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

SC16

Salt Lake City, Utah | hpc matters.

Conference Program

http://sc16.supercomputing.org/

Salt Lake City Convention Center
Salt Lake City, Utah

Conference: November 13-18, 2016

Exhibition: November 14-17, 2016

Sponsors: IEEE computer society | acm | sighpc

SC16

Salt Lake City, Utah | hpc matters.

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

IEEE computer society

acm

sighpc

Sponsors

SC16 Salt Lake City, Utah | hpc matters.

Table of Contents

SC16 Salt Lake City, Utah | hpc matters.

Welcome to Salt Lake City

Welcome to Salt Lake City for SC16, the international conference for high performance computing, networking, storage and analysis. The city on the shores of the Great Salt Lake offers attendees the best of two worlds: a vibrant metropolis that is one of the fastest growing tech centers in the United States, and the adjacent Wasatch Mountains – a range offering a spectacular backdrop to this glittering jewel.

We're delighted to return to the Salt Palace Convention Center, site of SC12. There's no better place to see why high performance computing matters to our daily lives, our future, our communities and our world. For five days, SC16 will be a global high performance computing hub powered by one of the fastest networks on the planet.

SC provides a unique opportunity to study a complete picture of the state of the art in HPC today and to capture a glimpse into where our community is headed tomorrow. From providing clean water and sanitation to underdeveloped areas of the world and helping coastal cities save lives by better preparing for hurricanes, to developing new drugs for people with early onset Alzheimer's and designing helmets that better protect children and athletes, HPC is powering revolutionary advances that improve our quality of life.

In that vein, keynote speaker Katharine Frase will provide plenty of food for thought in her feature presentation, "Cognitive Computing: How Can We Accelerate Human Decision Making, Creativity and Innovation Using Techniques from Watson and Beyond." http://sc16.supercomputing.org/2016/08/02/sc16-selects-industry-veteran-katharine-frase-keynote-speaker/ Other invited speakers, who also are thought leaders in their fields, will deliver presentations on current topics ranging from the development of new supercomputing capabilities to creating a richer and more diverse HPC workforce.

As in other computing fields, women and other minorities are significantly under-represented in supercomputing, and changing this is vital to ensuring the vitality of the HPC community. Dr. Maria Klawe, president of Harvey Mudd College, will discuss successful strategies for increasing the number of women and students of color in supercomputing in a presentation entitled, "Diversity and Inclusion in Supercomputing." http://sc16.supercomputing.org/2016/09/06/sc16-invited-talk-spotlight-dr-maria-klawe-presents-diversity-inclusion-supercomputing/

SC's own commitment to a diverse workforce and attract young professionals to our field goes back many years, and is anchored by programs such as the Student Cluster Competition, Student Volunteers, the Mentor–Protégé program and the Student Job Fair. Additionally, SC supports members of the HPC community just starting their computing careers through the Early Career Program, which provides an opportunity for special sessions of interest to early career researchers, including getting funding, publishing venues, establishing long-term mentor relationships, and time management.

We are also very excited to add a new facet to these efforts: to better accommodate those with family obligations, SC is offering onsite childcare at the conference for children 6 months to 12 years old. There's no reason SC shouldn't be a family affair! For more information, see http://sc16.supercomputing.org/attendees/on-site-child-care/

The technical program at the heart of the SC is internationally recognized for its breadth and depth, and brings together scientists, engineers, programmers, researchers, system developers and administrators who are leaders in their disciplines. SC's Technical Program is highly competitive and one of the broadest of any HPC conference. SC16 will offer one of the most comprehensive technical programs to date, including presentations, topical panel discussions, papers, tutorials, timely research posters, and Birds-of-a-Feather sessions.

The depth and breadth of topics covered by technical papers has evolved with the conference over the years and selecting those that meet SC's high standard has correspondingly become more difficult. This year, papers cover nine focus areas: Applications; Algorithms; Architectures and Networks; Clouds and Distributed Computing; Data Analytics, Visualization and Storage; Performance Measurement, Modeling and Tools; Programming Systems; State of the Practice; and System Software. Out of 442 technical papers submitted to SC16, only 81 were accepted and of these, seven have been selected as finalists for the conference's Best Paper Award. Reflecting the caliber of talent submitting papers, one of the 'best paper' candidates also is an ACM Gordon Bell Prize finalist.

Popular Technical Program offerings include the panel discussions covering topics of current interest to the HPC community. SC Panels don't just showcase panelists' expertise; they enable participants to engage with key thinkers and producers as more than listeners. Panelists engage in lively, rapid-fire exchanges, often using real-time audience polling, video feeds, and social media to help make panels not only rich with insight, but great fun. The emphasis at SC is on enabling the exchange of ideas through participation and interaction.

Attendees also will have the opportunity to interact with the future of HPC in the Salt Palace's 515,000 sq.ft. exhibition hall, which will feature the latest technologies and accomplishments from the world's leading vendors, research organizations and universities. SC16 will be the first opportunity to learn about many of the technologies that will shape the future of large-scale technical computing and data-driven science.

Innovation is what drives HPC and SC. It is what empowers our community to push the frontiers of science and engineering. A goal of the conference is to serve as a vehicle for advancing the state-of-the-art and broadening the use of HPC by demonstrating how "HPC Matters." SC16 marks the beginning of a multi-year emphasis on advancing our community's 'state-of-the-practice.' A new program track will include all aspects related to the practical use of HPC, including infrastructure, services, facilities, and large-scale application executions. The state-of-the-practice track will include special sessions, papers, panels, workshops, and tutorials that develop and share best practices. Submissions are encouraged from all areas of practice (e.g., system and center determination).

The richness of the conference – with all the opportunities it offers to learn, exchange ideas, and network – can be overwhelming and challenging to navigate, especially to first time attendees. If this is your first SC, welcome! Please take advantage of the resources we've developed especially for you http://sc16.supercomputing.org/attendees/first-time-attendees/. And don't hesitate to stop by one of the information booths or consult our volunteers if you need direction or have questions. We want this to be just the first of many productive years for you at SC.

John West
SC16 General Chair

General Information

Registration and Conference Store

The registration area and conference store are located in the South Foyer, Lower Concourse. They will be open during the following days and hours:

Saturday, November 12	1pm – 6pm
Sunday, November 13	7am – 6pm
Monday, November 14	7am – 9pm
Tuesday, November 15	7:30am – 6pm
Wednesday, November 16	7:30am – 6pm
Thursday, November 17	7:30am – 5pm
Friday, November 18	8am – 11am

Registration Pass Access

See page 7.

Exhibition Hall Hours

Tuesday, November 15	10am-6pm
Wednesday, November 16	10am-6pm
Thursday, November 17	10am-3pm

SC16 Information Booths

Need up-to-the-minute information about what's happening at the conference. Need to know where to find your next session? What restaurants are close by? Where to get a document printed? These questions and more can be answered by a quick stop at one of the SC Information booths. There are two booth locations for your convenience: the Main Booth is on the Lower Concourse, South Foyer, just near registration and the conference store; the Satellite Booth is located on the Upper Concourse Lobby near Room 251.

SC16 Information Booth Hours

Dates	Main Booth	Satellite Booth
Saturday, Nov. 12	1pm-5pm	Closed
Sunday, Nov. 13	8am-6pm	8am-5pm
Monday, Nov. 14	8am-7pm	8am-5pm
Tuesday, Nov. 15	8am-6pm	10am-5pm
Wednesday, Nov. 16	8am-6pm	8am-4pm
Thursday, Nov. 17	8am-6pm	Closed
Friday, Nov. 18	8:30am-12:30pm	Closed

SC17 Preview Booth

Members of next year's SC committee will be available in the SC17 preview booth (located in the South Foyer of the convention center) to offer information and discuss next year's SC conference in Denver, Colorado. 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 15	10am-3pm
Wednesday, November 16	10am-3pm
Thursday, November 17	10am-3pm

Transportation

There are many hotels within walking distance of the Convention Center as well as hotels serviced by the light rail service. Additionally, SC16 will provide four bus loops providing transportation from the farther hotels to the Convention Center. The following table provides information on the transportation recommendation for each of the conference hotels.

Hotel	Transportation Method
Comfort Inn Salt Lake City Airport	SC16 Bus: Route #3
Courtyard by Marriott	SC16 Bus: Route #2
Courtyard Salt Lake City Downtown	Walk
Crystal Inn Salt Lake City Downtown	SC16 Bus: Route #1
Doubletree Suites by Hilton	SC16 Bus: Route #1
Doubletree by Hilton Salt Lake City Airport	SC16 Bus: Route #3
Fairfield Inn & Suites Salt Lake City	SC16 Bus: Route #2
Grand America Hotel & Towers	Light Rail – Courthouse Station
Hampton Inn Downtown	SC16 Bus: Route #1
Hampton Inn & Suites Salt Lake City Airport	SC16 Bus: Route #3
Hilton Salt Lake City Center	Walk
Holiday Inn Salt Lake City – Airport West	SC16 Bus: Route #3
Holiday Inn Express Salt Lake City Downtown	Walk
Homewood Suites	SC16 Bus: Route #2
Hotel Monaco	Walk
Hyatt House Salt Lake City Downtown	Light Rail – Arena Station
Hyatt Place Salt Lake City Airport	SC16 Bus: Route #3
Hyatt Place Salt Lake City Downtown	Light Rail – Arena Station
La Quinta Inn and Suites Salt Lake City Airport	SC16 Bus: Route #3
Little America	Light Rail Courthouse Station
Peery Hotel	Walk
Radisson Hotel Salt Lake City Downtown	Walk
Red Lion Hotel Salt Lake City Downtown	SC16 Bus: Route #1
Residence Inn City Center	SC16 Bus: Route #2
Salt Lake City Marriott Downtown	Walk
Salt Lake City Marriott City Center	Walk
Salt Lake City Marriott University Park	SC16 Bus: Route #4
Salt Lake Plaza Hotel	Walk
Sheraton	Light Rail

SC16 Shuttle Bus Information

Operational Hours:

Saturday, November 12	1pm- 6pm
Sunday, November 13	7am- 9pm
Monday, November 14	7am- 9pm
Tuesday, November 15	7am- 7:30pm
Wednesday, November 16	7am- 7:30pm
Thursday, November 17	7am- 9pm
Friday, November 18	7am- 12:30pm

Frequency:
Sunday 11/13 thru Thursday 11/17:
7:00 am - 11:00 am: Buses run approximately every 15-20 minutes
11:00 am – 4:00 pm: Buses run approximately every 30 minutes

Saturday 11/12 and Friday 11/18:
Buses will run every approximately every 30 minutes

Social Events

Exhibitors' Reception (21+ Event)
Sunday, November 13
6pm – 9pm
SKY SLC
149 Pierpont Ave,
Salt Lake City, UT 84101

SC16 will host the Exhibitors' Appreciation Party at SKY SLC, a brief walk from the Convention Center, for registered exhibitors. The Conference would not be what it is without the generous support and enthusiasm of the exhibitors and we look forward to an evening celebrating these partnerships. This year's reception will have an eclectic mix of activities, from cocktails and live music by the Groove Merchants, to local beer sampling and food pairings, to darts and ping-pong. The reception will offer a little something for everyone and lots of fun for all!

An Exhibitor badge, party ticket, and government-issued photo ID are required and you must be 21 or over to attend this event. Each exhibitor will receive two beverage tickets upon arrival to enjoy libations, followed by a cash bar. Complimentary coat check will be provided at the entry to SKY.

Shuttle transportation from the Convention Center back to hotels along the normal bus routes will be available from 6:30pm - 9:30pm.

Exhibits Gala Opening Reception
Monday, November 14
7pm-9pm

SC16 will host its annual Grand Opening Gala in the Exhibit Hall. This will be your first opportunity to see the latest high performance computing, networking, storage, analysis, and research products, services, and innovations. This event is open to all Technical Program, Exhibitors and Students@SC registrants.

Posters Reception
Tuesday, November 15
5:15pm-7pm
South Foyer & Hall E

The Posters Reception is an opportunity for attendees to interact with poster presenters and includes research and ACM Student Research Competition posters, Doctoral Showcase posters as well as the Scientific Visualization & Data Analytics Showcase. The reception is open to all attendees. Complimentary refreshments and appetizers are available.

Technical Program Conference Reception
Thursday, November 17
6pm-9pm
The Leonardo

The SC16 conference ensures that the science and innovation at the Conference soar to new heights each year. To thank our Technical Program attendees and to encourage continued interactions, SC16 will host a conference reception for all Technical Program attendees at the Leonardo, an interactive museum now featuring a flight exhibit and additional displays geared to stimulate the mind. With a full-sized C-131 aircraft dangling from the rafters on the ground floor, it would be easy to spend your time at the Reception where you entered, especially if you happen to run into Amelia Earhart or the Wright Brothers. Attendees are encouraged to explore the all three floors of the Museum, as we know you will be pleased with the extensive food, exhibit, and entertainment offerings found throughout the space.

A Tech Program badge, event ticket, and government-issued photo ID are required to attend this event. As the event is open to all ages, please be aware that all attendees over 21 years of age who wish to consume alcoholic beverages must check-in and receive a wristband to receive bar service.

Shuttle transportation to and from the event will run starting at 5:45 pm from the Convention Center. Starting at 7:00 pm, shuttles will provide service from the venue to the standard hotel routes as well as service back to the Convention Center.

Conference Services/Convention Center Facilities

ATMs

Two ATM cash machines are located inside the convention center. One is located on the Lower Concourse next to the Parking Garage elevators near the South Lobby. The other is located on the Upper Concourse next to the Business Center and Room 254.

Business Center

The Salt Palace Business Center is an on-site, full service print, copy and shipping center located inside the convention center on the north end of the upper concourse. For more information, contact the Business Center at 385-468-2228 or businesscenter@saltpalace.com.

The business center will be open from 8:00 a.m. – 5:00 p.m. from Saturday, November 12th through Friday, November 18th.

City and Dining Information

Stop by the Visit Salt Lake kiosk in the South Foyer of the convention center to discover dining and city options.

Sunday, November 12	10am-6pm
Monday, November 13	10am-7pm
Tuesday, November 14	10am-6pm
Wednesday, November 15	10am-6pm
Thursday, November 16	10am-6pm

Coat & Bag Check

There are two coat and bag check stations. One is located near the South Entrance in Exhibit Hall "E" and the other is just inside the East Entrance. The hours for both locations are as follows:

Saturday, November 12	Noon-8:00pm
Sunday, November 13	7am-6pm
Monday, November 14	7am-10pm
Tuesday, November 15	7am-9pm
Wednesday, November 16	7am-8pm
Thursday, November 17	7am-6pm
Friday, November 18	7am-1pm

Emergency Contact

For an onsite emergency, please dial 3040 on any house phone.

Family Day

Family Day is Wednesday, November 16, 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

There are two first aid offices in the convention center. One is on the east side, located next to Room 150A; the other is on the west side lobby, outside of Hall 4.

Lost Badge

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

Lost & Found

Lost and found is located in Room 258.

On-Site Child Care

Room 257

On-site childcare services in the convention center is being provided by KiddieCorp for children between ages 6 months through 12 years old. The deadline for sign-up was Friday, October 14, 2016. No drop-in childcare is available. All children using the SC16 childcare services must be registered in advance by a parent or legal guardian. Contact KiddieCorp at 858-455-1718.

Hours:

Sunday, November 13th	8:00am-6:00pm
Monday, November 14th	8:00am-9:00pm
Tuesday, November 15th	8:00am-7:00pm
Wednesday, November 16th	8:00am-7:00pm
Thursday, November 17th	8:00am-6:00pm

Parents with Infants Room

SC16 has provided a lactation room with privacy areas for nursing mothers, located in Room 252A. Each private space will have a plush chair and small table. A sofa with end tables, refrigerator, changing tables, lockers, and diaper disposal will be provided in the common space. The room will be open during the following dates and hours:

Saturday, November 12	noon – 5pm
Sunday, November 13	8am – 6pm
Monday, November 14	8am – 9pm
Tuesday, November 15	8am – 9pm
Wednesday, November 16	8am – 7pm
Thursday, November 17	8am – 6pm
Friday, November 18	8am – 1pm

Parking

The Salt Palace Convention Center offers two convenient underground garages that are open 7 days a week:

- South Lot – Enter off of 200 South between West Temple and 200 West (600 stalls)
- West Lot – Enter off of 300 West between 100 South and South Temple (400 stalls)
- Handicap parking available in both lots
- Sorry, no in/out privileges or overnight parking; height/length restrictions
- Bicycles racks are located in the garages and in various places outside the facility
- Continuous security surveillance

Prayer & Meditation Room

The prayer and meditation room is located in Room 260-B and is open Sunday-Thursday, 9am-5pm.

Restrooms

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

Lower Level:

- Halls A-E (located in the back)
- Hall 1
- Halls 4&5 (west side)
- North and South Foyers
- Outside Room 155

Upper level:

- Across from 254B and near 255
- Upper Mezzanine (on left-hand side)

Visitor's Center

The Visitor's Center is located near the East entrance. It is open daily from 9am-5pm.

Wheelchair/Scooter Rental

Wheelchairs and medical mobility scooters can be rented from the Salt Palace Business Center. Equipment is located on-site with hourly, daily and event rates.

Registration Pass Access

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

Type of Registration	Technical Program	Technical Program + Workshops	Tutorials	Workshop Only	Exhibitor 24-hour Access	Exhibit Hall Only Tue-Thur
Awards (Thursday/12:45)	*	*			*	*
Birds-of-a-Feather	*	*			*	*
Conference Reception (Thursday)	*	*				
Doctoral Showcase	*	*		*	*	*
Emerging Technologies	*	*	*	*	*	*
Exhibit Floor	*	*			*	*
Exhibitor Forum	*	*			*	*
Exhibits Gala Opening (Monday)	*	*			*	
Exhibitors Reception					*	
HPC Impact Showcase	*	*	*	*	*	*
HPC Matters Plenary	*	*	*	*	*	*
Invited Talks (Non-Plenary)	*	*				
Invited Talks (Plenary)	*	*			*	
Keynote (Tuesday)	*	*	*	*	*	*
Panels (Tue-Thur)	*	*				
Panels (Friday Only)	*	*			*	
Papers	*	*				
Posters	*	*	*	*	*	*
Poster Reception (Tuesday)	*	*				
Scientific Visualization Showcase	*	*	*	*	*	*
Tutorial Lunch (Sun/Mon ONLY)			*			
Tutorial Sessions			*			
Student Cluster Competition	*	*			*	*
Workshops		*		*		

SCinet

The Fastest Network Connecting the Fastest Computers

The Fastest Network Connecting the Fastest Computers
SC16 is once again hosting one of the most powerful and advanced networks in the world – SCinet. Created each year for the conference, SCinet brings to life a very high-capacity network that supports the revolutionary applications and experiments that are a hallmark of the SC conference. SCinet will link the convention center to research and commercial networks around the world. In doing so, SCinet serves as the platform for exhibitors to demonstrate the advanced computing resources of their home institutions and elsewhere by supporting a wide variety of bandwidth-driven applications including supercomputing and cloud computing.

Volunteers from academia, government and industry work together to design and deliver the SCinet infrastructure. Industry vendors, service providers and carriers donate millions of dollars in equipment and services needed to build the local and wide area networks. Planning begins more than a year in advance of each SC conference and culminates in a high-intensity installation in the days leading up to the conference.
For SC16, SCinet is exploring up-and-coming topics in the high-performance networking community of SC through the Network Research Exhibition (NRE) and returning for a third year, the INDIS workshop.

In addition, SCinet is continuing its Diversity Program, Women in IT Networking at SC (WINS). It is funded by the NSF and administered by a collaboration between the University Corporation for Atmospheric Research (UCAR), the Department of Energy's Energy Sciences Network (ESnet) and the Keystone Initiative for Network Based Education and Research (KINBER). The ongoing activities center around funding U.S. women in their early to mid-careers to participate in SCinet to get hands-on training and build their professional network. Three alumna of the program are returning to continue their participation, and help mentor the new WINS participants.

SCinet Network Research Exhibition (NRE)

The NRE is SCinet's forum for displaying new or innovative demonstrations in network testbeds, emerging network hardware, protocols, and advanced network intensive scientific applications that will stretch the boundaries of a high-performance network. This year showcases SCinet's inclusion of Optical Transport Network (OTN) components and Data Center Interconnects (DCI) technologies in the showfloor infrastructure which deliver services to our exhibitors. These technologies allow for increased densities and advanced diagnostic capabilities.

SCinet INDIS Workshop

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

SCinet Collaborators

SCinet is the result of the hard work and significant contributions of many government, research, education and corporate contributors. Contributors for SC16 include:

Platinum

CenturyLink Government

ciena

Cisco

Coriant

CORSA

ESnet ENERGY SCIENCES NETWORK

infinera

Internet2

Juniper Networks

zayo GROUP

Gold

ARISTA

BROCADE

CloudLab

UEN UTAH EDUCATION NETWORK UEN.ORG

SCinet Collaborators

Silver

SCinet Collaborators

Bronze

SC16

Salt Lake City, Utah | hpc matters.

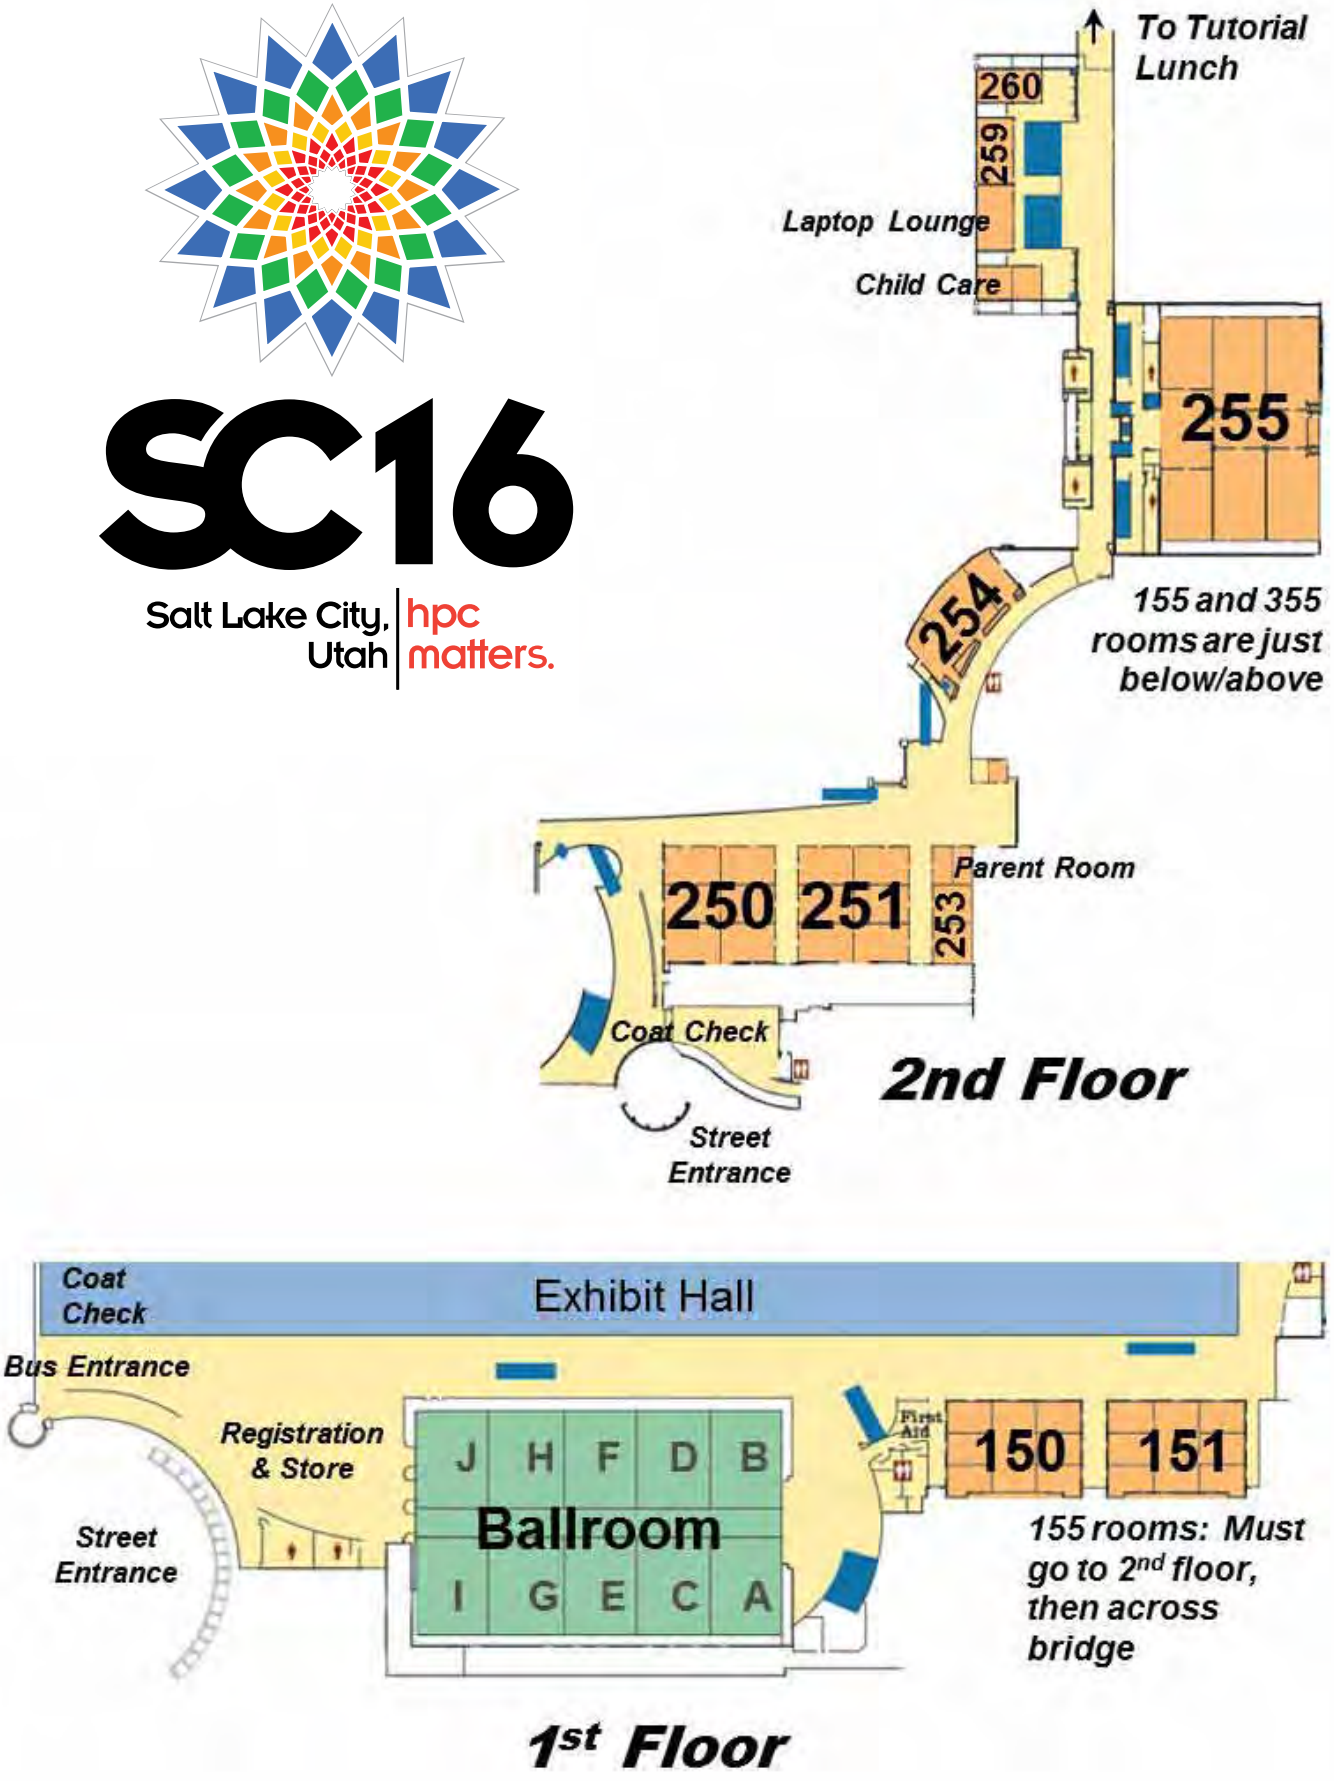

Salt Lake Downtown

Visit Salt Lake
90 South West Temple, Salt Lake City, Utah
801-534-4900 info@visitsaltlake.com VisitSaltLake.com

