The hardware used for this demonstration is a 1U, full-length rack-mount R-Scope appliance requiring 490W, 120V or 230V AC power. The R-Scope appliance will connect to the SCinet NRE infrastructure via 10GBASE-SR optical media. Additionally, we will use a 1U rack-mount system for additional control plane activities and situational awareness visualization.

Enhancement of Globus GridFTP over SmartNIC User-Programmable 10GigE NIC

Rajkumar Kettimuthu (Argonne National Laboratory), Gerald Sabin (RNET Technologies)

GridFTP is a high performance and secure protocol widely utilized for bulk file transfer in scientific and commercial grid environments. Globus GridFTP, as the de-facto implementation of GridFTP, plays a significant role in grid-based scientific research and collaboration. Therefore, its efficiency, performance and scalability over high speed networks is critical to the overall success of scientific research and collaboration. In this work, we have used several techniques including transport protocol offload, security processing offload, and network reservation to improve the file transfer experience of Globus GridFTP. We will demonstrate the throughput and host utilization benefits of using such techniques for long-haul file transfers.

Dynamic Monitoring and Adaptation of Data Driven Scientific Workflows Using Federated Cloud Infrastructure

Paul Ruth (University of North Carolina at Chapel Hill)

This demonstration will showcase a novel, dynamically adaptable cloud infrastructure driven by the demand of a scientific data-driven workflow. It will use resources from ExoGENI - a Networked Infrastructure-as-a-Service (NIaaS) testbed funded through NSF's Global Environment for Network Innovation (GENI) project. The demo will connect compute and data resources in the RENCI SC13 booth to a large dynamically provisioned `slice' spanning multiple ExoGENI cloud sites that is interconnected using dynamically provisioned connections from I2, NLR and ESnet. The slice will be used to execute a scientific workflow driven from a computer in the RENCI SC13 booth connected to the slice via SCinet. A closed-loop control mechanism linked to monitoring infrastructure will adapt the slice to the demands of the workflow as it executes.

Application-Aware Flow Optimization in Wide Area Networks Using Distributed OpenFlow Controller

Michael Bredel (California Institute of Technology)

Scientific collaborations on global scale, such as the LHC, rely on the presence of high performance, high availability networks. The efficiency in data movement translates directly into the capability to reach scientific goals in a timely manner. With the increasing scale and complexity of international research collaborations as well as the network infrastructures, new methods of providing efficient data transport are necessary. Software Defined Networking and OpenFlow offer themselves as a new paradigm and a protocol to achieve this goal. In this demo we address the challenge of traffic optimization for large flows in Wide Area OpenFlow Networks. We present a distributed controller implementation used to enable application aware networking. In particular, we present an API that can be used by data movement applications to provide additional information to the OpenFlow controller. Using this information, we demonstrate efficient automatic on-demand path selection and traffic engineering in a multipath network, aiming at optimized network utilization.

Firewalling Science DMZ without Bottlenecks: Using Application-Aware Traffic Steering

RajaRevanth Narisetty, Deniz Gurkan (University of Houston)

Our demonstration of firewalling science DMZ without bottlenecks using Application aware traffic steering leverages the Software Defined Networking paradigm to offload the science DMZ traffic after the firewall's application identification by deep packet inspection.In this respect, as soon as a science data transfer session has been positively identified by the firewall, a flow rule can be written to offload the remainder of the data transfer from the firewall to a fast path on the switch. All sessions are inspected by the firewall, but not all packets of every session need to pass through the firewall. Thus the security is preserved and the throughput constraint is removed.

For this demonstration, we will utilize GENI resources. We would require a high performance compute resource to host firewall virtual appliance(PC with Ubuntu OS, 4 physical NICs, minimum of 8G RAM & 4 cores of CPU), an OVS (PC with Ubuntu OS, 2G RAM & 4 physical NICs), Controller (PC with Ubuntu OS, 2G RAM & 2 physical NICs), AppServer (PC with Ubuntu OS, 2G RAM & 2 physical NICs), Client (Ubuntu VM) and Server (Ubuntu VM).

SDN Innovation Framework Enabling Programmability of Flows within the Network

Levent Dane, Deniz Gurkan (University of Houston)

We demonstrate a middlebox on the path of an OpenFlow network flow in the GENI infrastructure. Our middlebox, called as Split Data Plane (SDP), has the capability of separating manipulating flows at a network node rather than the traditional end-site origination and termination. The middlebox enables application deployment on the path of OpenFlow flows in a dynamic fashion. A different programmability in the networking has been achieved where end points are transmitting their data and network nodes are programmable towards increasing performance, adding functionality in security, etc., and further services that may be possible.