RESTAURANTS & BARS

1	Bambara	801-363-5454
2	Bayleaf Bar & Grub	801-359-8490
3	The Bayou	801-961-8400
4	Beehive Tea Room & Wedding Library	801-328-4700
5	Beerhive Pub	801-364-4268
6	Benihana	801-322-2421
7	Biaggi's Ristorante Italiano	801-596-7222
8	Blue Iguana Restaurant	801-533-8900
9	Blue Lemon	801-328-2583
10	Buca di Beppo	801-575-5262
11	Cafe SuperNatural	801-363-1000
12	Cafe Trang	801-359-1638
13	Caffe Molise	801-364-8833
14	California Pizza Kitchen	801-456-0075
15	Canyon Creek Cafes Food Court	
16	Cedars of Lebanon	801-364-4096
17	The Cheesecake Factory	801-532-4706
18	Christopher's Seafood & Steakhouse	801-519-8515
19	City Creek Center Food Court	
20	Clouseau's	801-359-7800
21	Club Piastra	801-961-8700
22	Copper Bowl Indian Cuisine	801-309-7543
23	Copper Canyon Grill House & Tavern	801-521-7800
24	Copper Onion	801-355-3282
25	Cucina Toscana	801-328-3463
26	The Depot	801-456-2800
27	Desert Edge Brewery at the Pub	801-521-8917
28	Destinations	801-531-0800
29	Elevations Restaurant	801-537-5019
30	Eva	801-359-8447
31	Faustina	801-746-4441
32	5th Street Grill	801-323-7575
33	First Press	801-401-2000
34	Fleming's Prime Steakhouse & Wine Bar	801-355-3704
35	Frida Bistro	801-983-6692
36	The Garden Cafe	801-258-5708
37	The Garden Restaurant	801-539-3170
38	Ginza Japanese Cuisine & Sushi Bar	801-322-2224
39	Gracie's	801-819-7565
40	Grilla Bites	801-456-2425
41	Happy Sumo At Gateway	801-456-7866
42	Himalayan Kitchen	801-328-2077
43	The Hotel Bar & Nightclub	801-487-4310
44	Ichiban Sushi & Japanese Cuisine	801-532-7522
45	Iggy's Sports Grill	801-532-9999
46	Inferno Cantina	801-883-8838
47	JB's Family Restaurant	801-328-8344
48	J. Wong's Asian Bistro	801-350-0888
49	Keys On Main	801-363-3638
50	Kneader's Bakery and Cafe	801-428-3051
51	Kristauf's Martini Bar	801-366-9490
52	Lamb's Grill Cafe	801-364-7166
53	Legends Pub & Grill	801-355-3598
54	The Lion House Pantry Restaurant*	801-363-5466
55	Little America Coffee Shop	801-596-5700
56	Little America Steak House	801-596-5700
57	Lumpys Downtown	801-938-3070
58	Market Street Grill - Downtown	801-322-4668
59	Market Street Oyster Bar - Downtown	801-531-6044
60	Martine	801-363-9328
61	Maxwell's East Coast Eatery	801-328-0304
62	McDonald's	801-364-1614
63	The Melting Pot	801-521-6358
64	Murphy's Bar and Grill	801-359-7271
65	Naked Fish Japanese Bistro	801-595-8888
66	Nauvoo Cafe	801-539-3346
67	New Yorker	801-363-0166
68	Oasis Cafe	801-322-0404
69	Pallet Bistro	801-935-4431
70	P. F. Chang's China Bistro	801-539-0500
71	Piastra on Gallivan Plaza	801-961-8700
72	The Plum Restaurant	801-359-7800
73	Poplar Street Pub	801-532-2715
74	Raw Bean Coffee House	801-990-2326
75	The Red Door	801-363-6030
76	Red Rock Brewing Company L.C.	801-521-7446
77	Red's Café	801-521-7373
78	Rio Grande Cafe	801-364-3302
79	Rodizio Grill	801-220-0500
80	Romano's Macaroni Grill	801-521-3133
81	The Roof Restaurant	801-539-1911
82	Royal Eatery	801-532-4301
83	Ruth's Chris Steak House	801-363-2000
84	Sage's Cafe	801-322-3750
85	Salt Lake Roasting Company	801-363-7572
86	Settebello Pizzeria Napoletana	801-322-3556
87	Siegfried's Delicatessen, Inc.	801-355-3891
88	Spencer's For Steaks & Chops	801-238-4748
89	Squatters Pub Brewery	801-363-2739
90	Star of India Restaurant	801-363-7555
91	Takashi	801-519-9595
92	Tavernacle Social Club	801-519-8900
93	Texas de Brazil Churrascaria	385-232-8070
94	The Tin Angel Cafe	801-328-4155
95	Toaster's	801-328-2928
96	Trofi	801-238-4877
97	Tucanos Brazilian Grill	801-456-2550
98	Tucci's Cucina Italiana	801-533-9111
99	Vinto	801-539-9999
100	Vosen's Bread Paradise	801-322-2424
101	Z'Tejas Southwestern Grill	801-456-0450
102	ZY Restaurant	801-779-4730

ATTRACTIONS

1	Beehive House	801-240-2681
2	Brewvies Cinema Pub	801-355-5500
3	Brigham Young Historic Park	
4	Brigham Young Monument	
5	Broadway Centre Cinemas	801-321-0310
6	Capitol Theatre	801-355-2787
7	Cathedral Church of St. Mark	801-322-3400
8	Cathedral of the Madeleine	801-328-8941
9	City and County Building	801-533-0858
10	City Creek Center – *Shopping Area*	801-238-5320
11	City Creek Park	
12	**Clark Planetarium** CP	**801-456-7827**
13	Council Hall	801-538-1900
14	**Discovery Gateway** CP	**801-456-5437**
15	Eagle Gate	
16	EnergySolutions Arena	801-325-2000
17	Family History Library	801-240-2584
18	First Presbyterian Church	801-363-3889
19	Gallivan Center	801-535-6110
20	The Gateway – *Shopping Area*	801-456-0000
21	Historic Trolley Square – *Shopping Area*	801-521-9877
22	Holy Trinity Cathedral Greek Orthodox Church	801-328-9681
23	Joseph Smith Memorial Building, FamilySearch™ Center	801-240-1266
24	LDS Church Office Building	801-240-2190
25	LDS Conference Center	801-240-0075
26	**The Leonardo** CP	**801-531-9800**
27	Library Square	801-524-8200
28	**Lion House** CP	**801-363-5466**
29	Maurice Abravanel Hall	801-355-2787
30	Megaplex 12 at the Gateway	801-304-4553
31	Memory Grove Park	801-972-7800
32	Mormon Pioneer Memorial Monument	
33	Museum of Church History & Art	801-240-3310
34	The Off Broadway Theatre	801-355-4628
35	Olympic Legacy Plaza	
36	Pioneer Memorial Museum	801-532-6479
37	Rio Grande Depot Utah State Historical Society	801-533-3500
38	Rose Wagner Performing Arts Center	801-323-6800
39	Salt Lake Temple	
40	Salt Palace Convention Center	801-534-4777
41	Simply Salt Lake Gift Shop	801-534-4906
42	Social Hall Heritage Exhibit	801-321-8745
43	Tabernacle	801-240-4872
44	Union Pacific Depot	
45	Utah Museum of Contemporary Art	801-328-4201
46	Utah State Capitol	801-538-3074
47	Visitor Information Center	801-534-4900

CP *Visit Salt Lake Connect Pass attractions*

SATURDAY, NOVEMBER 12

Event	Tine	Session Title	Room
Child Care	12:00pm-5:00pm	Room for parents with infants	252-A
Coat/Bag Check 1	12:00pm-8:00pm	Coat and bag check	East Entrance
Coat/Bag Check 2	12:00pm-8:00pm	Coat and bag check	Hall E
Registration	1:00pm-6:00pm	Attendee Registration	South Foyer
Store	1:00pm-6:00pm	Conference Store	South Foyer
Students@SC	5:30pm-7:30pm	Building Professional Networks Workshop	Exhibit Hall 5

SUNDAY, NOVEMBER 13

Event	Tine	Session Title	Room
Break	10:00am-10:30am	Refreshments (Tutorials/Workshops)	255-Pre+255-A
Break	10:00am-10:30am	Refreshments (Tutorials/Workshops)	355-Pre+355-A
Break	10:00am-10:30am	Refreshments (Tutorials/Workshops)	Hallway 250/251
Break	3:00pm-3:30pm	Refreshments (Tutorials/Workshops)	255-Pre+255-A
Break	3:00pm-3:30pm	Refreshments (Tutorials/Workshops)	355-Pre+355-A
Break	3:00pm-3:30pm	Refreshments (Tutorials/Workshops)	Hallway 250/251
Child Care	8:00am-6:00pm	Child care	257
Child Care	8:00am-6:00pm	Room for parents with infants	252-A
Coat/Bag Check 1	7:00am-6:00pm	Coat and bag check	East Entrance
Coat/Bag Check 2	7:00am-6:00pm	Coat and bag check	Hall E
Media/Press	1:00pm-5:00pm	Media Room and Press Office	151-DE
Press Interview Room	1:00pm-5:00pm	Interview Room for Press (reservations only)	151-F
Reception	6:00pm-9:00pm	Exhibitors' Reception	SKY SLC
Registration	7:00am-6:00pm	Attendee Registration	South Foyer
Store	7:00am-6:00pm	Conference Store	South Foyer
Students@SC	8:30am-10:00am	Diversity & Inclusion: Views from the Field	260
Students@SC	10:30am-12:00pm	Panel: Coming out on Behalf of Diversity	260
Students@SC	1:30pm-2:15pm	Education/Career Keynote	260
Students@SC	2:15pm-3:00pm	Pitch-It Workshop: Learn to Sell Yourself and your Research	260
Students@SC	3:30pm-4:15pm	Mentoring: Learning from the Pros	260
Students@SC	4:15pm-5:00pm	Life on the Other Side: Early Career Panel	260
Tutorial	8:30am-12:00pm	HPC Acquisition and Commissioning	355-F
Tutorial	8:30am-12:00pm	Solving and Sharing the Puzzle: Modeling and Simulation of Computer Architectures with SST and OCCAM	250-F
Tutorial	8:30am-12:00pm	Scalable HPC Visualization and Data Analysis Using VisIt	255-F

SUNDAY, NOVEMBER 13

Event	Tine	Session Title	Room
Tutorial	8:30am-12:00pm	Power-Aware High Performance Computing: Challenges and Opportunities for Application and System Developers	255-E
Tutorial	8:30am-5:00pm	Efficient Parallel Debugging for MPI, Threads, and Beyond	250-D
Tutorial	8:30am-5:00pm	Parallel Programming in Modern Fortran	250-B
Tutorial	8:30am-5:00pm	Productive Programming in Chapel: A Computation-Driven Introduction	250-A
Tutorial	8:30am-5:00pm	Programming Intel's 2nd Generation Xeon Phi (Knights Landing)	255-B
Tutorial	8:30am-5:00pm	Parallel I/O In Practice	355-B
Tutorial	8:30am-5:00pm	Building the Modern Research Data Portal Using the Globus Platform	250-C
Tutorial	8:30am-5:00pm	High Performance Python for Scientific Computing	355-E
Tutorial	8:30am-5:00pm	Kokkos: Enabling Manycore Performance Portability for C++ Applications	250-E
Tutorial	8:30am-5:00pm	Node-Level Performance Engineering	255-C
Tutorial	8:30am-5:00pm	Parallel Computing 101	255-D
Tutorial Lunch	12:00pm-1:30pm	Tutorial Attendees' Lunch	Exhibit Hall 5
Tutorial	1:30pm-5:00pm	Insightful Automatic Performance Modeling	250-F
Tutorial	1:30pm-5:00pm	Essential HPC Finance Practice: Total Cost of Ownership (TCO), Internal Funding, and Cost-Recovery Models	355-F
Tutorial	1:30pm-5:00pm	Managing HPC Software Complexity with Spack	255-F
Tutorial	1:30pm-5:00pm	Big Data Meets HPC: Exploiting HPC Technologies for Technologies for Accelerating Apache Hadoop, Spark, and Memcached	355-C
Workshop	9:00am-12:30pm	WOLFHPC: Sixth International Workshop on Domain-Specific Languages and High-Level Frameworks for HPC	251-F
Workshop	9:00am-12:30pm	3rd International Workshop on HPC User Support Tools (HUST-16)	155-E
Workshop	9:00am-12:30pm	HPGDMP'16: First International Workshop on High Performance Graph Data Management and Processing	251-C
Workshop	9:00am-5:30pm	IA^3 2016 - 6th Workshop on Irregular Applications: Architectures and Algorithms	251-D
Workshop	9:00am-5:30pm	ISAV 2016: In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization	251-A
Workshop	9:00am-5:30pm	Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems	251-E
Workshop	9:00am-5:30pm	ExaMPI16	155-A
Workshop	9:00am-5:30pm	Innovating the Network for Data Intensive Science (INDIS 2016)	155-F
Workshop	9:00am-5:30pm	Computational Approaches for Cancer	251-B
Workshop	9:00am-5:30pm	7th International Workshop on Performance Modeling, Benchmarking & Simulation of HPC Systems (PMBS16)	155-B

SUNDAY, NOVEMBER 13

Event	Tine	Session Title	Room
Workshop	9:00am-5:30pm	Seventh Annual Workshop for the Energy Efficient HPC Working Group (EE HPC WG)	155-C
Workshop	9:00am-5:30pm	Women in HPC: Diversifying the HPC Community	355-D
Workshop	2:00pm-5:30pm	Workshop on Extreme-Scale Programming Tools	155-E
Workshop	2:00pm-5:30pm	The Fourth International Workshop on Software Engineering for HPC in Computational Science and Engineering	251-C

MONDAY, NOVEMBER 14

Event	Tine	Session Title	Room
Break	10:00am-10:30am	Refreshments (Tutorials/Workshops)	255-Pre+255-A
Break	10:00am-10:30am	Refreshments (Tutorials/Workshops)	355-Pre+355-A
Break	10:00am-10:30am	Refreshments (Tutorials/Workshops)	Hallway 250/251
Break	3:00pm-3:30pm	Refreshments (Tutorials/Workshops)	255-Pre+255-A
Break	3:00pm-3:30pm	Refreshments (Tutorials/Workshops)	355-Pre+355-A
Break	3:00pm-3:30pm	Refreshments (Tutorials/Workshops)	Hallway 250/251
Child Care	8:00am-9:00pm	Child care	257
Child Care	8:00am-9:00pm	Room for parents with infants	252-A
Coat/Bag Check 1	7:00am-10:00pm	Coat and bag check	East Entrance
Coat/Bag Check 2	7:00am-10:00pm	Coat and bag check	Hall E
Early Career Program	8:30am-10:00am	Introduction & Peer Networking/ Time Management	251-D
Early Career Program	10:00am-10:30am	Morning Break	Hallway 250/251
Early Career Program	10:30am-12:00pm	Planning Your Career/Career Networking	251-D
Early Career Program	12:00pm-1:30pm	Lunch (program participants only)	Exhibit Hall 5
Early Career Program	1:30pm-3:00pm	Funding Your Research/ Publications Venues & the Review Process	251-D
Early Career Program	3:00pm-3:30pm	Afternoon Break	Hallway 250/251
Early Career Program	3:30pm-5:00pm	Speed Mentoring	251-D
Inclusivity	4:30pm-5:15pm	Introduction to SC16 for First Time Attendees	260
Media/Press	9:00am-5:00pm	Media Room and Press Office	151-DE
Plenary	5:45pm-6:45pm	HPC Matters Plenary Panel Discussion: HPC and Precision Medicine	Ballroom
Press Briefing	12:00pm-1:30pm	Top500	151-F
Press Interview Room	9:00am-5:00pm	Interview Room for Press (reservations only)	151-F
Reception	7:00pm-9:00pm	Gala Opening Reception	Exhibit Halls A-E, 1-5
Registration	7:00am-9:00pm	Attendee Registration	South Foyer
Store	7:00am-9:00pm	Conference Store	South Foyer
Students@SC	8:30am-10:00am	Finding Your Way: The Possibilities, Pitfalls and Practice of HPC Research	260

MONDAY, NOVEMBER 14

Event	Tine	Session Title	Room
Students@SC	10:30am-12:00pm	How To: Getting Started in the HPC Community	260
Students@SC	1:30pm-3:00pm	Navigating SC16; Getting Involved with SC17	260
Students@SC	2:00pm-3:00pm	Experiencing HPC for Undergrades Reception	151-B
Students@SC	3:00pm-5:00pm	Experiencing HPC for Undergrades Orientation	151-B
Students@SC	3:30pm-5:00pm	Mentor-Protégé	Exhibit Hall 5
Students@SC	7:30pm-9:00pm	Student Cluster Competition Kickoff	Exhibit Hall 1
Tutorial	8:30am-12:00pm	Harnessing the Power of FPGAs with Altera's SDK for OpenCL	355-F
Tutorial	8:30am-12:00pm	Container Computing for HPC and Scientific Workflows	355-B
Tutorial	8:30am-12:00pm	Programming Your GPU with OpenMP: A Hands-On Introduction	255-B
Tutorial	8:30am-12:00pm	Secure Coding Practices and Automated Assessment Tools	250-C
Tutorial	8:30am-12:00pm	InfiniBand and High-Speed Ethernet for Dummies	355-C
Tutorial	8:30am-5:00pm	Advanced MPI Programming	355-E
Tutorial	8:30am-5:00pm	Hands-On Practical Hybrid Parallel Application Performance Engineering	250-B
Tutorial	8:30am-5:00pm	Large Scale Visualization with ParaView	250-D
Tutorial	8:30am-5:00pm	Debugging and Performance Analysis on Native and Offload HPC Architectures	250-A
Tutorial	8:30am-5:00pm	Application Porting and Optimization on GPU-Accelerated POWER Architectures	255-F
Tutorial	8:30am-5:00pm	How to Analyze the Performance of Parallel Codes 101	255-E
Tutorial	8:30am-5:00pm	Linear Algebra Libraries for HPC: Scientific Computing with Multicore and Accelerators	250-E
Tutorial	8:30am-5:00pm	Advanced OpenMP: Performance and 4.5 Features	255-C
Tutorial	8:30am-5:00pm	Fault-Tolerance for HPC: Theory and Practice	250-F
Tutorial	12:00pm-1:30pm	Tutorial Attendees' Lunch	Exhibit Hall 5
Tutorial	1:30pm-5:00pm	InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage	355-C
Tutorial	1:30pm-5:00pm	Data Center Design and Planning for HPC Folks	355-F
Tutorial	1:30pm-5:00pm	Testing of HPC Scientific Software	255-B
Tutorial	1:30pm-5:00pm	Vectorization Strategies for Intel's 2nd Generation Intel Xeon® Phi™ Architecture Named Knights Landing	355-B
Tutorial	1:30pm-5:00pm	Programming Irregular Applications with OpenMP: A Hands-On Introduction	250-C
Workshop	9:00am-12:30pm	MTAGS16: 9th Workshop on Many-Task Computing on Clouds, Grids, and Supercomputers	251-B
Workshop	9:00am-12:30pm	The Seventh International Workshop on Data-Intensive Computing in the Cloud	251-E
Workshop	9:00am-12:30pm	Second International Workshop on Heterogeneous Computing with Reconfigurable Logic	251-C
Workshop	9:00am-12:30pm	Machine Learning in HPC Environments	355-D

MONDAY, NOVEMBER 14

Event	Time	Session Title	Room
Workshop	9:00am-12:30pm	Third SC Workshop on Best Practices for HPC Training	155-F
Workshop	9:00am-5:30pm	Energy Efficient Supercomputing (E2SC)	255-D
Workshop	9:00am-5:30pm	Taking Supercomputing to the Clinic: Medical Image Analysis and Visualization	251-A
Workshop	9:00am-5:30pm	Joint International Workshop on Parallel Data Storage and Data Intensive Scalable Computing Systems (PDSW-DISCS)	155-C
Workshop	9:00am-5:30pm	7th SC Workshop on Big Data Analytics: Challenges and Opportunities	155-A
Workshop	9:00am-5:30pm	LLVM-HPC2016: Third Workshop on the LLVM Compiler Infrastructure in HPC	251-F
Workshop	9:00am-5:30pm	The 1st International Workshop on Post-Moore Era Supercomputing (PMES)	155-B
Workshop	9:00am-5:30pm	PyHPC 2016: 6th Workshop on Python for High-Performance and Scientific Computing	155-E
Workshop	2:00pm-5:30pm	HPC Systems Professional Workshop	155-F
Workshop	2:00pm-5:30pm	Third International Workshop on Accelerator Programming Using Directives (WACCPD)	251-C
Workshop	2:00pm-5:30pm	Workshop on Education for High Performance Computing (EduHPC)	251-E
Workshop	2:00pm-5:30pm	The 11th Workshop on Workflows in Support of Large-Scale Science	355-D
Workshop	2:00pm-5:30pm	PAW: PGAS Applications Workshop	251-B

TUESDAY, NOVEMBER 15

Event	Tine	Session Title	Room
Break	10:00am-10:30am	Refreshments (Tech Program)	255-Pre+255-A
Break	10:00am-10:30am	Refreshments (Tech Program)	355-Pre+355-A
Break	3:00pm-3:30pm	Refreshments (Tech Program)	255-Pre+255-A
Break	3:00pm-3:30pm	Refreshments (Tech Program)	355-Pre+355-A
Birds of a Feather	10:30am-12:00pm	Meeting of the SIGHPC - Big Data Chapter	155-A
Birds of a Feather	12:00pm-1:30pm	13th Graph500 List	155-A
Birds of a Feather	12:15pm-1:15pm	SAGE2: Scalable Amplified Group Environment for Global Collaboration	250-E
Birds of a Feather	12:15pm-1:15pm	Burst Buffers: Early Experiences and Outlook	255-D
Birds of a Feather	12:15pm-1:15pm	Next Generation of Co-Processors Emerges: In-Network Computing	255-EF
Birds of a Feather	12:15pm-1:15pm	The 2016 HPC Challenge Awards	155-C
Birds of a Feather	12:15pm-1:15pm	Lustre Deployments for the Next 5 Years	255-BC
Birds of a Feather	12:15pm-1:15pm	HPC Education: Meeting of the SIGHPC Education Chapter	355-D

TUESDAY, NOVEMBER 15

Event	Tine	Session Title	Room
Birds of a Feather	12:15pm-1:15pm	New Technologies, Platforms and Services for Sharing Research Data	250-D
Birds of a Feather	12:15pm-1:15pm	On-Demand Infrastructure for Data Analytics and Storage	250-C
Birds of a Feather	12:15pm-1:15pm	SIGHPC Annual Member Meeting	355-F
Birds of a Feather	1:30pm-3:00pm	OpenStack for HPC: Best Practices for Optimizing Software-Defined Infrastructure	155-A
Birds of a Feather	5:15pm-7:00pm	Emerging Trends in HPC Systems and Application Modernization	255-D
Birds of a Feather	5:15pm-7:00pm	OpenMP: Where Is It Now and Where Is It Going?	355-E
Birds of a Feather	5:15pm-7:00pm	Reconfigurable Supercomputing	250-F
Birds of a Feather	5:15pm-7:00pm	MPICH: A High-Performance Open-Source MPI Implementation	250-E
Birds of a Feather	3:30pm-5:00pm	The Future of NSF Advanced Cyberinfrastructure	155-A
Birds of a Feather	5:15pm-7:00pm	TOP500 Supercomputers	255-EF
Birds of a Feather	5:15pm-7:00pm	Energy Efficiency Considerations and HPC Procurement	255-BC
Birds of a Feather	5:15pm-7:00pm	The 2016 Ethernet Roadmap	155-C
Birds of a Feather	5:15pm-7:00pm	Distributed Machine Intelligence Using Tensorflow	155-E
Birds of a Feather	5:15pm-7:00pm	How to Build Diverse Teams for More Effective Research	250-C
Birds of a Feather	5:15pm-7:00pm	Today's Hot Technology: The Growing Necessity of Liquid Cooled HPC	355-D
Birds of a Feather	5:15pm-7:00pm	Experimental Infrastructure and Methodology for HPC Cloud Research	250-D
Child Care	8:00am-7:00pm	Child care	257
Child Care	8:00am-9:00pm	Room for parents with infants	252-A
Coat/Bag Check 1	7:00am-9:00pm	Coat and bag check	East Entrance
Coat/Bag Check 2	7:00am-9:00pm	Coat and bag check	Hall E
Doctoral Showcase	10:30am-12:00pm	Doctoral Showcase 1	155-C
Doctoral Showcase	1:30pm-3:00pm	Doctoral Showcase 2	155-C
Doctoral Showcase	3:30pm-5:00pm	Doctoral Showcase 3	155-C
Doctoral Showcase	5:15pm-7:00pm	Poster Presentations	Exhibit Hall E
Emerging Technologies	10:00am-6:00pm	Emerging Technologies Exhibits	155-B
Exhibitor Forum	10:30am-12:00pm	Tools	155-E
Exhibitor Forum	10:30am-12:00pm	Burst Buffers & GPUs	155-F
Exhibitor Forum	3:30pm-5:00pm	Networking	155-F
Exhibitor Forum	3:30pm-5:00pm	Workflow	155-E
Exhibits	10:00am-6:00pm	Exhibits	Exhibit Halls A-E, 1-5
HPC Impact Showcase	1:00pm-3:30pm	HPC Impact Showcase	155-E

TUESDAY, NOVEMBER 15

Event	Time	Session Title	Room
Invited Talks	10:30am-12:00pm	• Understanding Cities through Computation, Data Analytics and Measurement Charles Catlett (Argonne National Laboratory) • The Materials Project—A Google of Materials Kristin Persson (Lawrence Berkeley National Laboratory)	Ballroom
Invited Talks	1:30pm-2:30pm	Test of Time Award Special Lecture: Automatically Tuned Linear Algebra Software Clint Whaley, Jack Dongarra (University of Tennessee)	Ballroom
Keynote	8:30am-10:00am	Dr. Katharine Frase Cognitive Computing: How Can We Accelerate Human Decision Making, Creativity and Innovation Using Techniques from Watson and Beyond?	Ballroom
Media/Press	8:00am-5:00pm	Media Room and Press Office	151-DE
Panel	10:30am-12:00pm	Different Architectures, Different Times: Reproducibility and Repeatability in HPC	255-BC
Panel	1:30pm-3:00pm	National Strategic Computing Initiative Update	255-BC
Panel	3:30pm-5:00pm	HPC Workforce Development: How Do We Find Them, Recruit Them, and Teach Them to Be Today's Practitioners and Tomorrow's Leaders?	255-BC
Papers	10:30am-12:00pm	Systems and Networks I	355-BC
Papers	10:30am-12:00pm	Molecular Dynamics Simulation	355-D
Papers	10:30am-12:00pm	State-of-the-Practice: Advanced Applications Development	255-EF
Papers	1:30pm-3:00pm	Numerical Algorithms I	355-E
Papers	1:30pm-3:00pm	Resilience and Error Handling	355-BC
Papers	1:30pm-3:00pm	Scientific Data Management and Visualization	355-D
Papers	3:30pm-4:30pm	Topics in Distributed Computing	355-E
Papers	3:30pm-5:00pm	Resilience	355-D
Papers	3:30pm-5:00pm	Tensor and Graph Algorithms	355-BC
Posters	8:30am-5:00pm	Research Posters Exhibit	Lower Lobby
Posters	8:30am-5:00pm	Scientific Visualization and Data Analytics Showcase Posters	South Foyer
Press Interview Room	8:00am-5:00pm	Interview Room for Press (reservations only)	151-F
Reception	5:15pm-7:00pm	Research Posters Reception	South Foyer
Reception	5:15pm-7:00pm	ACM Student Research Competition Posters Reception	Exhibit Hall E Booth #104
Reception	5:15pm-7:00pm	Scientific Visualization and Data Analytics Showcase Posters Reception	South Foyer
Reception	5:15pm-7:00pm	Doctoral Showcase Posters Reception	Exhibit Hall E Booth #122

TUESDAY, NOVEMBER 15

Event	Time	Session Title	Room
Registration	7:30am-6:00pm	Attendee Registration	South Foyer
Store	7:30am-6:00pm	Conference Store	South Foyer
Student Research	10:00am-6:00pm	ACM Student Research Competition Posters Competition	Exhibit Hall E
			Booth # 104
Students@SC	10:00am-6:00pm	Student Cluster Competition	Exhibit Hall 1

WEDNESDAY, NOVEMBER 16

Event	Time	Session Title	Room
Awards Presentations	8:30am-10:00am	Cray/Kennedy/Fernbach Award Recipient Talks	Ballrooms EFGHIJ
Break	10:00am-10:30am	Refreshments (Tech Program)	255-Pre+255-A
Break	10:00am-10:30am	Refreshments (Tech Program)	355-Pre+355-A
Break	3:00pm-3:30pm	Refreshments (Tech Program)	255-Pre+255-A
Break	3:00pm-3:30pm	Refreshments (Tech Program)	355-Pre+355-A
Birds of a Feather	10:30am-12:00pm	European Exascale Projects and their International Collaboration Potential	155-A
Birds of a Feather	12:15pm-1:15pm	PGAS: The Partitioned Global Address Space	355-E
Birds of a Feather	12:15pm-1:15pm	Charting the PMIx Roadmap	250-F
Birds of a Feather	12:15pm-1:15pm	Women in HPC: Intersectionality	155-C
Birds of a Feather	12:15pm-1:15pm	IEEE TCHPC Community Meeting	355-BC
Birds of a Feather	12:15pm-1:15pm	Special Interest Group on HPC in Resource Constrained Environments (SIGHPC-RCE)	250-D
Birds of a Feather	12:15pm-1:15pm	HDF5: State of the Union	250-C
Birds of a Feather	12:15pm-1:15pm	Unleashing the Power of the Intel® Xeon Phi and Beyond	355-F
Birds of a Feather	12:15pm-1:15pm	The Green500: Trends for Energy-Efficient Supercomputing	255-D
Birds of a Feather	12:15pm-1:15pm	Charm++ and AMPI: Adaptive and Asynchronous Parallel Programming	250-E
Birds of a Feather	12:15pm-1:15pm	The Message Passing Interface: On the Road to MPI 4.0 and Beyond	255-EF
Birds of a Feather	1:30pm-3:00pm	OpenHPC Community	155-A
Birds of a Feather	3:30pm-5:00pm	Big Data and Exascale Computing (BDEC) Community Report	155-A
Birds of a Feather	5:15pm-7:00pm	Report: An NSF-Supported Workshop on the Profession: Community-Building and Next Steps	250-C
Birds of a Feather	5:15pm-7:00pm	HPCG Benchmark Update	355-E
Birds of a Feather	5:15pm-7:00pm	Open MPI State of the Union X	255-BC
Birds of a Feather	5:15pm-7:00pm	Impacting Cancer with HPC: Opportunities and Challenges	355-BC