Argonne NATIONAL LABORATORY
Alcatel·Lucent
AFRL DSRC DoD SUPERCOMPUTING RESOURCE CENTER
ARISTA
ARL DSRC A TRADITION OF INNOVATION ATTITUDE FOR SUCCESS
BERKELEY LAB
BROCADE
ca technologies
Caltech
ciena
cisco
CSC
CLEMSON UNIVERSITY
ESnet Energy Sciences Network
extreme networks
FAU FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
FLR Florida LambdaRail Florida's Research and Education Network
FUJITSU
GÉANT
Gigamon
INTERNET2
infinera
inMon
ixia Deliver On
JUNIPER NETWORKS
Los Alamos NATIONAL LABORATORY
Lawrence Livermore National Laboratory
LEARN LONESTAR EDUCATION AND RESEARCH NETWORK
Level(3) COMMUNICATIONS
LEVERAGE INFORMATION SYSTEMS
LONI
NCAR NATIONAL CENTER FOR ATMOSPHERIC RESEARCH
NCSA
NLR
MISSISSIPPI STATE UNIVERSITY
METAFLOWS
NERSC
NICS NATIONAL INSTITUTE FOR COMPUTATIONAL SCIENCES
NOAA
n-wave
OAK RIDGE National Laboratory
PENNSTATE 1855
PURDUE UNIVERSITY
Reservoir Labs
redhat
RICE Unconventional Wisdom
Sandia National Laboratories
Schneider Electric
SDSC SAN DIEGO SUPERCOMPUTER CENTER
SPIRENT
splunk>
STARLIGHT The Optical STAR TAP
SWITCH
uen UTAH EDUCATION NETWORK
THE UNIVERSITY OF WISCONSIN MADISON
UNIVERSITEIT VAN AMSTERDAM
UCSD Jacobs School of Engineering
University of Colorado Boulder
UF UNIVERSITY of FLORIDA The Foundation for The Gator Nation
ILLINOIS
THE UNIVERSITY OF UTAH
THE UNIVERSITY of TENNESSEE KNOXVILLE • CHATTANOOGA • MARTIN • MEMPHIS • TULLAHOMA
YarcData
zayo

SC Silver Anniversary Exhibits

Exhibitor	Booth	Title and Description
Cray Inc.	**1313**	**Original Cray-1 Module:** Cray will display a module from a Cray-1 supercomputer, which reigned as the world's fastest computer from 1976 to 1982.
Numerical Algorithms Group	**1716**	**25 years of Numerical Algorithms Group HPC Expertise and Numerical Libraries:** 25 years of NAG HPC Expertise and Numerical Libraries at SC demonstrates that getting the right answer still matters!
Japan Aerospace Exploration Agency	**1722**	**NWT, Numerical Wind Tunnel:** NWT, the world's first vector parallel computer and three-time Gordon Bell Prize winner. BiCMOS, ECL, and GaAs LSI were used.
University of Illinois	**1820**	**ILLIAC II – From the Dawn of Supercomputing:** With transistors replacing vacuum tubes and magnetic-core member, ILLIAC II (1963) led to smaller, faster, reliable, energy-efficient and commercially-successful computers.
Intel	**2701**	**Intel iPSC/860, Intel Paragon, and ASCI Red:** Intel's massively parallel supercomputers, the iPSC/860 and Paragon, were the basis for ASCI Red, the world's fastest supercomputer from 1997-2000.
IBM	**2713**	**IBM Celebrating 25 Years of Supercomputing and Beyond:** IBM is honored to present computer artifacts and a graphical timeline spanning over 67 years of HPC innovation.
Fujitsu Limited	**2718**	**Fujitsu's Contribution to HPC:** FACOM M-190, world's first fully LSI-based mainframe co-developed with Amdahl and NWT, revolutionary vector parallel supercomputer featuring distributed-memory architecture.
NEC	**3109**	**NEC SX Series Vector Supercomputers —Three Decades of Evolution:** NEC showcases the evolution of the SX Series vector supercomputers with photographs since its early model of the SX-2.
HPCWire	**3510**	**HPCwire - Yesterday, Today, and Tomorrow: HPCwire celebrates SC's Silver Anniversary! Visit #3510** for your copy of the 1st Supercomputing Program Guide, compliments of HPCwire! tci.taborcommunications.com/HPC-wire Reports
Cycle Computing	**3610**	**Cycle Computing Supercomputing History Fold-Out Poster:** Cycle Computing will hand out limited-edition Timeline of HPC History posters, celebrating six eras of computing in photos and words.

SC14
New Orleans, LA | hpc matters.

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

http://sc14.supercomputing.org/

Conference Dates:
Nov. 16 - 21, 2014

Exhibition Dates:
Nov. 17 - 20, 2014

IEEE computer society

Association for Computing Machinery

Follow us on Facebook and Twitter

Call for Participation

Ernest N. Morial Convention Center
New Orleans, LA

SC14: HPC Matters

HPC is helping to solve the hardest problems in the world. Innovations from our community have had a far-reaching impact in every corner of science from the discovery of new drugs to the precise prediction of the next superstorm. HPC is pushing our economy ahead by enabling significant improvements in automobile production, banking, and food manufacturing. For more than two decades, the SC Conference has been the place to build and share the innovations that are making these life-changing discoveries possible.

In November of 2014, SC is going back to New Orleans with new ideas and a fresh take on HPC. Spotlighting the most original and fascinating scientific and technical applications from around the world, SC14 will bring together the HPC community – an unprecedented array of scientists, engineers, researchers, educators, students, programmers, and developers. Together this community will participate in an exceptional program of technical papers, tutorials, timely research posters, and Birds-of-a-Feather (BOF) sessions.

The SC14 Exhibition Hall will feature exhibits of emerging technologies from industry, academia and government research organizations - many of these technologies will be seen for the first time in New Orleans. Mark your calendar and make your way to New Orleans. No city offers the same extraordinary mix of food, music, culture, and history; and no conference offers a better opportunity to experience why ***HPC matters.***

New Orleans, LA