WEDNESDAY, NOVEMBER 16

Event	Time	Session Title	Room
Birds of a Feather	5:15pm-7:00pm	Ceph in HPC Environments	355-F
Birds of a Feather	5:15pm-7:00pm	HPC Outreach: Promoting Supercomputing to the Next Generation	250-E
Birds of a Feather	5:15pm-7:00pm	Monitoring Large Scale HPC Systems: Understanding, Understanding, Diagnosis, and Attribution of Performance Variation and Issues	155-F
Birds of a Feather	5:15pm-7:00pm	Best Practices in Mentoring Undergraduate Research in Supercomputing	355-D
Birds of a Feather	5:15pm-7:00pm	OpenACC API User Experience, Vendor Reaction, Relevance, and Roadmap	155-C
Birds of a Feather	5:15pm-7:00pm	The U.S. Exascale Computing Project	255-EF
Birds of a Feather	5:15pm-7:00pm	High Performance Geometric Multigrid (HPGMG): an HPC Performance Benchmark	250-F
Birds of a Feather	5:15pm-7:00pm	Analyzing Parallel I/O	155-E
Child Care	8:00am-7:00pm	Child care	257
Child Care	8:00am-7:00pm	Room for parents with infants	252-A
Coat/Bag Check 1	7:00am-8:00pm	Coat and bag check	East Entrance
Coat/Bag Check 2	7:00am-8:00pm	Coat and bag check	Hall E
Doctoral Showcase	10:00am-6:00pm	Doctoral Showcase Posters Exhibit	Exhibit Hall E Booth #122
Emerging Technologies	10:00am-6:00pm	Emerging Technologoies	155-B
Exhibitor Forum	10:30am-12:00pm	Security	155-E
Exhibitor Forum	10:30am-12:00pm	Processors	155-F
Exhibitor Forum	3:30pm-5:00pm	GPUs	155-E
Exhibitor Forum	3:30pm-5:00pm	Software	155-F
Exhibits	10:00am-6:00pm	Exhibits	Exhibit Halls A-E, 1-5
Family Day	4:00pm-6:00pm	Family Day (Kids 12 Years and Older) Go to Registration for Badging	Exhibit Halls A-E, 1-5
Gordon Bell Finalists	10:30am-12:00pm	ACM Gordon Bell Finalists I	255-EF
HPC Impact Showcase	1:00pm-3:30pm	HPC Impact Showcase	155-E
Invited Talks	10:30am-12:00pm	• A Super High-Resolution Global Atmospheric Simulation by the Non-Hydrostatic Icosahedral Atmospheric Model, NICAM Masaki Satoh (University of Tokyo) • Reflecting on the Goal and Baseline for Exascale Computing, Thomas Schulthess (Swiss National Supercomputing Center)	Ballroom-EFGHJ

WEDNESDAY, NOVEMBER 16

Event	Time	Session Title	Room
Invited Talks	3:30pm-5:00pm	• Bias: From Overt to Unconscious and What Research Suggests Can Be Done Sharon McGrayne (McGrayne.com) • Memory Bandwidth and System Balance in HPC Systems John McCalpin (University of Texas at Austin)	Ballroom-EFGHIJ
Job Fair	10:00am-3:00pm	Student/Postdoc Job Fair	251-ABDE
Media/Press	9:00am-5:00pm	Media Room and Press Office	151-DE
Panel	10:30am-12:00pm	Bringing about HPC Open-Standards World Peace	255-BC
Panel	1:30pm-3:00pm	Post Moore's Era Supercomputing in 20 Years	255-BC
Panel	3:30pm-5:00pm	The End of Von Neumann? What the Future Looks Like for HPC Application Developers	255-BC
Papers	10:30am-12:00pm	Systems and Networks II	355-D
Papers	10:30am-12:00pm	Performance Measurement and Analysis	355-BC
Papers	1:30pm-3:00pm	Fluid Dynamics	255-EF
Papers	1:30pm-3:00pm	Compilation for Enhanced Parallelism	355-D
Papers	1:30pm-3:00pm	Performance Tools	355-E
Papers	1:30pm-3:00pm	Storage Systems	355-BC
Papers	3:30pm-5:00pm	Accelerator Programming Tools	355-D
Papers	3:30pm-5:00pm	Numerical Algorithms, Part II	355-E
Papers	3:30pm-5:00pm	Memory and Power	355-BC
Posters	8:30am-5:00pm	Research Posters Exhibit	Lower Lobby
Posters	8:30am-5:00pm	Scientific Visualization and Data Analytics Showcase Posters	South Foyer
Press Interview Room	9:00am-5:00pm	Interview Room for Press (reservations only)	151-F
Registration	7:30am-6:00pm	Attendee Registration	South Foyer
Store	7:30am-6:00pm	Conference Store	South Foyer
Student Research Competition	10:00am-6:00pm	ACM Student Research Competition Posters Exhibit	Exhibit Hall E Booth #104
Student Research Competition	3:30pm-5:00pm	ACM Student Research Competition Posters Presentations	155-C
Students@SC	10:00am-5:00pm	Student Cluster Competition	Exhibit Hall 1
Students@SC	10:30am-12:00pm	Experiencing HPC for Undergraduates: Graduate Student Perspective	250-D

THURSDAY, NOVEMBER 16

Event	Time	Session Title	Room
Awards Presentations	12:45pm-1:30pm	Awards Ceremony	Ballroom-EFGHIJ
Birds of a Feather	10:30am-12:00pm	Software Engineering for Computational Science and Engineering on Supercomputers	155-A
Birds of a Feather	12:15pm-1:15pm	Use Cases of Reconfigurable Computing Architectures for HPC	250-C
Birds of a Feather	12:15pm-1:15pm	The Virtual Institute for I/O and the IO-500 List	250-D
Birds of a Feather	12:15pm-1:15pm	SLURM User Group Meeting	355-E
Birds of a Feather	12:15pm-1:15pm	Intel QuickAssist User Gathering	355-D
Birds of a Feather	12:15pm-1:15pm	Multi-Kernel OSes for Extreme-Scale HPC	250-F
Birds of a Feather	1:30pm-3:00pm	Omni-Path User Group (OPUG) Meeting	155-A
Break	10:00am-10:30am	Refreshments (Tech Program)	255-Pre+255-A
Break	10:00am-10:30am	Refreshments (Tech Program)	355-Pre+355-A
Break	3:00pm-3:30pm	Refreshments (Tech Program)	255-Pre+255-A
Break	3:00pm-3:30pm	Refreshments (Tech Program)	355-Pre+355-A
Child Care	8:00am-6:00pm	Child care	257
Child Care	8:00am-6:00pm	Room for parents with infants	252-A
Coat/Bag Check 1	7:00am-6:00pm	Coat and bag check	East Entrance
Coat/Bag Check 2	7:00am-6:00pm	Coat and bag check	Hall E
Doctoral Showcase	10:00am-3:00pm	Doctoral Showcase Posters Exhibit	Exhibit Hall E Booth #122
Emerging Technologies	10:00am-5:00pm	Emerging Technologies	155-B
Exhibitor Forum	10:30am-12:00pm	Exascale & Cloud	155-E
Exhibitor Forum	10:30am-12:00pm	Storage	155-F
Exhibits	10:00am-3:00pm	Exhibits	Exhibit Halls A-E, 1-5
Gordon Bell Finalists	10:30am-12:00pm	ACM Gordon Bell Finalists II	255-EF
HPC Impact Showcase	1:30pm-3:30pm	HPC Impact Showcase	155-E
Invited Talks, Plenary	8:30am-10:00am	• Diversity and Inclusion in Supercomputing Maria Klawe (Harvey Mudd College) • Beyond Exascale: Emerging Devices and Architectures for Computing Thomas Theis (Columbia University)	Ballroom-EFGHIJ
Invited Talks	10:30am-12:00pm	• Co-Design 3.0 – Configurable Extreme Computing, Leveraging Moore's Law for Real Applications, Sadasivan Shankar (Harvard University) • Advances and Challenges in Wildland Fire Monitoring and Prediction, Janice Coen (National Center for Atmospheric Research)	Ballroom-EFGHIJ

THURSDAY, NOVEMBER 16

Event	Time	Session Title	Room
Invited Talks	1:30pm-3:00pm	• HPC Runtime System Software for Asynchronous Multi-Tasking, Thomas Sterling (Indiana University) • Parallel Multiway Methods for Compression of Massive Data and Other Applications Tamara Kolda (Sandia National Laboratories)	Ballroom-EFGHIJ
Media/Press	9:00am-3:00pm	Media Room and Press Office	151-DE
Panel	10:30am-12:00pm	Emerging Realities and New Protocols for Power Distribution: High Voltage and DC	255-BC
Panel	1:30pm-3:00pm	Virtualization, Partitioning, Cohabitation, Containers, Something Else? Approaches to Integrating HPC and Big Data Applications	255-BC
Panel	3:30pm-5:00pm	Leveraging the Architectures, Flexibilities, and Tools Emerging from Members of the OpenStack Scientific Community	255-BC
Papers	10:30pm-12:00pm	Performance Analysis of Network Systems	355-D
Papers	10:30pm-12:00pm	Data Analytics	355-BC
Papers	1:30pm-3:00pm	Manycore Architectures	255-EF
Papers	1:30pm-3:00pm	File Systems and I/O	355-D
Papers	1:30pm-3:00pm	Combinatorial and Multigrid Algorithms	355-E
Papers	1:30pm-3:00pm	Inverse Problems and Quantum Circuits	355-BC
Papers	3:30pm-4:30pm	State-of-the-Practice: System Characterization and Design	355-D
Papers	3:30pm-4:30pm	Task-Oriented Runtimes	355-E
Papers	3:30pm-5:00pm	Accelerating Science	255-EF
Papers	3:30pm-5:00pm	Clouds & Job Scheduling	355-BC
Posters	8:30am-5:00pm	Research Posters Exhibit	Lower Lobby
Posters	8:30am-5:00pm	Scientific Visualization and Data Analytics Showcase Posters	South Foyer
Press Interview Room	9:00am-3:00pm	Interview Room for Press (reservations only)	151-F
Reception	6:00pm-9:00pm	Technical Program Attendees Reception	The Leonardo
Registration	7:30am-5:00pm	Attendee Registration	South Foyer
SC17 Preview	8:25am-8:30am	SC17 Preview with Bernd Mohr	Ballroom-EFGHIJ
Store	7:30am-5:00pm	Conference Store	South Foyer
Student Research Competition	10:00am-3:00pm	ACM Student Research Competition Posters Exhibit	Exhibit Hall E Booth # 104
Students@SC	10:30am-12:00pm	Experiencing HPC for Undergraduates: Careers in HPC	250-D

FRIDAY, NOVEMBER 17

Event	Time	Session Title	Room
Break	10:00am-10:30am	Refreshments (Workshop/Tech Program)	255-Pre+255-A
Break	10:00am-10:30am	Refreshments (Workshop/Tech Program)	355-Pre+355-A
Coat/Bag Check 1	7:00am-1:00pm	Coat and bag check	East Entrance
Coat/Bag Check 2	7:00am-1:00pm	Coat and bag check	Hall E
Panel	8:30am-10:00am	Future of Memory Technology for Exascale and Beyond IV	255-BC
Panel	8:30am-10:00am	HPC File Systems, Challenges, and Use Cases	255-EF
Panel	10:30am-12:00pm	Data Analytics Support for HPC System Management	255-BC
Registration	8:00am-11:00am	Attendee Registration	South Foyer
Store	8:00am-11:00am	Conference Store	South Foyer
Workshop	8:30am-12:00pm	Exascale I/O: Challenges, Innovations and Solutions (ExaI0)	355-BC
Workshop	8:30am-12:00pm	Runtime Systems for Extreme Scale Programming Models and Architectures (RESPA)	155-A
Workshop	8:30am-12:00pm	3rd International Workshop on Visual Performance Analytics – VPA 2016	355-E
Workshop	8:30am-12:00pm	First International Workshop on Communication Optimizations in HPC	355-D
Workshop	8:30am-12:00pm	NRE2016: Numerical Reproducibility at Exascale	155-C
Workshop	8:30am-12:00pm	ESPM2 2016: Second International Workshop on Extreme Scale Programming Models and Middleware	155-E

SC16 Salt Lake City, Utah | hpc matters.

Plenary/Keynote/ Invited Talks

Monday, November 14

HPC Matters Plenary Panel Discussion

Moderator: *Steve Conway (IDC)*
5:45pm-6:45pm
Room: Ballroom-Full

HPC and Precision Medicine: Researchers Are on the Brink of Finding Cures to Cancer and Other Deadly Diseases within This Generation but Only with the Power of HPC

Mitchell Cohen (University of California, San Francisco), Marti Head (GlaxoSmithKline), Dimitri Kusnezov (US Department of Energy), Steve Scott (Cray Inc.)

Leading voices from the frontlines of clinical care, medical research, public policy and pharmaceutical R&D will share diverse perspectives on the future of precision medicine.

Emerging technology is transforming the practice of medicine from an analog field of paper notebooks, textbooks and charts into a digital domain. This transition may emerge as one the most fundamental important tipping points since the discovery of penicillin, molecular biology, and the sequencing of the human genome.

Recently, DOE Secretary Moniz, VA Secretary MacDonald, NCI Director Lowy and the GSK CEO Andrew Witty announced that the Nation's leading supercomputers would be applied to the challenge of the Cancer Moonshot initiative. This partnership of nontraditional groups, collectively see the path to unraveling the complexities of cancer through the power of new machines, operating systems, and applications that leverage simulations, data science and artificial intelligence to accelerate bringing precision oncology to the patients that are waiting. This initiative is one of many research efforts in the race to solve some of our most challenging medical problems.

The success of all of these research programs hinge on harnessing the power of HPC to analyze volumes of complex genomics and other biological datasets that simply can't be processed by humans alone. The challenge for our community will be to develop the computing tools and services needed to transform how we think about disease and bring us closer to the precision medicine future.

Tuesday, November 15

Keynote

Chair: *Dona Crawford (Lawrence Livermore National Laboratory)*
8:30am-10am
Room: Ballroom-Full

Cognitive Computing: How Can We Accelerate Human Decision Making, Creativity and Innovation Using Techniques from Watson and Beyond?

Katharine Frase (IBM)

Dr. Katharine Frase served as Chief Technology Officer of Public Sector and Vice President of Public Sector at International Business Machines Corporation (IBM) from March 2013 to September 2015. In that role, she provided thought leadership for IBM and its customers on innovation and strategic transformation specific to government; education; life sciences; healthcare and cities; and driving the creation of new solutions. Since October 2015, she has led the strategy and business development for IBM's Watson Education unit, providing transformative solutions to improve student outcomes and support instructors across the education spectrum. Prior to this role, she served as Vice President of Industry Solutions Research, working across IBM Research on behalf of IBM clients, to create transformational industry-focused solutions, including the application of IBM Watson technologies to business applications and the realization of Smarter Planet solutions. Earlier roles included technical and business strategy for IBM's software business, corporate assignments on technology assessment and strategy, and roles in IBM Microelectronics in the management of process development, design/modeling methodology and production of chip carriers, assemblies and test. In 2006, she was elected as a member of the (U.S.) National Academy of Engineering. Dr. Frase received an A.B. in chemistry from Bryn Mawr College and a Ph.D. in materials science and engineering from the University of Pennsylvania. She is a member of the IBM Academy of Technology and sits on numerous external committees and boards.

Invited Talks – Session I

10:30am-12pm
Room: Ballroom-Full

Understanding Cities through Computation, Data Analytics, and Measurement

Charles Catlett (Argonne National Laboratory)

Urbanization is one of the great challenges and opportunities of this century, inextricably tied to global challenges ranging from climate change to sustainable use of energy and natural resources, and from personal health and safety to accelerating innovation and education. There is a growing science community—spanning nearly every discipline—pursuing research related to these challenges. The availability of urban data has increased over the past few years, in particular through open data initiatives, creating new opportunities for collaboration between academia and local government in areas ranging from scalable data infrastructure to tools for data analytics, along with challenges such as replicability of solutions between cities, integrating and validating data for scientific investigation, and protecting privacy. For many urban questions, however, new data sources will be required with greater spatial and/or temporal resolution, driving innovation in the use of sensors in mobile devices as well as embedding intelligent sensing infrastructure in the built environment. Collectively these data sources also hold promise to begin to integrate computational models associated with individual urban sectors such as transportation, building energy use, or climate. Catlett will discuss the work that Argonne National Laboratory and the University of Chicago are doing in partnership with the City of Chicago and other cities through the Urban Center for Computation and Data, focusing in particular on new opportunities related to embedded systems and computational modeling.

Bio: Charlie Catlett is the founding director of the Urban Center for Computation and Data, UrbanCCD, which brings scientists, artists, architects, technologists, and policy makers together to use computation, data analytics, and embedded systems to pursue insight into the dynamics, design, and resilient operation of cities. He leads the NSF-funded Array of Things, establishing a network of 500 Argonne-developed intelligent sensor units in Chicago. He is a Senior Computer Scientist at Argonne National Laboratory, a Senior Fellow at the Computation Institute of the University of Chicago and Argonne National Laboratory, and a Senior Fellow at the Harris School of Public Policy at the University of Chicago. In previous roles he was Chief Information Officer for Argonne National Laboratory, Director of the NSF "TeraGrid" nationally distributed supercomputing facility, designer and director of the I-WIRE regional optical network, and Chief Technology Officer at the National Center for Supercomputing Applications. He has worked in Internet and supercomputing technologies since 1985. Recognized as one of 25 "Doers, Dreamers & Drivers" of 2016 by Government Technology magazine and in 2014 as one of Chicago's "Tech 50" technology leaders by Crain's Chicago Business, Charlie is a Computer Engineering graduate of the University of Illinois at Urbana-Champaign.

The Materials Project – A Google of Materials

Kristin Persson (Lawrence Berkeley National Laboratory)

Advanced materials are essential to economic and societal development, with applications in multiple industries, from clean energy, to national security, and human welfare. Traditional empirical and 'one-at-a-time' materials testing is unlikely to meet our future materials innovation challenges in a timely manner. Historically, novel materials exploration has been slow and expensive, taking on average 18 years from concept to commercialization. What is needed is a scalable approach that leverages the talent of the US materials research community and enables a rational design and synthesis of materials from atoms to functionality. The Materials Project (www.materialsproject.org) is harnessing the power of supercomputing together with state of the art quantum mechanical theory to compute the properties of all known inorganic materials and beyond, design novel materials and offer the data for free to the community together with online analysis and design algorithms. The current release contains data derived from quantum mechanical calculations for over 60,000 materials and millions of properties. The software infrastructure enables thousands of calculations per week – enabling screening and predictions - for both novel solid as well as molecular species with target properties. To exemplify the approach of first-principles high-throughput materials design, we will make a deep dive into some of the ongoing work, showcasing the rapid iteration between ideas, new materials development, computations, and insight as enabled by the Materials Project infrastructure and computing resources.

Bio: Kristin Aslaug Persson is an Assistant Professor at UC Berkeley with a joint appointment with the Lawrence Berkeley National Laboratory. She obtained her Ph.D in Theoretical Physics at the Royal Institute of Technology in Stockholm, Sweden in 2001. During the years 2001-2008 she held a postdoctoral appointment and later a research associate position at MIT. Persson is the Director of The Materials Project (www.materialsproject.org) which is an open 'Google' of materials data and – to date – has attracted more than 23,000 users worldwide. She is also the Director of the 2012 BES-funded 'Materials Project Center for Functional Electronic Materials Design' and leads one of the five thrusts in the BES-funded Joint Center for Energy Storage Research (JCESR) (www.jcesr.org). Kristin is also a PI of the Batteries for Advanced Transportation Technologies program (http://batt.lbl.gov) and leads the Theory Guiding Synthesis effort in the 2014 EFRC Center for Next Generation Materials Design (CNGMD). In 2009 she co-founded the clean-energy start-up Pellion Technologies Inc. (www.pelliontech.com), recipient of an ARPA-E award in 2010 for developing high-energy rechargeable magnesium batteries. She was awarded the 2013 LBNL Director's Award for Exceptional Scientific Achievement.

Wednesday, November 16

Invited Talks – Session II

10:30am-12pm
Room: Ballroom-EFGHIJ

A Super High-Resolution Global Atmospheric Simulation by the Non-Hydrostatic Icosahedral Atmospheric Model, NICAM

Masaki Satoh (University of Tokyo)

In the field of the atmospheric science, general circulation models (GCMs) have been used to simulate global atmospheric circulations both for numerical weather forecasts and climate change projections. GCMs use discretized equations for fluids and integrate them for winds, temperature, and humidity. GCMs are also coupled with ocean, land surface, and eco-system models to construct earth system models. Until very recently, horizontal resolution of GCMs is several tens of kilometers. As computer power increases, resolutions of GCMs become higher, and more complicated processes are introduced. Our research team has developed a new type of the global atmospheric model called Non-hydrostatic Icosahedral Atmospheric Model (NICAM) which covers the earth with a quasi-uniform mesh, whose horizontal interval can be a sub-kilometer (870 m) using the K computer. With this model, fine structures of global cloud distributions are reproduced well and can be viewed as a global cloud resolving model (GCRM). The target of GCRM is multi-scale and multi-physics atmospheric phenomena whose scales are from a sub-kilometers to several ten-thousand kilometers with interactive processes of cloud microphysics, radiation, and turbulence. This presentation overviews recent results from the super-high resolution simulations with NICAM using the K computer and outlooks the post K era toward new horizon simulations such as global large-scale eddies (global-LES) simulations.

Bio: Satoh earned his Ph.D. in Geophysics at the University of Tokyo in 1993 and also became a lecturer of Saitama Institute of Technology. In 1998, he was a senior visiting scholar at the Department of Applied Mathematics and Theoretical Physics, Cambridge University. Later, Satoh joined the Japan Agency for Marine-Earth Science and Technology when the development of NICAM first started and still contributes there today. In 2005, he moved to Center for Climate System Research at the University of Tokyo. Since 2011, he has served as professor at AORI and as a senior research scientist at the Japan Aerospace Exploration Agency. Satoh is a director of the Meteorological Society of Japan, a member of the American Meteorological Society, chief editor of Journal of the Meteorological Society of Japan, and section chief editor of "Progress in Earth and Planetary Science." Recently, he was elected as a delegate of the Japan Geoscience Union.

Reflecting on the Goal and Baseline for Exascale Computing

Thomas Schulthess (Swiss National Supercomputing Center)

Application performance is given much emphasis in discussions of exascale computing. A 100-fold increase in sustained performance over today's applications running on multi-petaflops supercomputing platforms should be the expected target for exascale systems deployed early next decade. In the present talk we will reflect on what this means in practice and how much these exascale systems will advance the state of the art. Experience with today's platforms show that there can be an order of magnitude difference in performance within a given class of numerical methods, depending only on choice of architecture and implementation. This bears the questions on what our baseline is, over which the performance improvements of exascale systems will be measured. Furthermore, how close will these exascale systems bring us to deliver on application goals, such as kilometer scale global climate simulations or high-throughput quantum simulations for materials design? We will discuss specific examples from meteorology and materials science.

Bio: Thomas Schulthess received his PhD in physics in 1994 from ETH Zurich. After a postdoctoral fellowship at Lawrence Livermore National Laboratory, he moved to Oak Ridge National Laboratory (ORNL) in 1997. In 2002 he became group leader of the Computational Materials Sciences Group in ORNL's Computer Science and Mathematics Division, and in 2005 was appointed to lead the Nanomaterials Theory Institute of ORNL's Center for Nanophase Materials Sciences. Since October 2008, Thomas holds a chair in Computational Physics at ETH Zurich and directs the Swiss National Supercomputing Center (aka CSCS) in Lugano. He currently holds a visiting distinguished scientist appointment at ORNL. Thomas' research interests are in condensed matter physics as well as the development of high-end computing systems for simulation-based sciences. He led the teams that were awarded the 2008 and 2009 ACM Gordon Bell Prizes, and that ran the first productive simulations with sustained petaflop/s performance.

Invited Talks – Session III

3:30pm-5pm
Room: Ballroom-EFGHIJ

Bias: From Overt to Unconscious and What Research Suggests Can Be Done

Sharon McGrayne (McGrayne.com)

When my book Nobel Prize Women in Science was published in 1993, the legal barriers against women in academic science seemed to be fading into the past. But now we realize that subtle barriers are also difficult to deal with. In my talk, I'll give some examples, past and present, and describe recent research on the subject. In particular, I will draw on what I've learned from the

book that Dr. Rita Colwell, former director of the National Science Foundation, and I are writing about women in science.

Bio: Sharon Bertsch McGrayne is the author of several books about the history of science and scientific discoveries. Her last book was a history of Bayesian probability, The Theory That Would Not Die: How Bayes' Rule Cracked the Enigma Code, Hunted Down Russian Submarines and Emerged Triumphant from Two Centuries of Controversy. It is published by Yale University Press. Her first book, Nobel Prize Women in Science, biographies of 15 leading women scientists up to about 1975, is published by National Academy of Sciences Press. She is currently working on a book about women in science today with Dr. Rita Colwell, former director of the National Science Foundation, entrepreneur, and Distinguished University Professor at the University of Maryland and at Johns Hopkins University Bloomberg School of Public Health. McGrayne is a former newspaper reporter and a graduate of Swarthmore College. She lives in Seattle, and her webpage is at McGrayne.com.

Memory Bandwidth and System Balance in HPC Systems

John McCalpin (University of Texas at Austin)

The "Attack of the Killer Micros" began approximately 25 years ago as microprocessor-based systems began to compete with supercomputers (in some application areas). It became clear that peak arithmetic rate was not an adequate measure of system performance for many applications, so in 1991 McCalpin introduced the STREAM Benchmark to estimate "sustained memory bandwidth" as an alternative performance metric.

STREAM apparently embodied a good compromise between generality and ease of use and quickly became the "de facto" standard for measuring and reporting sustained memory bandwidth in High Performance Computing systems.

Since the initial "attack", Moore's Law and Dennard Scaling have led to astounding increases in the computational capabilities of microprocessors. The technology behind memory subsystems has not enjoyed comparable performance improvements, causing sustained bandwidth to fall behind.

This talk reviews the history of the changing "balances" between computation, memory latency, and memory bandwidth in deployed HPC systems, then discusses how the underlying technology changes led to these market shifts. Key metrics are the increasing relative "cost" of memory accesses and the massive increases in concurrency that are required to obtain increasing memory throughput.

A review of new technologies (such as "stacked DRAM") shows that the difficulties of delivering increased memory bandwidth are not alleviated unless the underlying computer architectures are changed in fundamental ways. The combination of technology trends and economic factors suggest that system balances will continue to shift in the same directions - favoring workloads with increasingly high compute intensity and increasing available concurrency.

Bio: John D. McCalpin is a Research Scientist in the High Performance Computing Group and Co-Director of the Advanced Computing Evaluation Laboratory (ACElab) at the Texas Advanced Computing Center (TACC) of the University of Texas at Austin. At TACC, he works on performance analysis and performance modeling in support of both current users and future system acquisitions. McCalpin joined TACC in 2009 after a twelve-year career in performance analysis and system architecture in the computer industry. This included three years at SGI (performance analysis and optimization on the SGI Origin 2000 and performance lead on the architecture team for the SGI Altix 3000), six years at IBM (performance analysis for HPC, processor and system design for Power4/4+ and Power5/5+, high-level architecture for the Power7-based PERCS Prototype system), and three years at AMD (technology lead for the "Torrenza" program enabling third-party accelerated computing technologies). Prior to his industrial career, John was an oceanographer (Ph.D., Florida State), spending six years as an assistant professor at the University of Delaware engaged in research and teaching on numerical simulation of the large-scale circulation of the oceans. In 1991, McCalpin developed the STREAM Benchmark to provide a simple way to measure and report sustainable memory bandwidth, and for 25 years has been an advocate for a multidimensional approach to understanding performance in HPC. In 2015 he was named an "Intel Black Belt Software Developer" in recognition of his contributions to the Intel Software Developer communities.

Thursday, November 17

Invited Talks – Session IV

8:30am-10am
Room: Ballroom-EFGHIJ

Diversity and Inclusion in Supercomputing

Maria Klawe (Harvey Mudd College)

Like many other computing research areas, women and other minority groups are significantly under-represented in supercomputing. This talk discusses successful strategies for significantly increasing the number of women and students of color majoring in computer science and explores how these strategies might be applied to supercomputing.

Bio: Maria Klawe began her tenure as Harvey Mudd College's fifth president in 2006. Prior to joining HMC, she served as dean of engineering and professor of computer science at Princeton University. Klawe joined Princeton from the University of British Columbia where she served in various roles from 1988 to 2002. Prior to UBC, Klawe spent eight years with IBM Research in California and two years at the University of Toronto. She

received her Ph.D. (1977) and B.Sc. (1973) in mathematics from the University of Alberta. Klawe is a member of the nonprofit Math for America, the chair of the board of the nonprofit EdReports.org, a fellow of the American Academy of Arts & Sciences, and a trustee for the Mathematical Sciences Research Institute in Berkeley and a member of the Canada Excellence Research Chairs Selection Board. She is the recipient of the 2014 Women of Vision ABIE Award for Leadership and was ranked 17 on Fortune's 2014 list of the World's 50 Greatest Leaders.

Beyond Exascale: Emerging Devices and Architectures for Computing

Thomas Theis (Columbia University)

The continuing evolution of silicon CMOS technology is clearly approaching some important physical limits. Since roughly 2003, the inability to reduce supply voltages according to constant-field scaling rules, combined with economic constraints on areal power density and total power, has forced designers to limit clock frequencies even as devices have continued to shrink. Still, there is a plausible path to exascale, based on the continued evolution of silicon device technology, silicon photonics, 3D integration, and more. The immediate challenge is to execute. However, longer term research exploring entirely new devices and architectures is essential if we want to take high performance computing well beyond exascale. Recent years have brought a large increase in research funding and interest in new device concepts. Some of the devices explored to date, such as tunneling field-effect transistors (TFETs) based on III-V semiconductors, promise to open a new low-power design space which is inaccessible to conventional FETs. Nanomagnetic devices may allow memory and logic functions to be combined in novel ways. And newer, perhaps more promising device concepts continue to emerge. At the same time, research in new architectures has also grown. Indeed, at the leading edge, researchers are beginning to focus on co-optimization of new devices and new architectures. Despite the growing research investment, the landscape of promising research opportunities outside the "FET devices and circuits box" is still largely unexplored.

Bio: Thomas Theis is Executive Director of the Columbia Nano Initiative (CNI) and Professor in Electrical Engineering in the Fu Foundation School of Engineering, Arts and Sciences. He joined IBM at the T.J. Watson Research Center in 1978 to study electronic properties of materials, and held various senior management and executive positions from 1984 through 2015. In the late 1990's, as Senior Manager, Silicon Science and Technology, he coordinated the transfer of copper interconnection technology from IBM Research to the IBM Microelectronics Division. The replacement of aluminum chip wiring by copper was an industry first, the biggest change in chip wiring technology in thirty years. As IBM's strategist for exploratory research worldwide from 1998 to 2012 and as Director, Physical Sciences from 1998 to 2010, he conceived and initiated successful research programs in silicon nanophotonics and Josephson junction-based quantum computing, and championed research in nanoelectronics, exploratory memory devices, and applications of information technology to address societal needs in energy, infrastructure, and the environment. From 2010–2012, as Program Manager, New Devices and Architectures for Computing, he organized research projects aimed at greatly improved energy-efficiency in computing. On assignment from IBM to the Semiconductor Research Corporation (SRC) from 2012–2016, he led SRC's Nanoelectronics Research Initiative, a private-public partnership funding university research aimed at new devices and circuits for computing. He joined Columbia University in April of 2016 to manage and lead CNI operations and work with faculty and university research offices to identify and develop concepts for major new research programs.

Invited Talks – Session V

10:30am-12pm
Room: Ballroom-EFGHIJ

Co-Design 3.0 – Configurable Extreme Computing, Leveraging Moore's Law for Real Applications

Sadasivan Shankar (Harvard University)

In this talk, we will discuss Co-design 3.0, a more adaptable and scalable framework in which systems can be dynamically configured driven by the specific needs of the applications. We will examine 6 different scaling paradigms as basis of this co-design: combinatorial nature of scientific problems, multi-scale nature of systems, algorithms, complexity of applications, Moore's law, and economics of scaling. With the power of computing to solve problems for addressing societal needs, we need to focus on real applications as they evolve in time as opposed to standard benchmarks. For this to be practically viable with reduced costs, this should be done in a scalable manner. We think that major ongoing research and development centers of computational and physical sciences need to be formally engaged in the co-design of hardware, software, numerical methods, algorithms, and applications. As we will demonstrate with a few examples, this will help address grand scientific (technology) challenges associated with the societal problems: materials, chemistry, environment, health, and information processing. In addition, this will help in wider dispersion of the benefits of computing rather than to niche scientific communities. As part of this talk, we will also briefly illustrate a new class that we have developed in which students are taught hands-on about using extreme computing to address real applications.

Bio: Sadasivan (Sadas) Shankar is currently the first Margaret and Will Hearst Visiting Lecturer in Computational Science and Engineering at Harvard School of Engineering and Applied Sciences, where he is involved in teaching and research in the areas of large-scale computational methods, chemistry, materials, and in translational ideas. In fall 2013, as the first Distinguished

Scientist in Residence at the Institute of Applied Computational Sciences in Harvard, he co-instructed a graduate-level class on Computational Materials Design, which covered fundamental atomic and quantum techniques and practical applications for new materials by design. Dr. Shankar earned his Ph.D. in Chemical Engineering and Materials Science from University of Minnesota, Minneapolis. Sadasivan has initiated and led multiple efforts in the electronics industry, most recently the Materials Design Program in Intel. He is a co-inventor in over 20 patent filings covering areas in new chemical reactor designs, semiconductor processes, bulk and nano materials, device structures, and algorithms. He is also a co-author in publications and presentations in the areas of chemical synthesis, plasma chemistry and processing, Monte Carlo and molecular dynamics methods, non- equilibrium electronic, ionic, and atomic transport, chemical reactive systems, and methods for experimental characterization and high throughout calculations for material interfaces. Dr. Shankar has been also involved in several collaborative national and international efforts with Semiconductor Research Corporation, Sematech, National Institute of Standards and Technology, Department of Energy, and President's Materials Genome Initiative.

Advances and Challenges in Wildland Fire Monitoring and Prediction

Janice Coen (National Center for Atmospheric Research)

The culture of fire suppression reveres tradition, standardized training, and organizational hierarchy. Yet, the past two decades have seen the infusion of technology that has transformed the understanding, observation, and prediction of wildland fires and their behavior, as well as provided a much greater appreciation of its frequency, occurrence, and attribution in a global context. These advances arose through cross fertilization from numerical weather prediction - reflecting advances in numerical weather prediction, data assimilation, and high performance computing - as well as new remote sensing instrumentation and platforms from satellite to unmanned aerial systems, mathematical algorithms, computer science, ecology, and the access to and interpretation of vast datasets on fires, fuels, and weather. This wave of technology has brought many goals within sight – rapid fire detection, nearly ubiquitous monitoring, and recognition that much of the distinctive characteristics of fire events are reproducible and perhaps predictable in real time. Concurrently, these more complex innovations raise new challenges, as only marginal improvement may be generated by further refinements in algorithms and model resolution; legal and safety considerations limit observations of dangerous phenomena and introduction of new technology; predictions face limits to predictability; and the complexity of identifying and forecasting the statistically rare megafire event among the 40,000–100,000 yearly wildfire events in the U.S., of which only a few percent will grow beyond a hundred hectares. This talk will highlight current research in integrated weather – wildland fire computational modeling, fire detection and observation, and their application to understanding and prediction.

Bio: Dr. Janice Coen is a Project Scientist at the National Center for Atmospheric Research in Boulder, Colorado. She received a B.S. in Engineering Physics from Grove City College and an M.S. and Ph.D. from the Department of Geophysical Sciences at the University of Chicago. She is currently an Associate Editor for the International Journal of Wildland Fire and a member of the Editorial Board of Environmental Modelling and Software and has served has served as a member of the Board of Directors of the International Association of Wildland Fire and on the Colorado Aerial Firefighting Center of Excellence Vision/Focus Subcommittee, the Office of the Federal Coordinator for Meteorology National Wildland Fire Weather Needs Assessment Joint Action Group, and interagency Fire Research Coordination Council. She currently investigates wildland fire behavior and its interaction with weather using coupled weather-wildland fire computer simulation models and by analyzing infrared imagery of wildland fires. Her current work investigates the unfolding of large wildland fire events, distilling understanding to improve firefighter safety, drought and fuel mitigation impacts on fire behavior, assimilation of satellite active fire detection data, and the use of coupled weather-fire models as forecasting tools.

Invited Talks – Session VI

1:30pm-3pm
Room: Ballroom-EFGHIJ

HPC Runtime System Software for Asynchronous Multi-Tasking

Thomas Sterling (Indiana University)

Increasing sophistication of application program domains combined with expanding scale and complexity of HPC system structures is driving innovation in computing to address sources of performance degradation. Efficiency and scalability of parallel resource operation (now peaking at 10 million cores) demand that challenges of starvation, latency, overhead, and contention in supercomputing be resolved. Asynchronous multi-tasking through runtime system software is of increasing interest and is being explored within the US and internationally to complement more conventional approaches in pursuit of effective exascale computing. In principle, runtime systems may exploit compute time information about application state and system usage for dynamic and adaptive introspective control of resource management and task scheduling. The intent is to achieve superior resource utilization, load balancing, data migration and affinity, and parallelism discovery. Other motivations include reducing programmer burden and portability across systems of different types, scales, and generations. However, advanced scalable run-

time systems are experimental and impose additional problems such as increased system software complexity, added overheads, and uncertainty about programming interfaces, support for legacy codes, and workload interoperability. Further, early results suggest that not all applications will benefit significantly through runtime support, with instances of performance reduction observed. This presentation will provide a comprehensive review of driving challenges, strategies, examples of existing runtime systems, and experiences. One important consideration is the possible future role of advances in computer architecture to accelerate the likely mechanisms embodied within typical runtimes. The talk will conclude with suggestions of future paths and work to advance this possible strategy.

Bio: Dr. Thomas Sterling holds the position of Professor of Electrical Engineering at the Indiana University (IU) School of Informatics and Computing Department of Intelligent Systems Engineering (ISE) and serves as Director of the IU Center for Research in Extreme Scale Technologies (CREST). Since receiving his Ph.D from MIT in 1984 as a Hertz Fellow, Dr. Sterling has engaged in applied research in parallel computing system structures, semantics, and operation in industry, government labs, and academia. Dr. Sterling is best known as the "Father of Beowulf" for his pioneering research in commodity/Linux cluster computing for which he shared the Gordon Bell Prize in 1997. He led the HTMT Project sponsored by multiple agencies to explore advanced technologies and their implication for high-end computer system architectures. Other research projects in which he contributed included the DARPA DIVA PIM architecture project with USC-ISI, the DARPA HPCS program sponsored Cray-led Cascade Petaflops architecture, and the Gilgamesh high-density computing project at NASA JPL. Sterling is currently involved in research associated with the innovative ParalleX execution model for extreme scale computing to establish the foundation principles guiding the development of future generation exascale computing systems. ParalleX is currently the conceptual centerpiece of the XPRESS project as part of the DOE X-stack program and has been demonstrated via the proof-of-concept HPX-5 runtime system software. Dr. Sterling is the co-author of six books and holds six patents. He was the recipient of the 2013 Vanguard Award and is a Fellow of the AAAS.

Parallel Multiway Methods for Compression of Massive Data and Other Applications

Tamara G. Kolda (Sandia National Laboratories)

Scientists are drowning in data. The scientific data produced by high-fidelity simulations and high-precision experiments are far too massive to store. For instance, a modest simulation on a 3D grid with 500 grid points per dimension, tracking 100 variables, for 100 time steps yields 5TB of data. Working with this massive data is unwieldy, and it may not be retained for future analysis or comparison. Data compression is a necessity, but there are surprisingly few options available for scientific data. We propose to exploit the 5-way structure (3D spatial grid x time x variable) of the data by applying Tucker tensor decomposition to reveal a latent low-dimensional representation. By taking advantage of multiway structure, we are able to compress combustion science data by a factor of 10-1000 with negligible loss in accuracy. Additionally, we need not reconstruct the entire data set to extract subparts or down-sampled versions. However, compressing such massive data requires a parallel implementation of the Tucker tensor decomposition. We explain the data distribution and algorithm and accompanying analysis. We apply the algorithm to real-world data sets to demonstrate the speed, compression performance, and accuracy of the method. We also consider extensions of this work into functional representations (useful for hierarchical/irregular grids and reduced order models) as well as acceleration via randomized computations. This talk will highlight work by collaborators Woody Austin, Grey Ballard, Alicia Klinvex, Hemanth Kolla, and others.

Bio: Tamara G. Kolda is a Distinguished Member of the Technical Staff at Sandia National Laboratories in Livermore, CA. She holds a Ph.D. in applied mathematics from the University of Maryland at College Park and is a past Householder Postdoctoral Fellow in Scientific Computing at Oak Ridge National Laboratory. She has received several awards for her work including a 2003 Presidential Early Career Award for Scientists and Engineers (PECASE), an R&D 100 Award, and three best paper prizes. She is a Distinguished Scientist of the Association for Computing Machinery (ACM) and a Fellow of the Society for Industrial and Applied Mathematics (SIAM). She is currently a member of the SIAM Board of Trustees, Section Editor for the Software and High Performance Computing Section for the SIAM Journal on Scientific Computing, and Associate Editor for the SIAM Journal on Matrix Analysis and Applications.

SC16 Salt Lake City, Utah | hpc matters.

Award Presentations

Tuesday, November 15

Test of Time Award Special Lecture

1:30pm-3pm
Room: Ballroom

Automatically Tuned Linear Algebra Software

Clint Whaley (University of Tennessee), Jack Dongarra (University of Tennessee)

This paper describes an approach for the automatic generation and optimization of numerical software for processors with deep memory hierarchies and pipelined functional units. The production of such software for machines ranging from desktop workstations to embedded processors can be a tedious and time consuming process. The work described here can help in automating much of this process. We will concentrate our efforts on the widely used linear algebra kernels called the Basic Linear Algebra Subroutines (BLAS). In particular, the work presented here is for general matrix multiply, DGEMM. However much of the technology and approach developed here can be applied to the other Level 3 BLAS and the general strategy can have an impact on basic linear algebra operations in general and may be extended to other important kernel operations.

Wednesday, November 16

Cray/Kennedy/Fernbach Awards

8:30am-10am
Room: Ballroom-EFGHIJ

Featuring talks by the Cray, Fernbach and Kennedy Award Recipients

ACM Gordon Bell Finalist I

10:30am-12pm
Room: 255-EF

Towards Green Aviation with Python at Petascale

Peter Vincent (Imperial College London), Freddie Witherden (Imperial College London), Brian Vermeire (Imperial College London), Jin Seok Park (Imperial College London), Arvind Iyer (Imperial College London)

Accurate simulation of unsteady turbulent flow is critical for improved design of 'greener' aircraft that are quieter and more fuel-efficient. We demonstrate application of PyFR, a Python based computational fluid dynamics solver, to petascale simulation of such flow problems. Rationale behind algorithmic choices, which offer increased levels of accuracy and enable sustained computation at up to 58% of peak DP-FLOP/s on unstructured grids, will be discussed in the context of modern hardware. A range of software innovations will also be detailed, including use of runtime code generation, which enables PyFR to efficiently target multiple platforms, including heterogeneous systems, via a single implementation. Finally, results will be presented from a full-scale simulation of flow over a low-pressure turbine blade cascade, along with weak/strong scaling statistics from the Piz Daint and Titan supercomputers, and performance data demonstrating sustained computation at up to 13.7 DP-PFLOP/s.
Award: BP

Modeling Dilute Solutions Using First-Principles Molecular Dynamics: Computing More than a Million Atoms with Over a Million Cores

Jean-Luc Fattebert (Lawrence Livermore National Laboratory), Daniel Osei-Kuffuor (Lawrence Livermore National Laboratory), Erik W. Draeger (Lawrence Livermore National Laboratory), Tadashi Ogitsu (Lawrence Livermore National Laboratory), William D. Krauss (Lawrence Livermore National Laboratory)

First-Principles Molecular Dynamics (FPMD) methods, although powerful, are notoriously expensive computationally due to the quantum modeling of electrons. Traditional FPMD approaches have typically been limited to a few thousand atoms at most, due to O(N3) or worse solver complexity and the large amount of communication required for highly parallel implementations. Attempts to lower the complexity have often introduced uncontrolled approximations or systematic errors. Using a robust new algorithm, we have developed an O(N) complexity solver for electronic structure problems with fully controllable numerical error. Its minimal use of global communications yields excellent scalability, allowing very accurate FPMD simulations of more than a million atoms on over a million cores. At these scales, this approach provides multiple orders of magnitude speedup compared to the standard plane-wave approach typically used in condensed matter applications, without sacrificing accuracy. This will open up entire new classes of FPMD simulations such as dilute aqueous solutions.

Simulations of Below-Ground Dynamics of Fungi: 1.184 Pflops Attained by Automated Generation and Autotuning of Temporal Blocking Codes

Takayuki Muranushi (RIKEN), Hideyuki Hotta (Chiba University), Junichiro Makino (Kobe University), Seiya Nishizawa (RIKEN), Hirofumi Tomita (RIKEN), Keigo Nitadori (RIKEN), Masaki Iwasawa (RIKEN), Natsuki Hosono (RIKEN), Yutaka Maruyama (RIKEN), Hikaru Inoue (Fujitsu Ltd), Hisashi Yashiro (RIKEN), Yoshifumi Nakamura (RIKEN)

Stencil computation has many applications in science and engineering, thus many optimization techniques such as temporal blocking have been developed. They are, however, rarely used in real-world applications, since a large amount of careful programming is required for even the simplest of stencils. We introduce Formura, a domain specific language that provides easy access to optimized stencil computations. Higher-order integration schemes can be defined using mathematical notations. Formura generates C code with MPI calls and performs autotuning. Hence its performance is portable to most distributed-memory computers. We show the scientific applicability of Formura by performing magnetohydrodynamics (MHD) and below-ground biology simulations. Ability to reach bytes-per-flops ratio only attainable by temporal blocking is demonstrated. We also demonstrate scaling up to the full nodes of the K computer, with 1.184 Pflops, 11.62% floating-point-operation efficiency, and 31.26% memory throughput efficiency.

Thursday, November 17

ACM Gordon Bell Finalist II

10:30am-12pm
Room: 255-EF

Extreme-Scale Phase Field Simulations of Coarsening Dynamics on the Sunway TaihuLight Supercomputer

Jian Zhang (Chinese Academy of Sciences), Chunbao Zhou (Chinese Academy of Sciences), Yangang Wang (Chinese Academy of Sciences), Lili Ju (University of South Carolina), Qiang Du (Columbia University), Xuebin Chi (Chinese Academy of Sciences), Dongsheng Xu (Chinese Academy of Sciences), Dexun Chen (National Research Center of Parallel Computer Engineering and Technology), Yong Liu (National Research Center of Parallel Computer Engineering and Technology), Zhao Liu (National Supercomputing Center in Wuxi)

Many important properties of materials such as strength, ductility, hardness, and conductivity are determined by the microstructures of the material. During the formation of these microstructures, grain coarsening plays an important role. The Cahn-Hilliard equation has been applied extensively to simulate the coarsening kinetics of a two-phase microstructure. It is well accepted that the limited capabilities in conducting large scale, long time simulations constitute bottlenecks in predicting microstructure evolution based on the phase field approach. We present here a scalable time integration algorithm with large step-sizes and its efficient implementation on the Sunway TaihuLight supercomputer. The highly nonlinear and severely stiff Cahn-Hilliard equations with degenerate mobility for microstructure evolution are solved at extreme scale, demonstrating that the latest advent of high performance computing platform and the new advances in algorithm design are now offering us the possibility to simulate the coarsening dynamics accurately at unprecedented spatial and time scales.

A Highly Effective Global Surface Wave Numerical Simulation with Ultra-High Resolution

Fangli Qiao (First Institute of Oceanography), Wei Zhao (First Institute of Oceanography), Xunqiang Yin (First Institute of Oceanography), Xiaomeng Huang (Tsinghua University), Xin Liu (National Research Center of Parallel Computer Engineering and Technology), Qi Shu (First Institute of Oceanography), Guansuo Wang (First Institute of Oceanography), Zhenya Song (First Institute of Oceanography), Xinfang Li (First Institute of Oceanography), Guangwen Yang (Tsinghua University), Haixing Liu (First Institute of Oceanography), Yeli Yang (First Institute of Oceanography)

Surfaces waves are the most energetic motion in global ocean and are crucially important to marine safety and climate change. High-resolution global wave model plays a key role for accurate wave forecasting. However, the parallel efficiency with large amount of computation is a big barrier for this type of model. In this work, breakthroughs in the design and application of irregular quasi-rectangular grid decomposition, master-slave cooperative computing workflow, and pipelining schemes for high resolution global wave model have been achieved. Based on these innovations, ultra-high horizontal resolution of (1/60)◦ by (1/60)◦ global wave model is implemented on the Sunway TaihuLight Supercomputer with 123.6 PFlops peak performance. Preliminary results show peak performance of our model can reach 35.97 PFlops with a full-scale system consisting of 10,495,680 cores. Final results before performance update deadline should be higher. These innovations provide good scalability and high efficiency for the ultra-high resolution global wave model.

10M-Core Scalable Fully-Implicit Solver for Nonhydrostatic Atmospheric Dynamics

Chao Yang (Chinese Academy of Sciences), Wei Xue (Tsinghua University), Haohuan Fu (Tsinghua University), Hongtao You (National Research Center of Parallel Computer Engineering and Technology), Xinliang Wang (Tsinghua University), Yulong Ao (Chinese Academy of Sciences), Fangfang Liu (Chinese Academy of Sciences), Lin Gan (Tsinghua University), Ping Xu (Tsinghua University), Lanning Wang (Beijing Normal University), Guangwen Yang (Tsinghua University), Weimin Zheng (Tsinghua University)

An ultra-scalable fully-implicit solver is developed for stiff time-dependent problems arising from the hyperbolic conservation laws in nonhydrostatic atmospheric dynamics. In the solver, we propose a highly efficient hybrid multigrid domain decomposition preconditioner that can greatly accelerate the convergence of the solver at the extreme scale. For solving the overlapped subdomain problems, a physics-based multi-block asynchronized incomplete LU factorization method is customized to further exploit the on-chip fine-grained concurrency. We perform systematic optimizations on different hardware levels to achieve best utilization of the heterogeneous computing units and a substantial reduction of data movement costs. The fully-implicit solver enables fast and accurate atmospheric simulations on the new Sunway TaihuLight supercomputer in China, scaling to over ten million heterogeneous cores and achieving a sustained performance of over two petaflops.

Awards Ceremony

12:45pm-1:30pm
Room: Ballroom

The SC16 conference awards, as well as selected ACM and IEEE awards, will be presented. The awards include: Best Paper, Best Student Paper, Best Poster, Best Scientific Visualization, ACM Student Research Competition, ACM Gordon Bell Prize, ACM/IEEE-CS George Michael Memorial HPC Fellowship, ACM SIGHPC / Intel Computational & Data Science Fellowships, IEEE TCHPC Award for Excellence for Early Career Researchers in High Performance Computing, and Student Cluster Competition.

SC16 Salt Lake City, Utah | hpc matters.

Birds of a Feather

Tuesday, November 15

Meeting of the SIGHPC - Big Data Chapter

10:30am-12pm
Room: 155-A

Stratos Efstathiadis (New York University), Suzanne McIntosh (New York University)

The goal of the BoF is to gather for the first time members and non-members of the SIGHPC Big Data Virtual Chapter who are interested in learning about the challenges of converging Big Data and HPC. The BoF will give people the opportunity to hear about existing challenges and openly discuss solutions, tools and new approaches on how to best utilize available Big Data and HPC resources.

Presentations: *Scott Yockel (Harvard Research Computing): "Big Data, where doesn't it come from, and how I deal with it?" Harry Mangalam (UC Irvine Research Computing): "BeeGFS in real life"*

13th Graph500 List

12pm-1:30pm
Room: 155-A

Richard Murphy (Micron Technology Inc), David Bader (Georgia Institute of Technology), Peter Kogge (University of Notre Dame)

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 13th 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.

Burst Buffers: Early Experiences and Outlook

12:15pm-1:15pm
Room: 255-D

Deborah Bard (Lawrence Berkeley National Laboratory), David Paul (Lawrence Berkeley National Laboratory), Wahid Bhimji (Lawrence Berkeley National Laboratory), Bilel Hadri (King Abdullah University of Science and Technology), Robert Ross (Argonne National Laboratory), Cornell Wright (Los Alamos National Laboratory)

The long-awaited Burst Buffer technology is now being deployed on major supercomputing systems, including new machines at NERSC, LANL, ANL, and KAUST. In this BOF, we discuss early experience with Burst Buffers from both a systems and a user's perspective, including challenges faced and perspectives for future development. Short presentations from early adopters will be followed by general discussion with the audience. We hope that this BOF will attract interest and participation from end-users and software/hardware developers.

HPC Education: Meeting of the SIGHPC Education Chapter

12:15pm-1:15pm
Room: 355-D

Richard Coffey (Argonne National Laboratory), Steven Gordon (Ohio State University), Scott Callaghan (University of Southern California), David Halstead (National Radio Astronomy Observatory)

The purpose of the SIGHPC Education chapter is to encourage interest in and knowledge of the applications of HPC by supporting educational activities. This BOF will bring together those interested in promoting HPC education through the formal and informal activities of the chapter. The current officers will present information on the chapter organization, membership, and a review of current and proposed activities. They will then lead an open discussion from participants to solicit their ideas and feedback on chapter activities.

Lustre Community BOF: Lustre Deployments for the Next 5 Years

12:15pm-1:15pm
Room: 255-BC

Stephen Simms (Indiana University), Frank Baetke (Hewlett Packard Enterprise)

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 world-wide contributors. Lustre is now more widely used and in more mission-critical installations than ever. Lustre currently supports many HPC infrastructures scientific research and financial services, oil and gas, advanced manufacturing, and visual effects. At the Lustre BOF, the worldwide community of developers, administrators, and solution providers will gather to discuss recent developments, such as Progressive File Layouts (PFL) and the ability to scale metadata servers (Distributed Namespace), new challenges, and corresponding opportunities.

New Technologies, Platforms and Services for Sharing Research Data: Data Commons, Distributed Clouds, and Distributed Data Services

12:15pm-1:15pm
Room: 250-D

Robert Grossman (University of Chicago), Allison Heath (University of Chicago)

Data commons collocate data, storage, and computing infrastructure with core data services and commonly used tools and applications for managing, analyzing, and sharing data to create an interoperable resource for the research community. This session will discuss practical experiences designing, building, and operating data commons for the research community. It will also discuss key services that data commons require, such as index services, metadata services, etc.

Next Generation of Co-Processors Emerges: In-Network Computing

12:15pm-1:15pm
Room: 255-EF

Gilad Shainer (HPC Advisory Council), Richard Graham (Mellanox Technologies), Scott Atchley (Oak Ridge National Laboratory)

The latest revolution in HPC is the move to a co-design architecture, a collaborative effort among industry, academia, and manufacturers to reach exascale performance by taking a holistic system-level approach to fundamental performance improvements. Co-design architecture exploits system efficiency and optimizes performance by creating synergies between the hardware and the software.

Co-design recognizes that the CPU has reached the limits of its scalability and offers an intelligent network as the new "co-processor" to share the responsibility for handling and accelerating application workloads. By placing data-related algorithms on an intelligent network, we can dramatically improve the data center and applications performance.

On-Demand Infrastructure for Data Analytics and Storage

12:15pm-1:15pm
Room: 250-C

Franck Cappello (Argonne National Laboratory), Gabrielle Allen (University of Illinois)

Extreme-scale simulations and experiments can generate much more data than can be stored and analyzed at a single site. The scientific community needs significant improvements on data access for complex and accurate analyses. The BOF will provide a forum for groups to discuss approaches, successes, and challenges in distributed computation, storage and analysis. The focus is on developing this community by leading by example, and speakers will report on the SC16 "on-demand infrastructure for scientific data analytics and storage" experiment. The expected outcome is more groups experimenting with their own distributed computing/data projects. A report will summarize the BOF outcomes.

SAGE2: Scalable Amplified Group Environment for Global Collaboration

12:15pm-1:15pm
Room: 250-E

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

SAGE2 is an innovative, user-centered, web-based platform that enables local and/or distributed teams to display, manage, share and investigate large-scale datasets on tiled display walls in order to glean insights and discoveries. SAGE2 enables users to easily display and juxtapose digital media (2D/3D charts, graphs, images, movies or documents) from a variety of sources. SAGE2, announced at SC14, is the next-generation SAGE (Scalable Adaptive Graphics Environment), which was the defacto operating system for managing Big Data content on tiled display walls for almost a decade.

SAGE2 provides the scientific community with persistent visualization and collaboration services for global cyberinfrastructure.

SIGHPC Annual Member Meeting

12:15pm-1:15pm
Room: 355-F

The annual business meeting of SIGHPC is your opportunity to hear about and discuss the status of SIGHPC and its chapters. We will also be discussing upcoming plans for the year. All of the elected officers and many of the other volunteers will be present to answer your questions about SIGHPC. Representatives from our chapters: Education, Big Data, and Resource Constrained Environments (RCE) will also be available. Come tell us how SIGHPC can meet your needs.

The 2016 HPC Challenge Awards

12:15pm-1:15pm
Room: 155-C

Piotr Luszczek (University of Tennessee), Jeremy Kepner (Massachusetts Institute of Technology), Jack Dongarra (University of Tennessee)

The 2016 HPC Challenge Awards BOF is the 10th 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.

OpenStack for HPC: Best Practices for Optimizing Software-Defined Infrastructure

1:30pm-3pm
Room: 155-A

Jonathan Mills (NASA), Mike Lowe (Indiana University)

OpenStack is becoming increasingly capable for management of HPC infrastructure and support of HPC workloads. However, performance overheads, network integration, and system complexity all combine to pose a daunting challenge. The optimal outcome is to achieve all the benefits of software-defined infrastructure without paying any of the penalties. How can we get closer to achieving this?

This BOF is aimed at architects, administrators and software engineers who are interested in designing and deploying OpenStack infrastructure for HPC, but are unsure how to begin. Speakers from OpenStack's Scientific Working Group will share their experiences and engage the audience in discussions.

The Future of NSF Advanced Cyberinfrastructure

3:30pm-5pm
Room: 155-A

William Miller (National Science Foundation), Rudi Eigenmann (National Science Foundation)

The National Science Foundation's vision and investment plans for cyberinfrastructure address the growing needs of the science and engineering research community. Program Directors from the NSF Division of Advanced Cyberinfrastructure will discuss new funding opportunities across the spectrum of computing, software, data infrastructure, networking, cybersecurity, and learning and workforce development. Presentations will also cover new interdisciplinary opportunities associated with NSF domain science and priority areas, such as the National Strategic Computing Initiative, Understanding the Brain, and Innovations at the Nexus of Food, Energy and Water Systems (INFEWS). Ample time will be provided for audience questions and engagement with NSF staff.

Distributed Machine Intelligence Using Tensorflow

5:15pm-7pm
Room: 155-E

Karan Bhatia (Google), Kevin Kissel (Google)

TensorFlow is the second-generation machine learning system from the Google Brain team, using tensor notation to describe neural networks as stateful dataflow graphs in ways that are both concise and flexible. This BOF session brings together developers of TensorFlow, users of TensorFlow, and members of the HPC community curious about how to develop and exploit deep learning systems for data analysis and autonomous operation.

Emerging Trends in HPC Systems and Application Modernization

5:15pm-7pm
Room: 255-D

Hans-Christian Hoppe (Intel Corporation), Marie-Christine Sawley (Intel Corporation)

This BOF discusses the application evolution required by clearly apparent HPC system trends: complex memory hierarchies with fast limited-size memory, addition of large capacity non-volatile memory, substantial increase in cores, and integration of high-performance interconnects. Applications need to reshape to achieve good performance, taking advantage of these architecture trends, while limiting complexity to allow efficient code development and maintenance. Developers of significant HPC applications will discuss their approaches and experiences in code evolution, preparing a discussion with the audience. This BOF aims to foster a global community of practitioners across funding organizations and schemes and combining academic and industrial participants.

Energy Efficiency Considerations and HPC Procurement

5:15pm-7pm
Room: 255-BC

*Natalie Bates (Lawrence Livermore National Laboratory), Ladina Gilly (Swiss National Supercomputing Center), James Laros (Sandia National Laboratories), James Rogers (Oak Ridge National Laboratory), Anna Maria Bailey (Lawrence Livermore National Laboratory), Daniel Hackenberg (Technical University Dresden), Marek Michalewicsz (A*STAR Computational Resource Centre), Bilel Hadri (King Abdullah University of Science and Technology), Thomas Ilsche (Technical University Dresden)*

The predominant goal for procurement of HPC systems is to identify the optimal solution to both technical and financial targets that maximizes the contribution of that system to the organization's mission. Beyond the acquisition cost of the system, it is also important to consider the total costs of ownership, including the improvements necessary to host the system, the infrastructure that supports its operation, and the significant operational costs associated with that new HPC system. In this BoF, HPC leaders will discuss and debate key procurement requirements and lessons learned that can contribute to greater energy efficiency and reduced operational costs.

Experimental Infrastructure and Methodology for HPC Cloud Research

5:15pm-7pm
Room: 250-D

Kate Keahey (Argonne National Laboratory)

The confluence of cloud computing, HPC, and Big Data is attracting increasingly more research interest with challenges ranging from systems research, networking, through resource and power management to new algorithms and innovative applications. New testbeds have been established to support this research and have attracted an active community of users. The objective of this BOF is to provide a forum that brings together operators of experimental testbeds for HPC and cloud computing as well as their current and prospective users to discuss current and future capabilities, user requirements, and experimental methodology.

How to Build Diverse Teams for More Effective Research

5:15pm-7pm
Room: 250-C

Toni Collis (EPCC at the University of Edinburgh), Lorna Rivera (University of Illinois)

Most of us recognize that diverse teams are good for productivity and output. But do you know how to improve diversity and build a more inclusive environment? Have you ever heard of unconscious bias, stereotype threat, or imposter syndrome? Do you ever feel like you aren't good enough to be in the community or feel like a 'fraud'? This BoF will discuss the real effects of these three topics on the workplace, providing the audience with an introduction to each theme, how they may affect you, and how they impact employers, employees, advisors, managers, or your peers.

MPICH: A High-Performance Open-Source MPI Implementation

5:15pm-7pm
Room: 250-E

Pavan Balaji (Argonne National Laboratory), Ken Raffenetti (Argonne National Laboratory)

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

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

OpenMP: Where Is It Now and Where Is It Going?

5:15pm-7pm
Room: 355-E

Jim Cownie (Intel Corporation), Michael Klemm (Intel Corporation)

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

This lively interactive BOF will change your mind, showcasing modern OpenMP ("The language that lets you use all your compute resources") by presenting OpenMP 4.5 and the vision for OpenMP 5.0.

Experts will give short presentations on key OpenMP 4.5 features, and then answer your questions on its technical aspects. Michael Klemm (OpenMP's new CEO) will present our vision for OpenMP 5.0 and we'll finish with an audience led discussion with a panel including members of the OpenMP Architecture Review Board.

Reconfigurable Supercomputing

5:15pm-7pm
Room: 250-F

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

Reconfigurable Supercomputing (RSC) is characterized by hardware that adapts to match the needs of each application, offering unique advantages in performance per unit energy for high-end computing. 2016 is a breakout year for RSC, with datacenter deployment by Microsoft, acquisition of Altera and new devices by Intel, RSC innovations by IBM, new devices and tools by Altera and Xilinx, and success of Novo-G# (world's first large-scale RSC) in the NSF CHREC Center. This BOF introduces architectures of such systems, describes applications and tools being developed, and provides a forum for discussing emerging opportunities and issues for performance, productivity, and sustainability.

The 2016 Ethernet Roadmap

5:15pm-7pm
Room: 155-C

John D'Ambrosia (Ethernet Alliance), Greg McSorley (Amphenol), David Chalupsky (Intel Corporation), David Rodgers (Teledyne LeCroy)

Ethernet Alliance developed the 2016 Ethernet Roadmap as a good overview of the latest developments in Ethernet standards. This BoF will discuss the latest developments in Ethernet and show how 2.5, 5 and 25GbE standards were completed in 2016 while 50 and 200GbE standards are under development with 400GbE. The BoF will also discuss Flex Ethernet that is considered next-generation LAG and creates speeds from 75GbE to 500GbE. The BoF will discuss the latest developments in new optical module form factors for 100-400GbE including the microQSFP, QSFP-DD and on-board optics (OBO). Physical copies of the roadmap will be handed out.

Today's Hot Technology: The Growing Necessity of Liquid Cooled HPC

5:15pm-7pm
Room: 355-D

Geoff Lyon (CoolIT Systems Inc), Wade Doll (Cray Inc.), Michael Patterson (Intel Corporation), Steve Hammond (National Renewable Energy Laboratory), Sammy Lee Zimmerman (Hewlett Packard Enterprise)

This BoF surrounds the necessity of liquid cooling in data center and HPC environments as increases in chip processing power propels the industry toward liquid in order to manage heat loads that are not possible with air. A panel of industry experts, vendors, and users will discuss the responsible management of global liquid cooling integration. Topics of interest include education of customers, training for and management of installations and upgrades, international service requirements, and the ongoing responsibility for every step of the innovation to be as efficient and environmentally friendly as possible despite the incredible processing power.

TOP500 Supercomputers

5:15pm-7pm
Room: 255-EF

Erich Strohmaier (Lawrence Berkeley National Laboratory), Jack Dongarra (University of Tennessee), Horst Simon (Lawrence Berkeley National Laboratory), Martin Meuer (ISC Group)

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 48th TOP500 list will be published in November 2016 just in time for SC16.

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

Wednesday, November 16

European Exascale Projects and Their International Collaboration Potential

10:30am-12pm
Room: 155-A

Jean-François Lavignon (European Technology Platform for High Performance Computing), Marcin Ostasz (European Technology Platform for High Performance Computing), Jean-Philippe Nominé (European Technology Platform for High Performance Computing), Jesus Carretero (Charles III University of Madrid)

The European HPC technology and application projects will showcase their work, aiming to seed new collaborations with international partners, to accelerate progress in exascale R&D.

Within the European HPC effort, a number of projects are running plus funding has been guaranteed for further investments in a variety of areas covering the entire HPC system stack and application expertise (see http://www.etp4hpc.eu/en/euexascale.html for a full description).

We will highlight the most important areas of these projects, discuss opportunities for international collaboration, and identify key mechanisms required to facilitate this process. Project representatives will be present to facilitate networking.

Charm++ and AMPI: Adaptive and Asynchronous Parallel Programming

12:15pm-1:15pm
Room: 250-E

Phil Miller (CharmWorks Inc), Laxmikant (Sanjay) Kale (University of Illinois), Sam White (University of Illinois)

A community gathering about parallel programming using Charm++, Adaptive MPI, the many applications built on them, and associated tools. This session will cover recent advances in Charm++ and the experiences of application developers with Charm++. There will also be a discussion on the future directions of Charm++ and opportunities to learn more and form collaborations.

Charm++ is a production-grade many-tasking programming framework and runtime system for modern HPC systems. It offers high productivity and performance portability through features such as multicore and accelerator support, dynamic load balancing, fault tolerance, latency hiding, interoperability with MPI and OpenMP, and online job-resizing.

Charting the PMIx Roadmap

12:15pm-1:15pm
Room: 250-F

Ralph Castain (Intel Corporation), David Solt (IBM), Artem Polyakov (Mellanox Technologies)

The PMI Exascale (PMIx) community will be concluding its second year of existence this fall that included the release of several implementations in both commercial and open source resource managers. We'll discuss what PMIx has accomplished over the past year and present a proposed roadmap for next year.

The PMIx community includes viewpoints from across the HPC runtime community. To that end, we solicit feedback and suggestions on the roadmap in advance of the session, and will include time for a lively discussion at the meeting. So please join us at the BOF to plan the roadmap. New contributors welcome!

HDF5: State of the Union

12:15pm-1:15pm
Room: 250-C

Quincey Koziol (Lawrence Berkeley National Laboratory), David Pearah (HDF Group)

A forum for HDF developers and users to interact. HDF developers will describe the current status of HDF5 and discuss future plans – including connectors for Big Data and Cloud – followed by an open discussion and our first annual community award for best use of HDF5 tech.

IEEE TCHPC Community Meeting

12:15pm-1:15pm
Room: 355-BC

The goal of this meeting is to introduce the recently established IEEE Computer Society Technical Consortium on High Performance Computing (TCHPC) and to explore how TCHPC can meet the needs of the community. The purpose of the TCHPC is to advance and coordinate the work in the field of high performance computing networking, storage, and analysis concepts, technologies and applications, which is carried throughout the IEEE and to expand the IEEE CS' and IEEE's role in this interdisciplinary and pervasive field. The TCHPC shall provide a forum for the discussion and the exchange of information to advance the theory and the practice in this field of interest. The TCHPC will promote technical activities, sponsor technical meetings and sessions, promote periodicals, educational activities, promote standards, and engage in any other activity defined within its field of interest. This work will be accomplished through the TCHPC's support and management of the IEEE CS' shared interest in the SC conference and related SC activities. Its current member Technical Committees are Technical Committee on Parallel Processing (TCPP) and Technical Committee on Computer Communications (TCCC).

Optimizing Performance on Many-Core Processors: Unleashing the Power of the Intel® Xeon Phi and Beyond

12:15pm-1:15pm
Room: 355-F

Richard Gerber (National Energy Research Scientific Computing Center), Thomas Steinke (Zuse Institute Berlin), Kent Milfeld (University of Texas at Austin), Michael Lysaght (Irish Centre for HighEnd Computing)

The first HPC systems based on the next-generation many-core Intel Xeon Phi processor (KNL) are becoming available to the community this fall. This BOF, conducted by the Intel Xeon Phi Users Group (IXPUG), will provide a forum for application and tool developers, HPC center staff, and industry experts to discuss their early successes and challenges. This BOF will showcase code optimization successes, highlighting tuning methodologies, tools features, and real-world impact and community codes. IXPUG is an independent users group for anyone interested in application performance on the Intel Xeon Phi. See http://ixpug.org

PGAS: The Partitioned Global Address Space

12:15pm-1:15pm
Room: 355-E

Tarek El-Ghazawi (George Washington University), Lauren Smith (National Security Agency)

PGAS, or the partitioned global address space programming model, strikes a balance between the ease of programming due to its global address view and performance and due to locality awareness. The latter can even help in reducing power as data movement is a major contributor to power consumption. Current active PGAS efforts include Chapel, UPC, and X10. However, there are common PGAS problems that must be addressed. The PGAS BOF at SC16 will bring together researchers and practitioners from vendors as well open efforts for cross-fertilization of new ideas and to address common issues and common infrastructures for PGAS.

Special Interest Group on HPC in Resource Constrained Environments (SIGHPC-RCE)

12:15pm-1:15pm
Room: 250-D

Hensley Omorodion (University of Benin), Elizabeth Leake (STEM-Trek)

This BOF is hosted by the Special Interest Group on High Performance Computing in Resource Constrained Environments (SIGHPC-RCE). Information about the SIG will be shared, and six topics of interest to attendees will be presented.

The SC15 BOF was attended by 30 delegates from the US, Europe, the Southern African Development Community (SADC) HPC Forum, SADC Advisers, and others. Group objectives were discussed as were common challenges for those who support research in resource constrained environments. Since SC15, SIGHPC-RCE and others have collaborated with STEM-Trek to define a number of projects that are in various stages of development.

The Green500: Trends for Energy-Efficient Supercomputing

12:15pm-1:15pm
Room: 255-D

Wu Feng (Virginia Polytechnic Institute and State University), Erich Strohmaier (Lawrence Berkeley National Laboratory), Natalie Bates (Lawrence Livermore National Laboratory), Thomas Scogland (Lawrence Livermore National Laboratory)

With power a first-order design constraint on par with performance, it is important to measure and analyze trends for energy-efficient supercomputing. This BOF will discuss trends across the Green500 and highlights from the latest Green500 list. The Green500 and TOP500, in collaboration with the Energy-Efficient HPC WG, have a newly integrated submission process designed to streamline future submissions and to provide a consistent set of data for the historical record. It will also provide a forum for community review of the integrated submission process. The BOF will close with an awards presentation, recognizing the most energy-efficient supercomputers in the world.

The Message Passing Interface: On the Road to MPI 4.0 and Beyond

12:15pm-1:15pm
Room: 255-EF

Martin Schulz (Lawrence Livermore National Laboratory)

The MPI Forum, the standardization body for the Message Passing Interface (MPI), recently released version MPI 3.1, which features minor improvements and corrections over MPI 3.0. The focus now shifts to work on the next major version, MPI 4.0, which includes initiatives for fault tolerance, improved support for hybrid programming models and application hints to MPI libraries to enable optimizations. In this BOF, we will discuss changes in MPI 3.1 and highlight details behind the new initiatives for MPI 4.0, seek user community feedback, and continue our active discussion with the HPC community on features and priorities for MPI.

Women in HPC: Intersectionality

12:15pm-1:15pm
Room: 155-C

Rebecca Hartman-Baker (National Energy Research Scientific Computing Center), Fernanda Foertter (Oak Ridge National Laboratory), Toni Collis (EPCC at the University of Edinburgh)

There are many groups that are under-represented in the HPC community, including women and African-Americans, but particularly poorly represented are those that fall into the intersection of two or more underrepresented groups. In this BOF, we hear the stories of women of different minority backgrounds in the HPC field, and the complex intersection between gender, race, sexual orientation and more, and how this has shaped their experience in HPC. We ask these women for their advice on making the HPC field more inclusive for people of all backgrounds.

OpenHPC Community BOF

1:30pm-3pm
Room: 155-A

Karl Schulz (Intel Corporation), David Brayford (Leibniz Supercomputing Centre)

There is a growing sense within the HPC community for the need to have an open community effort to more efficiently build, test, and deliver integrated HPC software components and tools. Formed initially in November 2015 and formalized as a Linux Foundation project in June 2016, OpenHPC is endeavoring to address this need. At this BOF, speakers from the OpenHPC Technical Steering committee will provide a technical overview of the project and near-term roadmaps. We then invite open discussion giving attendees an opportunity to provide feedback on current conventions, packaging, request additional components and configurations, and discuss general future trends.

Big Data and Exascale Computing (BDEC) Community Report

3:30pm-5pm
Room: 155-A

Jack Dongarra (University of Tennessee), Mark Asch (European Extreme Data & Computing Initiative), Peter Beckman (Argonne National Laboratory)

The emergence of large scale data analytics and machine learning in a wide variety of scientific fields has disrupted the landscape on which emerging plans for exascale computing are developing. Participants in the international workshop series on Big Data and Extreme-scale Computing (BDEC) are systematically mapping out the ways in which the major issues associated with data intensive science interact with plans for achieving exascale computing. This meeting will present an overview of this road mapping effort and elicit community input on the development of plans for the convergence of currently bifurcated software ecosystems on a common software infrastructure.

Analyzing Parallel I/O

5:15pm-7pm
Room: 155-E

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

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

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

Further information: http://wr.informatik.uni-hamburg.de/events/2016/bof-monitoring

Best Practices in Mentoring Undergraduate Research in Supercomputing

5:15pm-7pm
Room: 355-D

Nancy Amato (Texas A&M University), Max Grossman (Rice University)

We present, discuss, and refine best practices on mentoring undergraduate researchers. We define "best" practices as those which 1)encourage student interest in high-performance computing, 2)produce high quality results, and 3) build student interest in the field. While some might cite the breakdown of Dennard scaling as the largest challenge facing the HPC community, scaling parallel computing education is both a more alarming and challenging issue.

The intended audience is graduate students, postdocs, faculty, and researchers who are interested in mentoring undergraduates. The session will comprise a presentation on best practices followed by a discussion of topics of interest to the audience.

Ceph in HPC Environments

5:15pm-7pm
Room: 355-F

Douglas Fuller (Red Hat Inc), James Wilgenbusch (University of Minnesota)

Ceph is an open-source distributed object store with an associated file system widely used in cloud and distributed computing. In addition, both the object store and file system components are seeing increasing deployments as primary data storage for traditional HPC. Ceph is backed by a robust, worldwide open source community effort with broad participation from major HPC and storage vendors.

High Performance Geometric Multigrid (HPGMG): an HPC Performance Benchmark

5:15pm-7pm
Room: 250-F

Mark Adams (Lawrence Berkeley National Laboratory), Samual Williams (Lawrence Berkeley National Laboratory), Jed Brown (University of Colorado, Boulder), Erich Strohmaier (Lawrence Berkeley National Laboratory), John Shalf (Lawrence Berkeley National Laboratory), Brian Van Straalen (Lawrence Berkeley National Laboratory)

This meeting facilitates community participation in the HPGMG project. HPGMG is a compact benchmark designed as a design tool for emerging architectures and a ranking metric. HPGMG is well balanced with respect to modern HPC applications and algorithms. We compile the HPGMG-FV list of the world's largest supercomputers with the metric, a multigrid solve of a fourth order accurate finite volume discretization of the Laplacian. We released our first list at ISC16, and continue with our next biannual release at SC16. We encourage community participation with submissions to the HPGMG-FV list, and contributed talks and discussion in the BOF.

HPC Outreach: Promoting Supercomputing to the Next Generation

5:15pm-7pm
Room: 250-E

Nick Brown (EPCC at the University of Edinburgh), Scott Callaghan (University of Southern California), Lorna Rivera (University of Illinois)

The goal of outreach is to promote HPC and our science research to the general public. In order to encourage the next generation of scientists and inform the public why "HPC matters", we need to carefully design outreach activities to have

maximum impact. This BoF hosts an interactive discussion of public engagement activities to learn from each other's experiences, explore what works well and provide attendees with a starting point for their own outreach. There will also be an opportunity to experiment with existing activities and demos. Issues surrounding encouraging diversity in public engagement and best practice will be explored.

HPCG Benchmark Update

5:15pm-7pm
Room: 355-E

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

The High Performance Conjugate Gradients (HPCG) Benchmark is a community metric for ranking HPC systems. The first list of results was released at ISC'14, including optimized results for systems built upon Fujitsu, Intel, Nvidia technologies. Lists have been announced at SC14, IS'14, SC15 and ISC16, with an increase from 15 to 25, 40, 64 and 82 entries, respectively.

In this BOF we present an update of HPCG 3.1, and opportunities for optimizing performance, with presentations from all vendors who have participated in HPCG optimization efforts. We spend the remaining time in open discussion.

Impacting Cancer with HPC: Opportunities and Challenges

5:15pm-7pm
Room: 355-BC

Eric Stahlberg (Frederick National Laboratory for Cancer Research), Patricia Kovatch (Icahn School of Medicine at Mount Sinai), Sean Hanlon (National Cancer Institute)

High-performance computing and HPC technologies have long been employed in multiple roles in cancer research and clinical applications. The importance of HPC in medical applications has been highlighted within the National Strategic Computing Initiative, several activities associated with the US Cancer Moonshot effort and collaborations between the US National Cancer Institute and the Department of Energy. This BoF session will focus on bringing together those with an interest in furthering the impact of HPC on cancer, providing overviews of application areas and technologies, opportunities and challenges, while providing for in-depth interactive discussion across interest areas.

Monitoring Large Scale HPC Systems: Understanding, Diagnosis, and Attribution of Performance Variation and Issues

5:15pm-7pm
Room: 155-F

Ann Gentile (Sandia National Laboratories), Jim Brandt (Sandia National Laboratories), Hans-Christian Hoppe (Intel Corporation), Mike Mason (Los Alamos National Laboratory), Mark Parsons (EPCC at the University of Edinburgh), Marie-Christine Sawley (Intel Corporation), Mike Showerman (University of Illinois)

This BOF addresses critical issues in large-scale monitoring from the perspectives of worldwide HPC center system administrators, users, and vendors. This year will be 100% facilitated audience interactive discussion on tools, techniques, experiences, and gaps in understanding, diagnosing, and attributing causes behind performance variation and poor performance. Causes include contention for shared network and I/O resources and system component problems. Our goal is to facilitate enhancement of community monitoring and analysis capabilities by identifying useful tools and techniques and encouraging the development of quickstart guides for these tools to be posted at the community web site: https://sites.google.com/site/monitoringlargescalehpcsystems/

Open MPI State of the Union X

5:15pm-7pm
Room: 255-BC

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

It's been another great year for Open MPI. We've added new features, improved performance, completed MPI-3.1, and are continuing to drive the state of the art in HPC. We'll discuss what Open MPI has accomplished over the past year and present a roadmap for next year.

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

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

OpenACC API User Experience, Vendor Reaction, Relevance, and Roadmap

5:15pm-7pm
Room: 155-C

Duncan Poole (NVIDIA Corporation), Sunita Chandrasekaran (University of Delaware), Fernanda Foertter (Oak Ridge National Laboratory)

Since SC11, the OpenACC BOF brings together the user and developer communities to gather valuable feedback on the specification and discuss newer features as required by the scientific community code owners.

Up for discussion are the OpenACC extensions used on the Sunway TaihuLight supercomputer. Also, AMD, Intel and NVIDIA released accelerators using high bandwidth memory, which present optimization opportunities for the OpenACC roadmap.

OpenACC adoption is driven by the member-users. This BOF allows users to discuss new experiences, and scaling developer engagement like Hackathons, to improve the use of directives in new and legacy code.

Report to the Facilitator/Cyberpractitioner Community of an NSF-Supported Workshop on the Profession: Community-Building and Next Steps

5:15pm-7pm
Room: 250-C

James Bottum (Internet2), Stephen Wolff (Internet2), Dustin Atkins (Clemson University)

This BoF is directed at the community of cyberpractitioners-facilitators-campus champions, and others with interest in this cadre of experts who are essential to data- and compute-intensive research.

The format is a brief report on an NSF-supported invitation-only workshop held in July 2016, organized around the results of a ten-topic survey of the 35 workshop participants, followed by audience discussion of, and contributions to, the survey questions and responses.

The BoF goals are to: illuminate a critical aspect of workforce development in the HPC and Data Analytics communities, enlarge the survey knowledge base, and suggest continuing activities to strengthen the profession.

The U.S. Exascale Computing Project

5:15pm-7pm
Room: 255-EF

Paul Messina (Argonne National Laboratory), Stephen Lee (Los Alamos National Laboratory)

The U.S. Exascale Computing Project (ECP) aims to develop at least two diverse, capable exascale architectures, a broad set of exascale applications that meet mission and science needs, and a software stack to support them. ECP is part of the National Strategic Computing Initiative launched by the Obama Administration in July 2015. This BOF will offer the community the opportunity to hear an update on the current scope, status, and activities of the project, as well as future plans, and to interact with and ask questions of the ECP leadership team.

Thursday, November 17

Software Engineering for Computational Science and Engineering on Supercomputers

10:30am-12pm
Room: 155-A

David Bernholdt (Oak Ridge National Laboratory), Jeffrey Carver (University of Alabama), Mike Heroux (Sandia National Laboratories), Neil Chue Hong (EPCC at the University of Edinburgh), Daniel Katz (University of Illinois), James Lin (Shanghai Jiao Tong University), Kengo Nakajima (University of Tokyo)

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

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

Intel QuickAssist User Gathering

12:15pm-1:15pm
Room: 355-D

Sven Karlsson (Technical University of Denmark), Pascal Schleuniger (Fachhochschule Nordwestschweiz)

In 2015, Intel made a number of experimental systems incorporating Intel QuickAssist technology available. The systems consist of an Intel processor tightly connected to an Altera FPGA via direct package pins.

The Intel QuickAssist user gathering aims at providing a forum for users of the aforementioned systems, or users of QuickAssist technology, to exchange experiences. The end goal is to build a community for the QuickAssist technology.

The gathering will have presentations intermixed with discussions on how to build the community. It is open to everyone interested in the QuickAssist technology, the Hardware Accelerator Research Program and FPGAs in general.

Multi-Kernel OSes for Extreme-Scale HPC

12:15pm-1:15pm
Room: 250-F

Rolf Riesen (Intel Corporation), Balazs Gerofi (RIKEN)

Lightweight multi-kernel operating systems -- where Linux and a lightweight kernel run side-by-side on a compute node -- have received attention recently for their potential to solve many of the challenges system software faces as we move toward ever more complex and parallel systems.

The four leading multi-kernel projects are at a stage where input from users, runtime developers, and programming model architects, is desired to guide further design and implementation. This BOF is an ideal venue to engage the community, foster participation, and initiate long term collaboration.

SLURM User Group Meeting

12:15pm-1:15pm
Room: 355-E

Morris Jette (SchedMD LLC)

Slurm is an open source workload manager used on many TOP500 systems. It 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 16.05, describe the Slurm roadmap, and solicit user feedback. Everyone interested in Slurm use and/or development is encouraged to attend.

The Virtual Institute for I/O and the IO-500 List

12:15pm-1:15pm
Room: 250-D

Julian Kunkel (German Climate Computing Center), Gerald Lofstead (Sandia National Laboratories), John Bent (Seagate Technology LLC)

Due to the increasing complexity of HPC data management, activities in the storage research community have increased over the last few years. The general purpose of this BoF is to foster this community and discuss the role of the international Virtual Institute for I/O (VI4IO, http://vi4io.org) in supporting, developing, and maintaining this community.

The speakers will give talks covering VI4IO and issues of benchmarking storage systems. The direction of these efforts is then discussed with the participants. A specific purpose of the BoF is to elaborate whether the community would be well-served by an IO-500 benchmark similar to the Top500.

Use Cases of Reconfigurable Computing Architectures for HPC

12:15pm-1:15pm
Room: 250-C

Marie-Christine Sawley (Intel Corporation), Hans-Christian Hoppe (Intel Corporation), John Shalf (Lawrence Livermore National Laboratory)

Use of reconfigurable execution units (mainly FPGAs) for HPC workloads has seen ups and downs, with the last heyday about a decade ago. Significant advances have been made recently in the field of FPGAs and their system integration: new generations provide highly efficient FP units, and fast cache-coherent interconnects to CPUs were announced. On the SW side, the momentum around OpenCL is lowering the entry barriers. This BOF assembles a distinguished panel of speakers, who will give an up-to-date view of existing proof points, specific potential of current and next-generation reconfigurable computing platforms, and remaining limitations to a wider take-up.

Omni-Path User Group (OPUG) Meeting

1:30pm-3pm
Room: 155-A

Nick Nystrom (Pittsburgh Supercomputing Center), Philip Murphy (Intel Corporation)

The goal of this BoF is to provide a forum for users and other interested parties to share their experiences and insights using the new Omni-Path interconnect fabric from Intel. The format will consist of an introduction, an OPA update, panel presentations from OPA site representatives, and group discussion, targeting 50% of the time to the short presentations and 50% to audience engagement.

The following panelists are confirmed: -Nick Nystrom, PSC: Bridges -Kim F. Wong, University of Pittsburgh -TBA, Cineca: Marconi -TBA, University of Colorado, Boulder

Additional panelists will be announced at https://www.psc.edu/index.php/opug.

Doctoral Showcase

Tuesday, November 15

Doctoral Showcase 1

10:30am-12pm
Room: 155-C

Characterizing and Improving Power and Performance of HPC Networks

Taylor L. Groves (University of New Mexico)

Networks are the backbone of modern HPC systems. They serve as a critical piece of infrastructure, tying together applications, analytics, storage, and visualization. Despite this importance, we have not fully explored how evolving communication paradigms and network design will impact scientific workloads. As networks expand in the race towards exascale, a principled approach should be taken to reexamine this relationship, so that the community better understands (1) characteristics and trends in HPC communication, (2) how to best design HPC networks and (3) opportunities in the future to save power or enhance the performance. Our thesis is that by developing new models, benchmarks, and monitoring techniques, we may better understand how to improve performance and power of HPC systems, with a focus on networks. This dissertation highlights opportunities for improving network performance and power efficiency, while uncovering pitfalls (and mitigation strategies) brought about by shifting trends in HPC communication.

Performance and Energy Aware Workload Partitioning on Heterogeneous Platforms

Li Tang (University of Notre Dame)

Heterogeneous platforms which employ a mix of CPUs and accelerators such as GPUs have been widely used in HPC. Such heterogeneous platforms have the potential to offer higher performance at lower energy cost than homogeneous platforms. However, it is rather challenging to actually achieve the high performance and energy efficiency promised by heterogeneous platforms. To address this issue, this work proposes a framework to assist application developers to partition workload on heterogeneous platforms for achieving high performance or energy efficiency before actual implementation. The framework includes both analytical performance/energy models and two sets of workload partitioning guidelines. Based on the design goal, application developers can obtain a workload partitioning guideline for a given platform and then follow it to partition workload. Then the performance/energy models can be used to estimate the performance or energy of the obtained workload partitions and help select the appropriate workload partition.

Dynamic Power Management for Hardware Over-Provisioned Systems

Daniel A. Ellsworth (University of Oregon)

To enable safe operation of hardware over-provisioned computational clusters, my work investigates how to engineer distributed power management software in systems where per component power capping hardware is available. My work contributes a formalization for the goal of a power management solution, a generalized model for roughly estimating the performance effect of processor power capping on application runtime, and a fully dynamic power management strategy tested through simulation and experimentation on an existing HPC system.

From Detection to Optimization: Understanding Silent Error's Impact on Scientific Applications

Jon Calhoun (University of Illinois)

As HPC systems are becoming more complex, they are becoming more vulnerable to failures. In particular, silent errors that lead to silent data corruption (SDC) are of particular concern. SDC refers to change in application state without any indication that a failure occurred, and can lead to longer simulation times or perturbations in results. Understanding how SDC impacts applications and how to create low cost SDC detectors is critical for large-scale applications. Full detection is cost prohibitive; therefore, we seek detection of SDC that impacts results. Acceptance of small perturbations in state allows for new optimizations such a lossy compression to mitigate memory bottlenecks.

In this talk, I discuss SDC impacts on HPC applications and detail a customized SDC detection and recovery scheme for algebraic multigrid linear solvers. I conclude showing how lossy compression can improve checkpoint-restart performance by adding small errors guided by a compression error tolerance selection methodology.

Doctoral Showcase 2

1:30pm-3pm
Room: 155-C

Parallel Storage Systems for Large-Scale Machines

Christos Filippidis (National and Kapodistrian University of Athens)

HPC has crossed the petaflop mark and is reaching the exaflop range quickly. The exascale system is projected to have millions of nodes, with thousands of cores for each node. At such an extreme scale, the substantial amount of concurrency can cause a critical contention issue for I/O system. This study proposes a dynamically coordinated I/O architecture for addressing some of the limitations that current parallel file systems and storage architectures are facing with very large-scale systems. The fundamental idea is to coordinate I/O accesses according to the topology/profile of the infrastructure, the load metrics, and the I/O demands of each application. The measurements have shown that by using IKAROS approach we can fully utilize the provided I/O and network resources, minimize disk and network contention, and achieve better performance.

High Performance File System and I/O Middleware Design for Big Data on HPC Clusters

Nusrat Islam (Ohio State University)

HDFS is the primary storage engine for Hadoop MapReduce, Spark, and HBase. HDFS, along with these Big Data middleware, is increasingly being used in HPC platforms for scientific applications. Modern HPC clusters are equipped with high performance interconnects (e.g. InfiniBand), heterogeneous storage devices, and parallel file systems. But HDFS cannot fully leverage these resources of modern HPC clusters. In this thesis, an RDMA (Remote Direct Memory Access)-Enhanced HDFS design is proposed that maximizes overlapping among different stages of HDFS operations. Data placement policies for HDFS are also devised to efficiently utilize the heterogeneous storage media, such as RAM Disk, SSD, HDD, and Parallel File System. A key-value store-based burst-buffer system to integrate Hadoop with Lustre has also been presented. Finally, advanced designs to exploit the byte-addressability of (Non-Volatile Memory) NVM for HDFS are proposed. Co-designs with MapReduce, Spark, and HBase offer significant performance benefits for the respective middleware and applications.

Low Design-Risk Checkpointing Storage Solution for Exascale Supercomputers

Nilmini Abeyratne (University of Michigan)

This work presents a checkpointing solution for exascale supercomputers that employs commodity DRAM and SSD devices that pose a low design risk compared to solutions that use emerging non-volatile memories.

The proposed local checkpointing solution uses DRAM and SSD in tandem to provide both speed and reliability in checkpointing. A Checkpoint Location Controller (CLC) is implemented to monitor the endurance of the SSD and the performance loss of the application and to decide dynamically whether to checkpoint to the DRAM or the SSD. The CLC improves both SSD endurance and application slowdown; but the checkpoints in DRAM are exposed to device failures. To design a reliable exascale memory, a low latency ECC is added to correct all errors due to bit/pin/column/word faults and also detect errors due to chip failures, and a second Chipkill-Correct level ECC is added to protect the checkpoints residing in DRAM.

Realizing a Self-Adaptive Network Architecture for HPC Clouds

Feroz Zahid (Simula Research Laboratory)

Clouds offer significant advantages over traditional cluster computing architectures including ease of deployment, rapid elasticity, and an economically attractive pay-as-you-go business model. However, the effectiveness of cloud computing for HPC systems still remains questionable. When clouds are deployed on lossless interconnection networks, challenges related to load-balancing, low-overhead virtualization, and performance isolation hinder full potential utilization of the underlying interconnect. In this work, we attack these challenges and propose a novel holistic framework of a self-adaptive IB subnet for HPC clouds. Our solution consists of a feedback control loop that effectively incorporate optimizations based on the multidimensional objective function using current resource configuration and provider-defined policies. We build our system using a bottom-up approach, starting by prototyping solutions tackling individual research challenges associated, and later combining our novel solutions into a working self-adaptive cloud prototype. All our results are demonstrated using state-of-the-art industry software to enable easy integration into running systems.

Doctoral Showcase 3

3:30pm-5pm
Room: 155-C

Automatic Discovery of Efficient Divide-and-Conquer Algorithms for Dynamic Programming Problems

Pramod Ganapathi (Stony Brook University)

We show that it is possible to develop algorithm(s) / framework(s) to automatically / semi-automatically discover algorithms that are simultaneously nontrivial, simple, fast, portable, and robust, which can be used to solve to a wide class of dynamic programming (DP) problems.

We present AutoGen -- an algorithm that given any blackbox implementation of a DP recurrence from a wide class of DP problems can automatically discover a fast recursive divide-and-conquer algorithm for solving that problem on a shared-memory multicore machine. [Published in PPoPP 2016]

We present AutoGen-Wave -- a framework for computer-assisted discovery of fast divide-and-conquer wavefront versions of the algorithms already generated by AutoGen. [Under review]

As a first step toward extending AutoGen to handle DP recurrences with irregular data-dependent dependencies, we design an efficient cache-oblivious parallel algorithm to solve the Viterbi recurrence. [Accepted for Euro-Par 2016]

Improving Fault Tolerance for Extreme Scale Systems

Eduardo Berrocal (Illinois Institute of Technology)

Mean Time Between Failures is expected to drop on exascale. It has been proved that combining checkpointing and failure prediction leads to longer checkpoint intervals, which in turn leads to fewer checkpoints. We present a new density-based approach for failure prediction based on the Void Search (VS) algorithm, and evaluate the algorithm using environmental logs from the Mira Blue Gene/Q supercomputer at Argonne National Laboratory. While moving to exascale, other problems will also arise as transistor size and energy consumption of future systems must be significantly reduced, steps that might dramatically impact the soft error rate (SER). When soft errors are not detected and corrected properly, either by hardware or software mechanisms, they have the potential to corrupt applications' memory state. In our previous work we leveraged the fact that datasets produced by HPC applications can be used effectively to design a general corruption detection scheme with relatively low overhead.

Cooperative Batch Scheduling for HPC Systems

Xu Yang (Illinois Institute of Technology)

The insatiable demand for computing power continues to drive the involvement of ever-growing HPC systems. The HPC systems with unprecedented scale are expected to face many challenges. Some of the most prominent challenges include energy and power, network contention and job interference, concurrency, and locality. These challenges demand great technical breakthroughs in many aspects of the HPC system's software and hardware stack. We identify three problems that we aim to solve from the perspective of batch scheduling in this research, namely, energy cost, network contention, and system fragmentation. The objective of this work is to provide new batch scheduling methodologies to solve these problems in a cooperative way.

NAM: Network Attached Memory

Juri Schmidt (University of Heidelberg)

Power and energy increasingly move into focus more, in particular, in HPC. The impact of the so-called memory wall continues to grow at a moderate rate. The idea of Processing in Memory, as one solution to overcome this issue, has been around for quite a long time. Only recently it has become technologically and economically available by heterogeneous die stacking. As one approach we have developed the Network Attached Memory (NAM). The NAM is a research vehicle to explore the possibilities of PIM in an affordable way. The first prototype is implemented as an FPGA that is connected to a Hybrid Memory Cube. Since these are two discrete devices we call it Near Data Computing. The NAM provides an interface to directly connect an EXTOLL high performance interconnection network interface controller. It is therefore a network-wide accessible and shared storage and compute node.

Doctoral Showcase Poster Reception

5:15pm-7pm
Room: Exhibit Hall E, Booth #122

The Doctoral Showcase Poster exhibition will be part of SC16's poster reception and is an opportunity for attendees to interact with the authors. The reception is open to all attendees. Complimentary refreshments and appetizers are available.

Doctoral Showcase Poster Exhibition

Tuesday, November 15: 10:00am-6pm
Wednesday, November 16: 10:00am-6pm
Thursday, November 16: 10:00am-3pm
Room: Exhibit Hall E, Exhibit #122

#DS1
Characterizing and Improving Power and Performance of HPC Networks
Taylor L. Groves (University of New Mexico)

#DS2
Performance and Energy Aware Workload Partitioning on Heterogeneous Platforms
Li Tang (University of Notre Dame)

#DS3
Dynamic Power Management for Hardware Over-Provisioned Systems
Daniel A. Ellsworth (University of Oregon)

#DS4
From Detection to Optimization: Understanding Silent Error's Impact on Scientific Applications
Jon Calhoun (University of Illinois)

#DS5
Parallel Storage Systems for Large-Scale Machines
Christos Filippidis (National and Kapodistrian University of Athens)

#DS6
High Performance File System and I/O Middleware Design for Big Data on HPC Clusters
Nusrat Islam (Ohio State University)

#DS7
Low Design-Risk Checkpointing Storage Solution for Exascale Supercomputers
Nilmini Abeyratne (University of Michigan)

#DS8
Realizing a Self-Adaptive Network Architecture for HPC Clouds
Feroz Zahid (Simula Research Laboratory)

#DS9
Automatic Discovery of Efficient Divide-and-Conquer Algorithms for Dynamic Programming Problems
Pramod Ganapathi (Stony Brook University)

#DS10
Improving Fault Tolerance for Extreme Scale Systems
Eduardo Berrocal (Illinois Institute of Technology)

#DS11
Cooperative Batch Scheduling for HPC Systems
Xu Yang (Illinois Institute of Technology)

#DS12
NAM: Network Attached Memory
Juri Schmidt (University of Heidelberg)

SC16 Salt Lake City, Utah | hpc matters.

Emerging Technologies

Tuesday, November 15
Wednesday, November 16
Thursday, November 17

10am-6pm
Room: 155-B

A Flexible Fabric for the Future of Supercomputing

John Peers (Lightfleet Corp.), Bill Dress (Lightfleet Corp.), Harold Cook (Lightfleet Corp.)

With the growth and increasing diversity of supercomputing applications have come some difficult challenges for the architects and users of the computer systems that analyze large amounts of data. Applying more processors to the problem of Big Data processing can make the processing problem worse, when bottlenecks develop in the inter-processor communication links as data are shared. Designing smarter network switches won't fix latency and reliability problems in the long run, however. Lightfleet has designed a continuously-adaptive, intelligent network fabric architecture that reconfigures itself in real-time as the processing workload changes, solving these problems. Rather than typical network architectures that are optimized for point-to-point transfers, and managed by a network control hierarchy, a more flexible architecture supports information transfers that are data-directed, requiring no control plane to intervene in setting up transfer paths. The result is a noticeable reduction in software size and complexity and dramatically-reduced latency and network congestion.

Global Energy Optimization: A Breakthrough in Overcoming the Exascale Power Wall

Jonathan Eastep (Intel Corporation)

The supercomputer industry is at risk of missing 2020 exascale performance goals by a factor of 2-3x if there are no major breakthroughs. Historical scaling techniques like improving process technology, improving system architecture, and increasing component integration will continue to apply but will be insufficient. Performance will be limited by power. To cover the gap, the industry needs a paradigm shift in the design of the power management layers in the system stack. To cover the gap, Intel is introducing a revolutionary new application-level runtime for energy optimization called GEO and making it available for free. Through online learning, GEO discovers application patterns then optimizes hardware control knobs to make applications run faster. For GEO, Intel is re-architecting aspects of x86 processor design and moving to a software-hardware co-designed solution that unlocks substantial additional performance. Our exhibit will provide an overview of GEO and share preliminary performance results with GEO.

HPC Power/Thermal/Performance Research with OpenPOWER

Todd Rosedahl (IBM), Bradley Bishop (IBM)

Reducing power/energy consumption without significantly impacting performance is a growing challenge in the HPC space. New tools and collaborative industry and academic support are required in order to further research and enable innovative solutions.

As a response to this challenge, the entire OpenPOWER firmware and software stack, including a new power/performance profiling tool, has been released as open source. System administrators and researchers can now collect over 144 power/thermal/performance sensors, add in their own sensor collection, understand exactly how the control algorithms work, investigate issues with detailed knowledge of the hardware hooks and firmware design, change the actual code at all levels of the software/firmware stack, and can provide solutions back to the HPC community at large in order to move the state of the art forward.

OpenHPC System Architect: An Open Toolkit for Building High Performance SoC's

David Donofrio (Lawrence Berkeley National Laboratory), Farzad Fatollahi-Fard (Lawrence Berkeley National Laboratory), John Leidel (Texas Tech University), Xi Wang (Texas Tech University), Yong Chen (Texas Tech University)

Given the recent difficulty in continuing the classic CMOS manufacturing density and power scaling curves, also known

as Moore's Law and Dennard Scaling, respectively, we find that modern complex system architectures are increasingly relying upon accelerators in order to optimize the placement of specific computational workloads. In addition, large-scale computing infrastructures utilized in HPC, data intensive computing and cloud computing must rely almost exclusively upon commodity device architectures provided by third-party manufacturers. The end result being a final system architecture that lacks specificity for the target software workload.

The OpenHPC System Architect infrastructure combines several open source design tools and methodologies into a central infrastructure for designing, developing and verifying the necessary hard- ware and software modules required to implement an application-specific SoC. The end result is an infrastructure the permits rapid development and deployment of of application-specific accelerators and SoC's, including an fully functional software development tool chain.

Parallware: Novel LLVM-Based Software Technology for Classification of Scientific Codes to Assist in Parallelization with OpenMP and OpenACC

Manuel Arenaz (Appentra Solutions)

HPC is a key enabling technology for the future of science, industry and society. As parallel hardware has already become a low-cost commodity, the high potential of HPC is being hindered by software issues, and porting software is one of the most significant costs in the adoption of breakthrough hardware technologies. Parallware tools offer a new parallel programming environment that helps to manage the complexity of developing parallel programs for large HPC facilities.

Parallware overcomes the limitations of the classical dependence analysis technology that is at the foundation of the tools to extract parallelism in scientific codes. Based on the production-grade LLVM compiler infrastructure, Parallware uses a fast, extensible hierarchical classification scheme to address dependence analysis. Published success stories have already shown the potential of the new technology with microbenchs and with the NAS Parallel Benchmark EP, covering fields such as finite elements, computational electromagnetics, and sparse codes.

Programming High-Performance Heterogeneous Computing Systems with the Radeon Open Compute Platform

Mayank Daga (Advanced Micro Devices Inc), Gregory Stoner (Advanced Micro Devices Inc)

The exigent demands of emerging applications to maximize performance in stringent power and thermal constraints have made heterogeneous computing ubiquitous. However, the current solutions that enable heterogeneous computing manifest several inefficiencies like (i) the lack of proper RDMA support (ii) dependence on a single programming language and (iii) an inadequate software stack. We have developed the Radeon Open Compute Platform (ROCm) which improves upon the current state of the art for GPU computing.

Some of the salient features of ROCm include enhanced multi-GPU computing with peer-sync RDMA support, a rich runtime with features like the platform atomics and user-mode queues. ROCm also provides support for several programming languages like C++, OpenCL, and Python via the Heterogeneous Compute Compiler (HCC). The basis of HCC lies in C++, the de-facto for performant programming. HCC eases the development process and provides full control of the platform to the programmer including direct-to-ISA compilation.

Reconfigurable Compile-Time Superscalar Computer

Earle Jennings (QSigma, Inc.)

Today's supercomputers use parallel processor chips that include instruction caches, superscalar instruction interpreters and multi-thread controllers, which increase size, complexity and energy consumption.

QSigma's non-von Neumann, co-designed, reconfigurable computer architecture transforms parallel programs into networks of state machines operating on typed data to perform integer, floating point and non-linear operations. This approach finds the instruction-level and thread-level parallelisms, at compile time in a supercomputer program, and implements these optimizations with a runtime target system that is reconfigured to become the supercomputer program. Instruction caches, superscalar instruction interpreters and multi-thread controllers are obsoleted, thereby removing 99% of the electronics in contemporary parallel chips. A QSigma chip, at the same number of transistors, provides 100X the computing performance of contemporary parallel processors, achieving the advantages of instruction caches, superscalar instruction processing with multi-thread control at compile time, saving the silicon and run-time energy consumption for data processing and communication.

Revolutionizing High Performance and Standard Computing Through Optical Processing Technology

Nick New (Optalysys Inc), Emma Blaylock (Optalysys Inc), Ananta Palani (Optalysys Inc)

Optalysys's groundbreaking, patented optical processing technology dramatically "turbo-charges" the speed of existing high performance and standard computers for computationally demanding tasks at a fraction of the energy consumption.

Our first device launching at the end of 2017 will be a PCIe-based optical co-processor that fits into a single unit of a rack. It is expected to be able to achieve equivalent processing rates of around 1 petaflop, but the technology has the potential to reach exascale levels by 2022. Low power laser light replaces electricity as the processing medium making the device incredibly energy efficient. The number of data points being processed can be scaled without affecting the process time, delivering truly parallel processing – an inherent feature of the optics.

The Path to Embedded Exascale

Kurt Keville (Massachusetts Institute of Technology), Anthony Skjellum (Auburn University), Mitch Williams (Sandia National Laboratories)

In June 2016, NVidia released JetPack 2.2 for the ARMv8 powered Jetson TX1. The release of unified 64 bit kernel, userspace, and CUDA 7.5 libraries significantly increased performance per watt over the previous JetPack, which was limited to 32-bits. Demonstrations on production codes and traditional benchmarks have shown the JTX1 to be on an aggressive path that will put us back on a Moore's Law trajectory as we approach the exascale era. The interplay of ARM commands, NEON, and UMA-enabled CUDA code has drastically increased embedded performance relative to their discrete analog. The authors have experimented with other TX1-based products which have had an unsupported 64-bit userspace since shortly after they shipped. Demonstrations of the latest prototypes will accompany the presentation.

The Nano Simbox: using virtual reality to interactively steer scientific simulations on high-performance computational architectures

Michael O'Connor (Interactive Scientific Ltd), Rebecca Sage (Interactive Scientific Ltd), Philip Tew (Interactive Scientific Ltd), Simon McIntosh-Smith (University of Bristol), David R. Glowacki (University of Bristol)

Developments in consumer gaming (e.g. general-purpose graphical processing units, or GP-GPUs) have played an important role in accelerating progress in scientific simulation and visualization. In this Emerging Technologies Showcase, we will demonstrate how the latest in video gaming technologies may be harnessed to accelerate HPC progress in scientific simulation. The "Nano Simbox" is a molecular research tool that exploits the latest in commodity virtual reality (VR), enabling a user to interactively steer real-time, research-grade biomolecular simulations run on a GPU-accelerated HPC back-end architecture. The simulation environment we will demo has been specifically designed for high-end commodity VR hardware like the HTC Vive, allowing users to intricately steer 3d simulations. The Nano Simbox allows users to interface with HPC capabilities in a way that has never before been possible, and opens up a powerful new paradigm for tackling the sorts of hyper-dimensional search problems often encountered in scientific simulation.

Toward Next-Generation HPC with HPE Moonshot

Aalap Tripathy (Hewlett Packard Enterprise), Sorin Cristian Cheran (Hewlett Packard Enterprise), Gerald Kleyn (Hewlett Packard Enterprise), Sai Rahul Chalamalasetti (Hewlett Packard Enterprise), Bill Whiteman (Hewlett Packard Enterprise), Mitch Wright (Hewlett Packard Enterprise)

HPE's Moonshot was created as infrastructure to host emerging application focused silicon (SoC's and accelerators). The advent of Intel Xeon E3 v3, 4, 5 processor based cartridges with integrated GPUs have made the next disruption possible. Benchmarking studies on diverse HPC workloads from molecular dynamics to life sciences show order of magnitude improvement in per core performance over staple Xeon E5. At a rack level this leads to extraordinary energy efficiency. Shared power, cooling, management infrastructure enables rapid development of cartridges of newer technology with more cores, memory speeds, storage capacity and heterogeneity. FPGAs offer the ability to escape the constraints of von Neumann architecture, the mainstay of computing over the past half-century. Moonshot is the only system in the industry today that enables FPGAs as standalone compute entities. Such heterogeneous offerings with Moonshot systems make it a future-proof investment for the next generation energy-efficient HPC.

Towards a 100% mechanical chiller free data center

Torsten Wilde (Leibniz Supercomputing Centre), Michael Ott (Leibniz Supercomputing Centre), Herbert Huber (Leibniz Supercomputing Centre), Walter Mittelbach (SorTech AG), Gregor Feig (SorTech AG)

As we move toward exascale computing, energy efficiency and power consumption is proving to be a major challenge for HPC data centers. Any power spent on cooling detracts from IT performance. After pioneering chiller-free, high-temperature, direct-liquid cooling with the SuperMUC HPC system, the Leibniz Supercomputing Centre (LRZ) demonstrates for the first time the feasibility of using adsorption chillers in a Top500 production environment to generate necessary cold water. Since adsorption chillers can produce cold water from the heat energy generated by the HPC system with little additional electrical energy, they are a very effective tool in reducing power and energy consumption. Initial efficiency data at LRZ indicates an electric Coefficient of Performance (COP) of more than 10, which is more than double compared to traditional mechanical chillers.

Towards "Write Once, Run Anywhere" HPC via Automated Translation

Paul Sathre (Virginia Polytechnic Institute and State University), Mark K. Gardner (Virginia Polytechnic Institute and State University), Wu Feng (Virginia Polytechnic Institute and State University)

Leveraging the vast array of already-written CUDA codes, we propose a "write once in CUDA, run anywhere in OpenCL" approach to enable functional portability across a spectrum of parallel computing devices, including CPUs from AMD, ARM, IBM, and Intel; GPUs from AMD, NVIDIA, and PowerVR; MIC co-processors from Intel; and FPGAs from Altera and Xilinx. This approach, encompassed by our CU2CL: Automated CUDA-to-OpenCL Source-to-Source Translator, ingests an application's CUDA source files and re-writes them into equivalent OpenCL host and kernel files. The emerging CU2CL tool has already supplanted the tedious translation of hundreds of thousands of lines of CUDA code (and 100+ CUDA-accelerated applications, including three major applications from neuroinformatics, molecular modeling, and molecular dynamics). Our Emerging Technology submission would showcase the following aspects of CU2CL: open-source release, ease of use, correctness, and instantiation across 100+ applications and many hardware devices (CPU/GPU/FPGA/MIC).

Exhibitor Forums

Tuesday, November 15

Burst Buffers and GPU's

10:30am-12pm
Room: 155-F

Best Practices for Architecting Performance and Capacity in the Burst Buffer Era

Laura Shepard (DDN Storage)

Now that Burst Buffers are commercially available, how will this game-changing technology disrupt the way high performance compute, file and storage systems are architected?

Incorporating a Burst Buffer inherently changes how you will select, procure and aggregate compute/storage/networking components in order to achieve performance and capacity goals. The positive impact is significant: - Performance and capacity that have been historically intertwined are finally bifurcated, now allowing you to scale each independently - Significant I/O, application acceleration and faster time to results can be realized - 20% or more computational time can be unlocked unleashing significant ROI - I/O bottlenecks and I/O bound applications that bring parallel file systems to their knees can be resolved without code modifications - Footprint and power planning metrics can now be substantially reduced - Exascale planning metrics can be modeled in greater, more accurate detail

Memory Hierarchies: Exposed!

Michael Wolfe (Portland Group)

This year, the Intel Xeon Phi x200 (Knights Landing) processor and NVIDIA Tesla P100 (Pascal) GPU both became available, each with an integrated high bandwidth (HBW) memory. There are several programming mechanisms for using HBW memory, ranging from special allocate routines with explicit data copies, to using directives to move data between system and HBW memory, to letting the hardware and system software automatically move pages or cache lines to the HBW memory. Future supercomputers are being designed with nonvolatile RAM, programmed either as a solid-state disk or as another level of the memory hierarchy. The days where a programmer can focus on locality and let the system manage data traffic across the memory hierarchy are vanishing. Here we explore the range of memory system designs and how the characteristics are likely to be exposed, virtualized or hidden in current and future programming models and languages.

Visualization and Deep Learning on Azure GPUs

Karan Batta (Microsoft Corporation)

Microsoft Azure's new N series virtual machines with NVIDIA GPUs opens up a range of new accelerated compute scenarios. Learn how you can take advantage of GPUs to stream content or remotely access engineering design, digital media or other graphics rich applications, and take advantage of CUDA or OpenCL accelerated applications for machine learning, analytics, and HPC applications. This session dives into the scenarios and best practices using customer case studies and partner solutions. We show you how to quickly get started and make GPU part of your HPC environment.

Tools

10:30am-12pm
Room: 155-E

Cross-Platform Automation using Open Source StackStorm

Chip Copper (Brocade)

Finding an automation tool that covers two or three domains is difficult. Finding one that covers more is nearly impossible. StackStorm provides automation across numerous platforms including vCenter, Cassandra, Microsoft Azure, Docker, Ansible, Github, Jira, Kubernetics, Google, Linux, generic email, mqtt, OpenStack, Puppet, Splunk, Yammer, MS Windows, Brocade fabrics, and more including Tesla cars. The open architecture allows one to easily build sensors, actions, and rules for specialized applications and other off-the-shelf products. In this session, we will present an overview of the tool, a few use cases, and directions on how to obtain and use this tool to make your own lives easier.

Moving, Managing and Storing Data with Data-Aware Intelligent Storage

Robert Murphy (General Atomics)

This presentation will describe General Atomics' hardware agnostic data aware storage software, Nirvana, that provides a global view of all data across an organization's multiple storage systems, including cloud storage. Nirvana implements intelligent, automated tiering across flash, parallel file systems, NAS, object and cloud storage - with comprehensive data awareness and rich analytics about the research data stored, and its contents. With Nirvana, research data can be stored and moved across multiple tiers with no vendor lock-in, reducing storage costs dramatically. Nirvana also provides comprehensive metadata extraction, creation and search capabilities through an intuitive graphical user interface, making research data easy to find, track, move, and manage throughout its lifecycle - while providing security, auditing and research data provenance across global, multi-organization collaborative workflows.

DataLogger – Data Tracking and Chain of Custody

Nate Schilling (Data In Science Technologies), Bill Pappas (Data In Science Technologies), Andrew Gauzza (Data In Science Technologies)

Providing interoperability and access to the various kinds of data in an organization can be daunting. Achieving common objectives between your HPC and non-HPC teams involves circumventing their vast environmental, technical, and cultural barriers. Opportunities to capitalize on complimentary intelligence across disciplines are rarely realized.

The DataLogger provides a comprehensive map of an organization's assets by extracting and aggregating all of the native metadata from each digital territory. A light-weight, versatile architecture allows it to be natively embedded on standard and high-performance file systems, such as GPFS. The extensibility of the platform provides attribute cataloging from any standard or proprietary file type. Reports reveal data location, creator, frequency of access and more.

Leveraging a wealth of data from the farthest reaches of your domain, the DataLogger's powerful dashboards, customizable analytics, and intuitive reports create the timely insight required for informed decision making.

Networking

3:30pm-5pm
Room: 155-F

Secure, Reliable, Congestion Free Transfer of Elephant Flows with Virtual Forwarding Contexts and Real-time Programmable Metering and QoS

Yatish Kumar (Corsa Technology)

For HPC network operators and architects, learn about a network architecture where single hardware switches can be virtualized into many internet-scale switching and routing instances that can individually offer secure self-serve "bandwidth reservation" so users can dynamically schedule and reserve bandwidth via separate class of service and meters. This bandwidth-on-demand reservation service is crucial for running reliable, fast Big Data workloads (reference this year's SCInet and also Developing Applications with Networking Capabilities via End-to-end SDN (DANCES) results presented at XSEDE 2016).

Performance SDN hardware responds to a request, sets up the appropriate circuits and securely deals with Big Data flows by only allowing passage of authorized flows. Real-time network conditions, such as traffic levels, congestion, and dynamically changing network policy, are dealt with by SDN metering and QoS which dynamically routes flows, rate-limits flows and reacts in real-time to Big Data events, such as node addition, node deletion and replication.

Evolution of 400 Gigabit Ethernet and Key Factors Driving Industry Adoption

Jeffery J. Maki (Juniper Networks)

Hyperscale data centers driving cloud computing, over the top content and 4K video have created a market inflection point. Capacity needs in the network are growing 100% year-over-year and there is a pressing need for higher density Ethernet speeds in the network. In this presentation, we explore historical trends in port speed transition from 10GE to 100GE and forecast the industry's transition to 400GE. We will also delve into key innovations that enable 400GE such as optical signaling and serdes technology, and factors crucial to the adoption such as form factors and interoperability. Finally, we will look at applications enabled by 400GbE in the supercomputing industry.

Next Generation Performance and Scalability with Mellanox Smart Interconnect

Scot Schultz (Mellanox Technologies), Gilad Shainer (Mellanox Technologies)

Pushing the frontiers of science and technology will require extreme-scale computing with machines that are 500-to-1000 times more capable than today's supercomputers. We will discuss the new era of co-design which enables all active system devices to effectively become acceleration devices, including the network. The emerging class of intelligent interconnect devices, such as Mellanox SwitchIB-2 EDR 100Gb/s InfiniBand Switching with SHArP (Scalable Hierarchical Aggregation Protocol) and ConnectX-5 100Gb/s Host Channel Adapters are now enabling in-network co-processing, and a more effective mapping of communication between devices in the system increases system performance by an order of magnitude. We will explore the benefits of the next generation capabilities of these smart-interconnect devices, off-load architecture, and in-network co-processing as we prepare to move toward HDR 200Gb/s and approach the next milestone to exascale computing.

Workflow

3:30pm-5pm
Room: 155-E

Penguin Computing: Tundra ES

Phil Pokorny (Penguin Computing)

Tundra ES: the Next Generation in High Performance and Scale-out Enterprise Computing. Tundra ES delivers the advantages of Open Computing in a single, cost-optimized, high-performance architecture. By integrating a wide variety of compute, accelerator, storage, network, software and cooling architectures in a vanity-free rack and sled solution, the solution can be customized with optimized Intel CPU, Phi, ARM or NVIDIA architectures with the latest Intel or Mellanox high-speed network technology for maximum performance. Add cloud computing options that extend the accessibility of critical HP/DCC resources with the density, cost-effectiveness and flexibility of an Open Compute architecture to provide maximum problem solving capability across an entire enterprise. Scale out with a solution that protects company investments for years to come, and includes the latest options for up to 100 Gigabit Software Defined Networks. With an array of storage options, Tundra ES combines dense compute and high-performance storage in a single rack.

Extract-Based CFD Workflows Reduce Bottlenecks and Cost

Steve M. Legensky (Intelligent Light)

Business and engineering leaders are eager to utilize High Performance Computing (HPC) to accelerate engineering while reducing design risk. Higher core counts provide faster turnaround and the ability to perform many design evaluations. However, two issues can block success from HPC: software costs and data bottlenecks.

Intelligent Light, makers of FieldView, the premier post-processing and workflow tool for CFD, have produced a solution to both of these issues by building upon DOE's VisIt, the premier visualization tool for ultrascale systems. Our VisIt Prime gives users a robust, reliable and fully supported software release backed by Intelligent Light's first class global support team. There is no per core or per user license cost for this product.

To reduce data bottlenecks, customers employ an extract database (XDB) workflow: on the HPC system, VisIt Prime extracts key information from results files or solver code memory into compact XDB files for post-processing with FieldView.

Workflow Management in Environmental Prediction

Kit Menlove (RedLine Performance Solutions, LLC), Terry McGuinness (RedLine Performance Solutions, LLC)

On the supercomputing systems at the National Centers for Environmental Prediction (NCEP), over 30,000 daily jobs are run operationally with tens of thousands more in testing and development. These jobs vary in size, complexity, and the degree to which they depend on and interact with other jobs. Two workflow management systems, NOAA's Rocoto and ECMWF's ecFlow, have been commissioned to manage these complex weather prediction suites and their time-, event-, and data-driven dependencies. Both have have proven effective in reliably managing the complexities of large-scale weather-related workflows with the former being used heavily by the research community and the latter being used primarily in the production environment. We will discuss their similarities and differences, how they approach workflow management from different angles, and the challenges of transitioning between them.

Wednesday, November 16

Processors

10:30am-12pm
Room: 155-F

Fujitsu HPC and the Development of the Post-K Supercomputer

Toshiyuki Shimizu (Fujitsu Ltd)

Fujitsu has been leading the HPC market for 40 years, and today offers a comprehensive portfolio of computing products: Fujitsu's high-end supercomputer PRIMEHPC series, as well as our x86-based PRIMERGY clusters, software, and solutions, continue to exceed wide-ranging HPC requirements.

Fujitsu also has been working together with RIKEN on Japan's next-generation supercomputer, and is currently in the detailed design phase. This "Post-K Development Project" aims to develop the successor to the K computer with 100 times more application performance. At SC16, Fujitsu will provide updates on the Post-K supercomputer development, especially on the recently announced adoption of the ARM instruction set architecture, the ARMv8-A with SVE (Scalable Vector Extension). On the development of the SVE, an extension for HPC applications, Fujitsu collaborated closely with ARM, as a "lead partner." Fujitsu will discuss the Post-K design and the benefits of the transition to this platform for the HPC community.

Unveiling NEC's Next Generation Vector Supercomputer Product

Shintaro Momose (NEC Corporation)

NEC unveils an overview of the next generation vector supercomputer product, which is a successor system to SX-ACE, our latest vector supercomputer. SX-ACE has demonstrated the world highest peak performance ratio with unmatched power-saving technologies for the HPCG benchmark program that is aimed at a more comprehensive metric for ranking HPC systems than the conventional High Performance LINPACK. The upcoming vector supercomputer will inherit these features for the enhanced performance of scientific applications. NEC is slated for launching this product with the focus on an outstanding sustained performance particularly for memory-intensive applications. It embraces a high-capability core strategy, featuring both the world top-level single-core performance and the world largest memory bandwidth per processor/core with unparalleled price competitiveness and thrifty power consumption.

Revolutionizing High Performance and Standard Computing through Optical Processing Technology

Nick New (Optalysys Ltd)

Optalysys's groundbreaking, patented optical processing technology dramatically "turbo-charges" the speed of existing high performance and standard computers for computationally demanding tasks at a fraction of the energy consumption.

Our first device launching at the end of 2017 will be a PCIe-based optical co-processor that fits into a single unit of a rack. It is expected to be able to achieve equivalent processing rates of around 1 petaFLOP but the technology has the potential to reach exascale levels by 2022. Low power laser light replaces electricity as the processing medium making the device incredibly energy efficient. The number of data points being processed can be scaled without affecting the process time, delivering truly parallel processing – an inherent feature of the optics.

Security

10:30am-12pm
Room: 155-E

Securing In-Flight Data in the Quantum Age

Chris Janson (Nokia)

Data is increasingly valuable to commerce, government and research- making networks an attractive target. Protecting data in-flight against theft and intrusion demands several mechanisms including intrusion detection, high network availability and data encryption. Encrypting at layer 1 and utilizing a strong, high quality key can yield highest threat protection while also maintaining application performance. Yet not all security solutions deliver adequate protection in the face of quantum computing based attacks. Use of symmetric encryption algorithms, shared public keys and centralized key management can ensure the highest quality effective key strength, exceeding NIST standards for classified information. This presentation will offer a comprehensive secure optical transport architecture, outline common misconceptions on encryption strength and offer methods to best ensure effective key strength.

Hybrid Cloud for HPC
Alan Settell (CSRA Inc)

CSRA has architected a hybrid cloud for HPC that offers a ubiquitous approach across a three tiered environment; performance, utility and scale tiers. The core objective is to maximize compute dollars by automatically applying workloads to the most appropriate tier and bursting to the next tier when thresholds are reached. Heavy MPI workloads are aligned to the Infiniband based performance (bare metal) and utility tiers and embarrassingly parallel workloads are aligned to either the utility tier (private cloud) or the scale tier (public cloud) if near capacity or spot pricing offers a cost advantage. Well suited to environments striving for 75% plus utilization and/or surge requirements.

HPC Security: risks, solutions and impacts
Jean-Olivier Gerphagnon (Bull)

HPC centers are tempting targets for attacks on storage capacity and computational power, and for the theft of sensitive data. The design of HPC solutions must incorporate security concerns in terms of overall architecture but also in terms of hardware and software components.

Network confinement is a classical step to achieve security. However, network is one part of the solution, and attacks such as denial-of-service or 0-day vulnerability must be managed in other layers. In addition, resource sharing brings additional complexity that must also be taken into account.

Regarding the in-depth security design of HPC infrastructures as critical must become a standard if we want to take HPC systems to a higher security level. Appropriate software architecture makes it possible to maximize system availability and reduce security breach impacts. We will present architecture and software solutions to manage common security requirements such as authentication, confinement, encryption, etc.

GPU's
3:30pm-5pm
Room: 155-E

Revolutionizing Large-Scale Heterogeneous HPC Systems with AMD's ROCm Platform
Gregory Stoner (Advanced Micro Devices Inc)

During this discussion, AMD will talk about its revolutionary approach to heterogeneous compute for HPC with their Radeon Open Compute Platform (ROCm). AMD's new approach to GPU computing is open source from the ground-up and optimized for node and rack scaling with multi-GPUs. ROCm provides users with a simplified path to GPU acceleration using main stream standard languages based on the rich single source ISO C++ 11/14 HCC compiler, which is optimized for scaling with accelerators. This talk will also cover how users can port codes developed under CUDA over to this new approach with AMD's HIP tool, providing users with an open path to choice in hardware. The stage is set with AMD's ROCm.

"Programming Massively Parallel Processors" Text and GPU Teaching Kit: New 3rd Edition
Wen-Mei Hwu (University of Illinois), Joe Bungo (NVIDIA Corporation)

Introducing the 3rd Edition of "Programming Massively Parallel Processors – a Hands-on Approach". This new edition is the result of the collaboration between GPU computing experts and contains state-of-the-art GPU programming algorithms, applications and technologies. Like the previous editions, the third edition begins with an introduction to CUDA, parallel patterns, case studies and related programming models. The 3rd edition reflects new GPU computing advances with new chapters on Deep Learning, graph search, sparse matrix computation, histogram and merge sort. It also contains significantly revised chapters on OpenACC, dynamic parallelism and scan.

The tightly-coupled GPU Teaching Kit for Accelerated Computing contains everything instructors need to teach university courses and labs with GPU parallel computing. This Teaching Kit is offered through NVIDIA's GPU Educators Program and is a collaboration between NVIDIA and the University of Illinois (UIUC).

Containers, Clouds, and Combos: Emerging Delivery Models for HPC

Doug Holt (NVIDIA Corporation), Martijn de Vries (Bright Computing)

This case study session showcases the NVIDIA Technology Center, a heterogenous computing resource and HPC cluster that leverages Bright Computing's platform to enable customers of NVIDIA's center to exploit the full computational power of GPUs. An innovative HPC delivery model – leveraging OpenStack and containers – is employed by NVIDIA to ensure that compute resources are allocated as quickly and efficiently as possible to those who need them. Join NVIDIA Solutions Architect Doug Holt and Bright Computing CTO Martijn de Vries for this lively discussion about how emerging delivery models for HPC are driving higher efficiency and productivity across clustered IT infrastructure.

Software

3:30pm-5pm
Room: 155-F

Measuring IT Success in Milliseconds on the F1 Track

Nathan Sykes (Red Bull Racing)

Formula One is an incredibly fast moving sport. There can be as little as a few hundredths of a second between first and second place. The speed of innovation is extreme – as many as 100,000 engineering changes are made to each car over the course of a season. And with Formula One regulations becoming more stringent each year, the role played by HPC for virtual analytics and simulation workloads is increasingly important. On the circuit, fast and highly accurate data-driven decision making is critical for success.

Join Nathan Sykes, Head of Numerical Tools and Technologies from Red Bull Racing to discuss the crucial role that IBM Software Defined Solutions plays – from the drawing board to race day! Learn how these solutions can help improve the competitiveness of your organization by accelerating simulation and design workloads up to 150X while controlling costs through superior resource utilization.

Approaches to Modernizing and Modularizing Fortran Codes Using Fortran 2003

Mark Potts (RedLine Performance Solutions, LLC)

The FORTRAN language is often derided by the software development community, but its use remains widespread throughout scientific computing. We will present several approaches that have been used successfully at NOAA to modularize and modernize existing Fortran code in heavy use within the weather forecasting community. Beginning with the Fortran 2003 standard, it is possible to encapsulate modules in a manner that is similar to object oriented languages such as Java and C++. This approach reduces code redundancy and ensures that shared interfaces remain synchronized by enforcing the calling protocol. We will show how it is possible to identify likely candidates for conversion and how to encapsulate modules and interfaces using both abstract and derived classes using Fortran 2003 standards without significantly changing the look or structure of the code.

Intel® HPC Orchestrator, a System Software Stack Providing Key Building Blocks of Intel® Scalable System Framework

Robert Wisniewski (Intel Corporation)

Building large scale technical computing systems is extremely challenging as HPC system software includes the integration and validation of dozens of components. Ensuring stable and reliable integration of the system software stack components is an enormous task due to multiple interdependencies and compatibility concerns.

The OpenHPC community project, and the Intel® HPC Orchestrator supported system software stack based on OpenHPC, provide a modular approach to simplifying system software stack development and maintenance. This session will demonstrate the hierarchical structure employed to allow tracking of dependencies across multiple OS, MPI, and compiler families. This design targets increasing the speed, reliability, and optimization of large scale system deployment, to deliver the Intel® Scalable System Framework.

Thursday, November 17

Exascale and Cloud

10:30am-12pm
Room: 155-E

From Today's HPC-Portfolio Toward "The Machine" and Exascale Systems

Scott Misage (Hewlett Packard Enterprise)

The talk will address HPE's portfolio of purpose-designed systems for HPC and Big Data and highlight the new trends of application-specific, modular system architectures. With the obvious end of Moore's Law and the ever increasing desire to minimize power the focus will then be on next generation HPC-system: HPE's "The Machine" is a long-term project that addresses those challenges with a design characterized as "Memory Centric Computing".

Expanding this architecture towards exascale requires system-wide non-volatile memory (NVM), ultra-efficient optical communications, and an innovative operation system with resiliency and security as key design objectives from the very beginning. Even at that scale, HPE will foster co-design, collaboration, and open standards that allow seamless communications between system components.

Navigating the Exascale Challenges with High Performance Storage

Robert Treindl (DDN Storage), Brett Costlow (Intel Corporation)

File systems are driving the high-performance workflows as HPC and verticals employing HPC-like technologies drive to Exascale. The path to Exascale will enable organizations to continue with the speed they need to realize a high level of productivity and performance to achieve the breakthroughs they seek in research, oil and gas, financial services, manufacturing and big data applications. In addition, the required massive throughput and highly parallel access are challenging traditional SAN and NAS storage technologies.

Using real-world examples, we will share specific techniques deployed in top HPC and customer sites that are helping them realize exascale speed. This discussion will focus on ways to optimize data access, storage and management at extreme scales and include examples of HPC sites.

The Convergence of Supercomputing and Big Data

Steve Scott (Cray Inc.)

Big data is changing how people use supercomputing, and supercomputing is changing how people handle big data. Traditional supercomputing users are increasingly using analytics to process the massive amounts of data coming out of large simulations and to guide simulations and experiments in real time. Analytics users are increasingly using HPC techniques to accelerate analytics applications. And we're seeing more hybrid simulation/analytics workflows.

With a more than 40-year history focused on HPC, Cray believes the convergence of supercomputing and big data is crucial to the success of the industry. We believe that big data analytics and traditional supercomputing benefit from similar system attributes, and that systems can and should be optimized to run both types of workloads.

Cray's Chief Technology Officer, Steve Scott, will share his views on how the convergence of supercomputing and big data continues to change the traditional world of supercomputing.

Storage

10:30am-12pm
Room: 155-F

HPC at Western Digital

Steve Phillpott (Western Digital Corporation)

Because of the vast quantity of mechanical components, lubricants, and internal factors involved in making HDD's, SSD's and storage arrays, Western Digital uses HPC to model and simulate all aspects of the products prior to actual manufacture, including (but not limited to):

- Model the air flow inside of an HDD for cooling purposes and particulate distribution.
- Evaluate an array of many HDDs and SSDs as found in large storage solutions for cooling.
- Model the interaction of the rotating disc(s) and the lubricant.
- Model the structural integrity of an HDD.
- Design the ASICs used in HDD and SSD's and associated solution arrays.
- Conduct "Air Bearing" simulations, allowing the HDD head and rotating media to safely interact.
- Model key microwave-assisted magnetic recording (MAMR) and heat-assisted magnetic recording (HAMR) technologies for materials selection, recording and writing sub-system designs, along with final design specifications.

Introducing Seagate ClusterStor Small Block Accelerator: Accelerate HPC Applications on Intel Lustre and IBM GPFS

Pradeep Balakrishnan (Seagate Technology LLC)

The common workload in High Performance Computing (HPC) consists of mixed block sizes from small to large. Parallel file systems such as Lustre and GPFS deliver high throughput with large block data but can become disk seek-limited with small block data.

In this talk, we will discuss the Seagate's new ClusterStor Small Block Accelerator (SBA) that uses our advanced caching technology to enhance performance for those mixed applications workloads.

We will discuss technology and share results.

The SBA/NytroXD solution suite is a collection of storage acceleration tools that combine spinning and solid state disk storage media with caching technologies to intelligently accelerate block and file-based data access. It is part of the Seagate ClusterStor Architecture that extends the functionality of the storage system's cache with SSDs and more persistent data with spinning media. This yields acceleration in the context of small blocks and also increases the throughput of the application.

Ensure Fast Access to HPC Archive Data: Both Now and Later

Donna Shawhan (Oracle Corporation)

In this session, you'll discover how to transform your HPC archive storage and address performance, management, and accessibility challenges with the latest automation advancements from Oracle. This presentation introduces the automated media migration capabilities, Memory-Assisted Copy and StorageTek Direct Copy, which accelerate HPC data movement and management. Learn how to efficiently store your HPC data, optimize performance, and ensure long-term access these latest enhancements from Oracle.

HPC Impact Showcase

Tuesday, November 15

HPC Impact Showcase I

1pm-3:30pm
Room: 155-E

Academic Supercomputing - More than just Theory

Dieter Kranzlmüller (Leibniz Supercomputing Centre)

Many HPC machines provide simulation capabilities to academic and public research organizations. Their statutes or other public regulations often seem to limit their availability to industrial and commercial users. However, there are many benefits of HPC and academic supercomputing, which help to advance the private sector into novel domains and business opportunities. This talk introduces the Leibniz Supercomputing Centre (LRZ) as part of the German Gauss Centre for Supercomputing (GCS), whose primary mission is to support science and research with computing capabilities at the highest level. Situated in the North of Munich, it supports a fruitful and diverse environment of sectors, which have been using HPC for various activities. Examples from the automotive and aerospace industry, from banking and stock exchanges, from bioinformatics, insurance, environment, and many other domains illustrate the potential collaboration and benefits between both worlds.

Using Supercomputers to Create Magic: HPC in Animation

Rasmus Tamstorf (Walt Disney Animation Studios)

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 computational challenges in an animation production system. On one hand, the sheer volume of work requires massive resources. On the other hand, it is also highly desirable to be able to leverage compute power 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 algorithms and systems. Many other problems can potentially benefit from such a transformation which makes animation and 3D graphics an exciting domain for developing HPC algorithms.

How Oil and Gas Industry is Leveraging High Performance Computing

Maxime Hugues (Total)

For several decades, Oil and Gas industry has been continuously challenged to produce more hydrocarbons in response to the growing world demand for energy. Finding new oil and gas bearing formations has become more challenging as those resources are no longer easily accessible. Hence, the industry has been pushed to explore deeper and further the subsurface that has an increased geological complexity and stronger environmental constraints.

HPC has an increasing key role in Oil and Gas to improve the effectiveness of seismic exploration and reservoir management. Highly-developed computational algorithms and more powerful computers have provided significantly a much better understanding of the distribution and description of complex geological structures, opening new frontiers to unexplored geological areas. Seismic depth imaging and reservoir simulation are the two main domains that take advantage of the rapid evolution of high performance computing. Seismic depth imaging provides invaluable and highly accurate subsurface images reducing the risk of deep and ultra-deep offshore seismic exploration. The improved computational reservoir simulation models optimize and increase the predictive capabilities and recovery rates of the subsurface assets. The growth of computing capability will open up the possibility to run higher resolution algorithms and to process richer data coming from the seismic acquisition and sensors from field and reservoir monitoring. In this presentation, we will review Total's experience with HPC and how it matters to our industry.

A CPU/GPU Implementation of Massively Parallel LES Solver Targeting Combustion Instabilities Predictions

Jin Yan (General Electric Company)

A LES code utilizing the GPU/CPU has been developed to predict the combustion inside an industry gas turbine combustor. The presentation will highlight the significant speed-up due to the utilization of the GPU/CPU and the power of high fidelity CFD.

Wednesday, November 16

HPC Impact Showcase II

1pm-3:30pm
Room: 155-E

Healthcare and Exascale

Patricia Kovatch (Icahn School of Medicine at Mount Sinai)

National U.S. efforts such as the Precision Medicine Initiative and the Cancer Moonshot aim to transform our understanding of human disease. High performance computing is already an essential element of biomedical research and clinical translational science: molecular dynamic simulations fuel the drug discovery process and computational analysis delivers results for genomic sequencing. Real-time, multi-scale cell, organ and body simulations will require not only massive computational capabilities but also new data science techniques. As computing performance slows with the end of Moore's Law, the demand for computing and data power for these initiatives will increase to exascale. Tremendous challenges await to enable these interdisciplinary efforts to better diagnose and treat human disease.

The Evolution of Supercomputing in the Creation of Computer Graphics Imagery: From DreamWorks Animation to DWA NOVA

Evan Smyth (Dreamworks)

Founded at DreamWorks Animation, DWA NOVA is commercializing the premium image generation, asset management, and versioning platform used to craft the groundbreaking imagery of DreamWorks Animation's feature films. In this talk, we provide a visual "tour" highlighting important innovations, culminating in the DreamWorks Animation Apollo Program for animation and lighting design with scalable data management. DWA NOVA builds on key components of the Apollo Program and presents them as a cloud SaaS platform for customers to exploit the benefits of CGI processing. We focus on product design, marketing, merchandizing, and sales. DWA NOVA enables a diverse range of enterprises to produce incredibly high-quality imagery as part of their internal and external business processes.

The DWA NOVA SaaS Platform enables customers to solve version-tracking challenges in complex assemblies and product lines with unparalleled clarity. Customers are also able to call up and configure virtual representations of product assets with ease as elaborate staging and assembly processes are abstracted behind the Platform APIs. This enables a large swath of a customer's business to avail itself of high-fidelity product imagery before physical products are accessible to inform earlier and better product life-cycle decision-making.

The DWA NOVA SaaS Platform will deliver the feature film rendering, animation, cloth simulation, compositing, and video encoding all scalable up to real-time with very low latency for situations that demand it. These capabilities are compose-able in the novel innovative computation assembly architecture enabling DWA NOVA to deliver scalable complex sequences of computations against strict SLAs as required by customers.

HPC and Scientific Computing with Amazon Web Services: Real-World Examples of HPC in the Cloud

Jamie Kinney (Amazon)

Amazon Web Services (AWS) gives organizations around the world access to effectively limitless compute capacity and a wide range of solutions that enable both traditional and novel HPC and high-throughput computing workloads to achieve unprecedented results. This session will describe the ways that scientists and engineers from commercial and public sector institutions have used HPC on AWS to accelerate time to results and conduct simulations and analyses at a scale that was previously inaccessible to most institutions. By examining real-world applications in the fields of high energy physics, manufacturing, computational chemistry, weather forecasting, and life sciences, we can observe how cloud computing is helping more organizations benefit from high-performance computing. Come away with a better understanding of how you can apply scalable, high-performance computing in the cloud for research and engineering, global collaboration, remote visualization, and big data analysis.

The past and future of HPC at United Technologies

Pete Bradley (Pratt & Whitney), Tim Ayer (Pratt & Whitney)

From the days when "computer" was a job description to the sophisticated 3D computational simulations of today, exponential advances in computing have ushered in revolutionary designs that have changed the way we live and work. From jet engines that deliver double digit reductions in fuel burn and greenhouse gas emissions to new technologies that combine computationally intensive simulations and big data to deliver value to our customers, High Performance Computing is an engine of innovation and growth. We will present a history of HPC and its impact at Pratt & Whitney, and a look at a future we've dubbed "HPC+", where hybrids of HPC with big data, Digital, and other systems create new opportunities for innovation and value.

Thursday, November 17th

HPC Impact Showcase III

1:30pm-3:30pm
Room: 155-E

Using machine learning to organize and contextualize the largest consumer DNA database in the world

Lei Wu (Ancestry.com)

Platforms like Facebook and LinkedIn use social graphs to help us find personal connections - however, at Ancestry, we're able to build a family history graph that reveals the complicated connections between billions of people, locations around the world, and tens of thousands of historical events.

Through the power of machine learning, Ancestry, the world's largest family history and consumer DNA database, is creating its Big Tree, a knowledge graph that stitches together 10 petabytes of structured and unstructured data from 18 billion records, more than 80 million family trees and eight billion people into one Big Tree. Utilizing the scalability of cloud computing, artificial intelligence, big data technology, AI based Search technology, and distributed stitching engines, the Big Tree is updated in real time - at a rate of 400 changes per second and 35 million changes a day - as users input data or make edits, becoming more powerful and unearthing new knowledge of familial connections with every update.

In this session, Ancestry will present an informative explanation of how the company is leveraging big data and machine learning capabilities to stitch together the largest family history graph and DNA database and how it will impact and reveal the organic relationships between people, locations and events.

The Role of High Performance Computing in Combustion Research at Cummins

John M. Deur (Cummins Inc)

In response to ever tightening environmental regulations over the past 30 years, emissions of NOx, soot, and other pollutants from Diesel engines have been reduced by 99%. Over that same period, rising fuel costs have also driven Diesel engine and truck manufacturers to significantly increase fuel efficiency as well. Future Diesel engines will continue both of these trends.

Cummins, the largest independent producer of Diesel engines in the world, has been at the forefront of these gains. To achieve them, the company has made extensive use of computational fluid dynamics to develop improved combustion system designs as part of what it calls an analysis led design (ALD) methodology. HPC has made this approach possible.

As current requirements in these areas are already at the limits of modern ALD tools, the descriptions of combustion chemistry, sprays, turbulence, and the interactions between them will all require significant improvement in the future. These improvements will require even more advanced computational resources and technologies to meet ever more demanding engine development cycles.

This presentation will describe the past, present, and future of Cummins' combustion system ALD process with examples of past successes, present challenges, and future developments being pursued. The growth in Cummins' HPC infrastructure to meet its growing computational needs will also be discussed.

HPC – A Key Enabler in Simulation Driven Design Engineering

Yoon Ho (Rolls-Royce)

In this talk, the role of High Performance Computing in Rolls-Royce simulation based engineering design will be presented. Current capability and achievements will be discussed, with future challenges identified. The use of HPC that enables many innovations in engineering product designs that otherwise would be impossible or unpractical will also be high-lighted and discussed.

Tuesday, November 15th

Different Architectures, Different Times: Reproducibility and Repeatability in High Performance Computing

10:30am-12pm
Room: 255-BC

Moderator: *Miriam Leeser (Northeastern University)*
Panelists: *Allison Baker (National Center for Atmospheric Research), Lorena Barba (George Washington University), James Demmel (University of California, Berkeley), Ganesh Gopalakrishnan (University of Utah), Michael Heroux (Sandia National Laboratories), Walid Keyrouz (National Institute of Standards and Technology), Jos Martin (Mathworks Inc.)*

The interest in reproducibility and repeatability in research has been growing. In this panel, we will address issues and approaches to help extend the adoption of reproducible research in high performance and supercomputing research. Specifically, the panelists will address the following questions: What does it mean to reproduce results in HPC? How do we get more researchers to adopt practices that will make their research reproducible? How do we make current approaches more scalable? What are the benefits and hindrances to researchers? What are good metrics for reproducibility beyond bit-for-bit equivalence?

National Strategic Computing Initiative Update

1:30pm-3pm
Room: 255-BC

Moderator: Keith Marzullo (University of Maryland)
Panelists: Steve Binkley (US Department of Energy Office of Advanced Scientific Computing Research), Gil Herrera (US Department of Defense), Irene Qualters (National Science Foundation), Erin Szulman (Office of Science and Technology Policy), Susan Gregurick (National Institutes of Health)

National Strategic Computing Initiative (NSCI) was launched in July 2015 by Presidential Executive Order to maximize the benefits of HPC research, development, and deployment for advancing science and ensuring economic competitiveness. The panel will provide an update on the NSCI to the HPC community with details about the NSCI implementation plan and allow a dialogue with key NSCI agencies and the White House. The panelists are composed of senior government personnel representing some of the key NSCI agencies and the White House. The panelists offer a long-term NSCI strategic vision and will address a broad range of policy, technology, and implementation details.

HPC Workforce Development: How Do We Find Them, Recruit Them, and Teach Them to Be Today's Practitioners and Tomorrow's Leaders?

3:30pm-5pm
Room: 255-BC

Moderator: *Henry Neeman (University of Oklahoma)*
Panelists: *Amy Apon (Clemson University), Toni Collis (EPCC at the University of Edinburgh), Sharon Broude Geva (University of Michigan), Thomas J. Lange (Technology Optimization and Management, LLC), Thomas Sterling (Indiana University), Valerie Taylor (Texas A&M University)*

This panel session is focused on gathering recommendations on mechanisms to expand the Cyberinfrastructure (CI) workforce via formal and informal education and training opportunities in CI, Computational and Data-enabled Science and Engineering (CDS&E), Data Science and related areas. The panel members have been selected to provide and elicit fresh, new, and controversial perspectives on strategies to prepare a larger and more diverse CI workforce that can advance research, discovery, scholarly studies, and economic competitiveness through the application of computational and data-enabled tools, resources, and methods across all sectors of society. The SC16 audience members will be encouraged to actively participate through oral and written communications during the session, and will be encouraged to contribute their suggestions via social media and public discussion forums after the conference ends. The panel discussions and subsequent submissions will be critical to developing a report documenting the community's recommendations, to be completed by mid-2017.

Wednesday, November 16th

Bringing About HPC Open-Standards World Peace

10:30am-12pm
Room: 255-BC

Moderator: *Andrew Richards (Codeplay Software)*
Panelists: *Alex Bourd (Qualcomm Inc), Bronis de Supinski (Lawrence Livermore National Laboratory), Martin Schulz (Lawrence Livermore National Laboratory), Greg Stoner (Advanced Micro Devices Inc), Michael Wolfe (NVIDIA Corporation), Michael Wong (Codeplay Software), Kelvin Li (IBM)*

HPC developers are bombarded with a plethora of programming models: open standards or proprietary; closed-source or open-source. Open standards enable greater code portability, however, proprietary models can be highly optimized for a specific processor core. Mixing these models can lead to incompatibility, but sometimes one standard can be built on top of another. This complexity can be a balancing act for developers, so they must clearly understand the trade-offs. How can HPC developers navigate these standards, and how can we bring about world open-standards peace?

This panel brings together members of standards bodies for OpenCL, SYCL, OpenMP, ISO C++, HSA, SPIR-V, and Vulkan. Topics will include benefits and comparisons of various standards, potential alignment between them, portability, ecosystem considerations, and more. In addition, audience members will be encouraged to give feedback to panelists on how to develop standards to be more useful for HPC software developers in the future.

Post Moore's Era Supercomputing in 20 Years

1:30pm-3pm
Room: 255-BC

Moderator: *Jeffrey S. Vetter (Oak Ridge National Laboratory)*
Panelists: *Keren Bergman (Columbia University), Tom Conte (Georgia Institute of Technology), Erik DeBenedictis (Sandia National Laboratories), Satoshi Matsuoka (Tokyo Institute of Technology), John Shalf (Lawrence Berkeley National Laboratory), George Michelogiannakis (Lawrence Berkeley National Laboratory), Jun Sawada (IBM), Matthias Troyer (ETH Zurich)*

With the approaching end of MOSFET technology scaling, an array of emerging technologies promises to preserve digital computing performance. Continuing progress of supercomputing beyond the scaling limits of Moore's Law is likely to require a comprehensive re-thinking of technologies, ranging from innovative materials and devices, circuits, system architectures, programming systems, system software, and applications. The goal of this panel is to explore the technological directions of supercomputing and in particular the range of emerging technologies over the next two decades, to prepare for this "Post Moore's Law" era. This panel will draw from the exciting results from the Post Moore's Era Supercomputing workshop (PMES) This panel will share results from the PMES workshop as well as the panelists' observations, in order to help the community set expectations for these diverse emerging technologies. Our discussion will include brief introductions by panelists, followed by a guided discussion with moderated and interactive-audience questions.

The End of Von Neumann? What the Future Looks Like for HPC Application Developers

3:30pm-5pm
Room: 255-BC

Moderator: *David Donofrio (Lawrence Berkeley National Laboratory)*
Panelists: *Gagan Gupta (Microsoft Corporation), Amir Khosrowshahi (Nervana Systems), Andrew Chien (University of Chicago), John Shalf (Lawrence Berkeley National Laboratory), Samuel Williams (Lawrence Berkeley National Laboratory)*

The slowing of Moore's Law means that in order to maintain performance scaling while maintaining or improving our current levels of energy efficiency, we must embrace new architectures and models of computation. These new architectures may depart from the traditional von Neumann model that opens a complex and multi-dimensional search space of optimizations that includes a combination of new devices, new models of computation, and specialized architectures. To focus the discussion, our panel will primarily focus on reconfigurable computing, neural inspired computing, and unconventional arrangement of architectural elements. The exploration of this search space makes for an exciting time for hardware architects and designers. However, application programmers must be involved from the beginning and throughout this transition to future computational models. This panel will discuss the evolution of computer architecture beyond the end of Moore's law and will include how these architectural shifts might impact application developers.

Thursday, November 17th

Emerging Realities and New Protocols for Power Distribution: High Voltage and DC

10:30am-12pm
Room: 255-BC

Moderator: *Anna Maria Bailey (Lawrence Livermore National Laboratory)*
Panelists: *Jack Pouchet (Emerson Network Power), Dan Stanzione (University of Texas at Austin), David Martinez (Sandia National Laboratories), Michael K. Patterson (Intel Corporation), Keiichi Hirose (Nippon Telegraph and Telephone Corporation)*

Power requirements for high-performance computers far surpass those of enterprise-servers, creating substantial loads and challenges. With enterprise racks drawing 4-15kW, their HPC counterparts can draw 40-150kW. Typical servers require 120volts, single-phase to operate, whereas high-performance-computers up the ante to 480volts, three-phase. This allows manageable wire sizes and improved efficiency. Alternating current (ac) is ubiquitous, but is direct current (dc) more efficient, reliable and less expensive? Will this trend accelerate with use of renewables? Many high-performance-computers use high-voltage dc internally; why not in the data center?

The implications for these new power distribution systems demand major shifts for both the infrastructure and operations of the data-center. Are we prepared for this shift? Will dc be compelling enough to drive the eco-system? What needs to be done to bridge the gap and ease the transition? This panel promises a lively discussion with experts in power distribution from the vendor and user community.

Virtualization, Partitioning, Cohabitation, Containers, Something Else? Approaches to Integrating HPC and Big Data Applications

1:30pm-3pm
Room: 255-BC

Moderator: *Ronald G. Minnich (Google)*
Panelists: *Ron Brightwell (Sandia National Laboratories), Rolf Riesen (Intel Corporation), Joefon Jann (IBM), Barret Rhoden (Google), Kamil Iskra (Argonne National Laboratory), Balazs Gerofi (RIKEN), Hermann Haertig (Technical University Dresden), Shane Canon (Lawrence Berkeley National Laboratory)*

In the last few years, the node OS picture for HPC and Data Intensive systems has gotten really complex. Our problems are not just pure HPC problems any more, and nodes do not just run computation-intensive code. A node must support a complex mix of computation and data processing to support activities such as visualization. From the other side, Big Data companies are now getting into computation in a big way, as Facebook, Google, and Baidu have shown.

What is the role of the node OS in this changed world? The node OS must juggle all these tasks; the main question is how. This international panel brings together representatives from national labs, academia, and industry, all working on this problem, all in different ways. We will have a stimulating discussion from the panel followed by Q&A from the audience.

HPC/Research Computing: Leveraging the Architectures, Flexibilities, and Tools Emerging from the Members of the OpenStack Scientific Community

3:30pm-5pm
Room: 255-BC

Moderator: *Bill Boas (System Fabric Works Inc)*
Panelists: *Robert Budden (Pittsburgh Supercomputing Center), Mike Lowe (Indiana University), Kate Keahey (Argonne National Laboratory), Jonathan Mills (NASA), Blair Bethwaite (Monash University), Stig Teller (Cambridge University), Paul Calleja (Cambridge University), Steve Quenette (Monash University)*

Clusters are widely used for production applications to many of the solution requirements in HPC worldwide. These clusters are usually homogeneous scalable architectures chosen by HPC systems specialists with optimal performance for a limited set of batch-scheduled codes, located to match their institutions facilities, and deployed so that they are most immediately available to co-located users, their data sources, and sinks. As HPC research and analytic computing capabilities are spreading in academia, industry, commerce, and government; users need greater flexibility, access to a wider variety of infrastructure, clusters, codes, interconnects, and filesystems with accessible yet with secure remote networking all geared to meet users' real-time-to-solution requirements. Members of the world wide OpenStack scientific computing community will debate and describe the efficacy of the methods, middleware, tools, and operation modes they are using to sustain users performance expectations and meet their time-to-solution requirements.

Friday, November 18th

Future of Memory Technology for Exascale and Beyond IV

8:30am-10am
Room: 255-BC

Moderator: Richard Murphy (Micron Technology Inc)
Panelists: Bill Dally (NVIDIA Corporation), Wendy Elasser (ARM), Mike Ignatowski (Advanced Micro Devices Inc), Doug Joseph (IBM), Peter Kogge (University of Notre Dame), Steve Pawlowski (Micron Technology Inc)

Memory technology is essential to building a successful exascale system at reasonable energy. 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 changed the relationship between the processor and memory system and the nature of the memory wall. 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 toward memory systems performance dominating compute capability. This panel will examine some of the critical exascale questions at SC13, SC14, and SC15, and will continue this discussion.

HPC File Systems, Challenges, and Use Cases

8:30am-10am
Room: 255-EF

Moderator: *Simon J. Thompson (University of Birmingham)*
Panelists: *James Coomer (Data Direct Networks), Sven Oehme (IBM), Sage Weil (Red Hat Inc), Martin Gasthuber (Deutsches Elektronen Synchrotron), Pamela Gilman (National Center for Atmospheric Research), Colin Morey (Hartree Centre)*

Not so long ago, high performance file systems for scientific applications were a technology that was understood by few, and expectations from the scientific community (aka the users) were low, the community had strong views on which filesystem was best for different workloads. High production stability required high expertise and administrator attention, and people accepted data loss as a fair possibility. Moving forward a few years, commercial offerings have become strong stable products, open source is widely accepted for enterprise solutions, and a new generation of researchers has arrived – those who grew up with Dropbox, Google Drive, Amazon's web services, and FlickR. They have much higher expectations. They expect more from a scale out storage infrastructure than the generation before. For them, capacity is normal, and durability is table stakes. With experienced technologists developing Spectrum Scale(formerly GPFS), Lustre, and Ceph, the panel will provide a unique and lively discussion.

Data Analytics Support for HPC System Management

10:30am-12pm
Room: 255-BC

Moderator: *Abani Patra (University at Buffalo)*
Panelists: *Matthew Ezell (Oak Ridge National Laboratory), William Scullin (Argonne National Laboratory), Michael Showerman (University of Illinois), Thomas Furlani (University at Buffalo), William Barth (University of Texas at Austin), James Brandt (Sandia National Laboratories)*

HPC systems are a complex combination of hardware and software. The personnel running them need to have the tools both to continuously ensure that the infrastructure is running with optimal efficiency and to proactively identify underperforming hardware and software. Furthermore, given that most HPC centers are oversubscribed, it is important that center personnel have access to appropriate technologies for monitoring all jobs that run on the cluster to determine their efficiency and resource consumption as well as to plan appropriately for future upgrades and acquisitions. In this panel, the current state-of-the-art in HPC system monitoring and management at various HPC centers, including ANL, NCSA, NNSA, ORNL, and TACC, are discussed. The discussion will focus on two use cases common to all HPC centers; (1) diagnosis of system errors (i.e., parallel file system crashes, network errors), and (2) diagnosis of poor performance or unexpected problems with user jobs.

Tuesday, November 15th

Molecular Dynamics Simulation

10:30am-12pm
Room: 355-D

The Vectorization of the Tersoff Multi-Body Potential: An Exercise in Performance Portability

Authors: *Markus Höhnerbach (RWTH Aachen University), Ahmed E. Ismail (West Virginia University), Paolo Bientinesi (RWTH Aachen University)*

Molecular dynamics simulations, an indispensable research tool in computational chemistry and materials science, consume a significant portion of the supercomputing cycles around the world. We focus on multibody potentials and aim at achieving performance portability. Compared with well-studied pair potentials, multibody potentials deliver increased simulation accuracy but are too complex for effective compiler optimization. Because of this, achieving cross-platform performance remains an open question. By abstracting from target architecture and computing precision, we develop a vectorization scheme applicable to both CPUs and accelerators. We present results for the Tersoff potential within the molecular dynamics code LAMMPS on several architectures, demonstrating efficiency gains not only for computational kernels, but also for large-scale simulations. On a cluster of Intel® Xeon Phi's, our optimized solver is between 3 and 5 times faster than the pure MPI reference.
Award: Best Student Paper Finalists

Increasing Molecular Dynamics Simulation Rates with an 8-Fold Increase in Electrical Power Efficiency

Authors: *W. Michael Brown (Intel Corporation), Andrey Semin (Intel Corporation), Michael Hebenstreit (Intel Corporation), Sergey Khvostov (Intel Corporation), Karthik Raman (Intel Corporation), Steven J. Plimpton (Sandia National Laboratories)*

Electrical power efficiency is a primary concern in designing modern HPC systems. Common strategies to improve CPU power efficiency rely on increased parallelism within a processor that is enabled both by an increase in the vector capabilities within the core and also the number of cores within a processor. Although many-core processors have been available for some time, achieving power-efficient performance has been challenging due to the offload model. Here, we evaluate performance of the molecular dynamics code LAMMPS on two new Intel® processors including the second generation many-core Intel® Xeon Phi™ processor that is available as a bootable CPU. We describe our approach to measure power consumption out-of-band and software optimizations necessary to achieve power efficiency. We analyze benefits from Intel® Advanced Vector Extensions 512 instructions and demonstrate increased simulations rates with over 9X the CPU+DRAM power efficiency when compared to the unoptimized code on previous generation processors.

Enhanced MPSM3 for Applications to Quantum Biological Simulations

Authors: *Alexander Pozdneev (IBM), Valery Weber (IBM), Teodoro Laino (IBM), Costas Bekas (IBM), Alessandro Curioni (IBM)*

Classical molecular dynamics simulations have been the preferred method to cope with the characteristic sizes and time scales of complex life-science systems. However, while classical methods have well known limitations, the practical use of far more accurate methods that rely on quantum Hamiltonians, has been limited by the current efficiency and scalability of sparse matrix-matrix multiplication algorithms used in the self-consistent field equations. In this paper, we show unprecedented massive scalability of a recently presented method, called MPSM3, for sparse matrix-matrix multiplication. Here we describe the algorithmic enhancements that allow us to claim at least one order of magnitude improvement in scalability and time to solution over the state-of-the-art. We achieve a time to solution for the multiplication of density matrices that is approaching the time needed to evaluate energy and forces with classical force-field methods and that is independent of the system size, provided proportional resources.

State-of-the-Practice: Advanced Applications Development

10:30am-12pm
Room: 255-EF

Development Effort Estimation in HPC

Authors: *Sandra Wienke (RWTH Aachen University), Julian Miller (RWTH Aachen University), Martin Schulz (Lawrence Livermore National Laboratory), Matthias S. Mueller (RWTH Aachen University)*

In order to cover the ever increasing demands for computational power, while meeting electrical power and budget constraints, HPC systems are continuing to increase in hardware and software complexity. As a direct consequence, this also leads to increased development efforts to parallelize, tune, or port applications. For an informed decision on how to spend available budgets, we therefore need quantitative metrics to estimate the development effort in HPC. While development effort estimation is widely used in software engineering, applying it to HPC, with its strong focus on performance, is not straightforward. In this paper, we first review existing approaches of effort estimation for general computing and then derive a novel methodology to estimate development effort specifically targeted at HPC. Further, we propose a concept to identify factors impacting development effort and encapsulate it in an effort log tool to collect data on development time.

MetaMorph: A Library Framework for Interoperable Kernels on Multi- and Many-Core Clusters

Authors: *Ahmed E. Helal (Virginia Polytechnic Institute and State University), Paul Sathre (Virginia Polytechnic Institute and State University), Wu-chun Feng (Virginia Polytechnic Institute and State University)*

To attain scalable performance efficiently, the HPC community expects future exascale systems to consist of multiple nodes, each with different types of hardware accelerators. In addition to GPUs and Intel MICs, additional candidate accelerators include embedded multiprocessors and FPGAs. End users need appropriate tools to efficiently use the available compute resources in such systems, both within a compute node and across compute nodes. As such, we present MetaMorph, a library framework designed to (automatically) extract as much computational capability as possible from HPC systems. Its design centers around three core principles: abstraction, interoperability, and adaptivity. To demonstrate its efficacy, we present a case study that uses the structured grids design pattern, which is heavily used in computational fluid dynamics. We show how MetaMorph significantly reduces the development time, while delivering performance and interoperability across an array of heterogeneous devices, including multicore CPUs, Intel MICs, AMD GPUs, and NVIDIA GPUs.

TrueNorth Ecosystem for Brain-Inspired Computing: Scalable Systems, Software, and Applications

Authors: *Jun Sawada (IBM), Filipp Akopyan (IBM), Andrew S. Cassidy (IBM), Brian Taba (IBM), Michael V. Debole (IBM), Pallab Datta (IBM), Rodrigo Alvarez-Icaza (IBM), Arnon Amir (IBM), John V. Arthur (IBM), Alexander Andreopoulos (IBM), Rathinakumar Appuswamy (IBM), Heinz Baier (IBM), Davis Barch (IBM), David J. Berg (IBM), Carmelo di Nolfo (IBM), Steven K. Esser (IBM), Myron Flickner (IBM), Thomas A. Horvath (IBM), Bryan L. Jackson (IBM), Jeff Kusnitz (IBM), Scott Lekuch (IBM), Michael Mastro (IBM), Timothy Melano (IBM), Paul A. Merolla (IBM), Steven E. Millman (IBM), Tapan K. Nayak (IBM), Norm Pass (IBM), Hartmut E. Penner (IBM), William P. Risk (IBM), Kai Schleupen (IBM), Benjamin Shaw (IBM), Hayley Wu (IBM), Brian Giera (Lawrence Livermore National Laboratory), Adam T. Moody (Lawrence Livermore National Laboratory), Nathan Mundhenk (Lawrence Livermore National Laboratory), Brian C. Van Essen (Lawrence Livermore National Laboratory), David P. Widemann (Lawrence Livermore National Laboratory), Qing Wu (US Air Force Research Laboratory), William E. Murphy (US Air Force Research Laboratory), Jamie K. Infantolino (US Army Research Laboratory), James A. Ross (US Army Research Laboratory), Dale R. Shires (US Army Research Laboratory), Manuel M. Vindiola (US Army Research Laboratory), Raju Namburu (US Army Research Laboratory), Dharmendra S. Modha (IBM)*

This paper describes the hardware and software ecosystem encompassing the brain-inspired TrueNorth processor – a 70mW reconfigurable silicon chip with 1 million neurons, 256 million synapses, and 4096 parallel and distributed neural cores. For systems, we present a scale-out system loosely coupling 16 single-chip boards and a scale-up system tightly integrating 16 chips in a 4x4 configuration by exploiting TrueNorth's native tiling. For software, we present an end-to-end ecosystem consisting of a simulator; a programming language; an integrated programming environment; a library of algorithms and applications; firmware; tools for deep learning; a teaching curriculum; and cloud enablement. For the scale-up systems, we summarize our approach to physical placement of neural network, to reduce intra- and inter-chip network traffic. The ecosystem is in use at over 30 universities and government/corporate labs. Our platform is a substrate for a spectrum of applications from mobile and embedded computing to cloud and supercomputers.

Systems and Networks I

10:30am-12pm
Room: 355-BC

Scheduling-Aware Routing for Supercomputers

Authors: *Jens Domke (Technical University Dresden), Torsten Hoefler (ETH Zurich)*

The interconnection network has a large influence on total cost, application performance, energy consumption, and overall system efficiency of a supercomputer. Unfortunately, today's routing algorithms do not utilize this important resource most efficiently. We first demonstrate this by defining the dark fiber metric as a measure of unused resource in networks. To improve the utilization, we propose scheduling-aware routing, a new technique that uses the current state of the batch system to determine a new set of network routes and so increases overall system utilization by up to 17.74%. We also show that our proposed routing increases the throughput of communication benchmarks by up to 17.6% on a practical InfiniBand installation. Our routing method is implemented in the standard InfiniBand tool set and can immediately be used to optimize systems. In fact, we are using it to improve the utilization of our production petascale supercomputer for more than one year.

Evaluating HPC Networks via Simulation of Parallel Workloads

Authors: *Nikhil Jain (Lawrence Livermore National Laboratory), Abhinav Bhatele (Lawrence Livermore National Laboratory), Sam White (University of Illinois), Todd Gamblin (Lawrence Livermore National Laboratory), Laxmikant V. Kale (University of Illinois)*

This paper presents an evaluation and comparison of three topologies that are popular for building interconnection networks in large-scale supercomputers: torus, fat-tree, and dragonfly. We propose a comprehensive methodology and we present a scalable packet-level network simulator, TraceR, which is used to perform this evaluation. Our methodology includes design of prototype systems that are being evaluated, use of proxy applications to determine computation and communication load, simulating individual proxy applications and multi-job workloads, and computing aggregated performance metrics. Using the proposed methodology, torus, fat-tree, and dragonfly based prototype systems with up to 730K endpoints (MPI ranks) executed on 46K nodes are compared in the context of multi-job workloads from capability and capacity systems. For the 180 Petaflop/s prototype systems simulated in the paper, we show that different topologies are superior in different scenarios.

Flexfly: Enabling a Reconfigurable Dragonfly Through Silicon Photonics

Authors: *Ke Wen (Columbia University), Payman Samadi (Columbia University), Sébastien Rumley (Columbia University), Christine P. Chen (Columbia University), Yiwen Shen (Columbia University), Meisam Bahadori (Columbia University), Jeremiah Wilke (Sandia National Laboratories), Keren Bergman (Columbia University)*

The Dragonfly topology provides low-diameter connectivity for high-performance computing with all-to-all global links at the inter-group level. Our traffic matrix characterization of various scientific applications shows consistent mismatch between the imbalanced group-to-group traffic and the uniform global bandwidth allocation of Dragonfly. Though adaptive routing has been proposed to utilize bandwidth of non-minimal paths, increased hops and cross-group interference lower efficiency. This work presents a photonic architecture, Flexfly, which "trades" global links among groups using low-radix Silicon photonic switches. With transparent optical switching, Flexfly reconfigures the inter-group topology based on traffic pattern, stealing additional direct bandwidth for communication-intensive group pairs. Simulations with applications such as GTC, Nekbone and LULESH show up to 1.8x speedup over Dragonfly paired with UGAL routing, along with halved hop count and latency for cross-group messages. We built a 32-node Flexfly prototype using a Silicon photonic switch connecting four groups and demonstrated 820 ns interconnect reconfiguration time.
Award: Best Student Paper Finalists

Numerical Algorithms I

1:30pm-3pm
Room: 355-E

PFEAST: A High Performance Sparse Eigenvalue Solver Using Distributed-Memory Linear Solvers

Authors: *James Kestyn (University of Massachusetts), Vasileios Kalantzis (University of Minnesota), Eric Polizzi (University of Massachusetts), Yousef Saad (University of Minnesota)*

The FEAST algorithm and eigensolver for interior eigenvalue problems naturally possesses three distinct levels of parallelism. The solver is then suited to exploit modern computer architectures containing many interconnected processors. This paper highlights a recent development within the software package that allows the dominant computational task, solving a set of complex linear system, to be performed with a distributed-memory solver. The software, written with a reverse-communication-interface, can now be interfaced with any generic MPI linear-system solver using a customized data

distribution for the eigenvector solutions. This work utilizes two common "black-box" distributed-memory linear-systems solvers (Cluster-MKL-Pardiso and MUMPS), as well as our own application-specific domain-decomposition MPI solver, for a collection of 3-dimensional finite-element systems. We discuss and analyze how parallel resources can be placed at all three levels simultaneously in order to achieve good scalability and optimal use of the computing platform.

Block Iterative Methods and Recycling for Improved Scalability of Linear Solvers

Authors: *Pierre Jolivet (French National Center for Scientific Research), Pierre-Henri Tournier (French Institute for Research in Computer Science and Automation)*

Contemporary large-scale partial differential equation simulations usually require the solution of large and sparse linear systems. Moreover, it is often needed to solve these linear systems with different or multiple right-hand sides. In this paper, various strategies will be presented to extend the scalability of existing multigrid or domain decomposition linear solvers using appropriate recycling strategies or block methods.

The scalability of this work is assessed by performing simulations on up to 8,192 cores for solving linear systems arising from various physical phenomena modeled by Poisson's equation, the system of linear elasticity, or Maxwell's equation. This work is shipped as part of on open-source software, readily available, and usable in any C/C++, Python, or Fortran code. In particular, some simulations are performed on top of a well established library, PETSc, and it is shown how our approaches can be used to decrease time to solution down by 30%.

Scalable Non-Blocking Preconditioned Conjugate Gradient Methods

Authors: *Paul R. Eller (University of Illinois), William Gropp (University of Illinois)*

The preconditioned conjugate gradient method (PCG) is a popular method for solving linear systems at scale. PCG requires frequent blocking allreduce collective operations that can limit performance at scale. We investigate PCG variations designed to reduce communication costs by decreasing the number of allreduces and by overlapping communication with computation using a non-blocking allreduce. These variations include two methods we have developed, non-blocking PCG and 2-step pipelined PCG, and pipelined PCG from Ghysels and Vanroose.

Performance modeling for communication and computation costs shows the expected performance of these methods. Weak and strong scaling experiments on up to 128k cores show that scalable PCG methods can outperform standard PCG at scale. We observe that the fastest method varies depending on the work per core, suggesting we need a suite of scalable solvers to obtain the best performance. Experiments with multiple preconditioners and linear systems show the robustness of these methods.

Resilience and Error Handling

1:30pm-3pm
Room: 355-BC

Pinpointing Scale-Dependent Integer Overflow Bugs in Large-Scale Parallel Applications

Authors: *Ignacio Laguna (Lawrence Livermore National Laboratory), Martin Schulz (Lawrence Livermore National Laboratory)*

We present a technique to pinpoint scale-dependent integer overflow bugs, a class of bugs in large-scale parallel applications that is hard and time-consuming to detect manually. Rather than detecting integer overflows when applications are deployed at large scale, as existing techniques do, our technique forecasts these overflows without requiring the application to be run at large scale. Our approach statically identifies integer variables that depend on the scale, and then in a refinement phase, uses data points from small-scale runs to forecast whether variables will actually overflow at large scale runs. We implement our technique in LLVM and evaluate it on several HPC benchmarks and the MPICH MPI implementation. Our tool finds five instances of previously unknown scale-dependent integer overflow bugs, including one in MPICH, and has few false positives, demonstrating its practical utility.

Compiler-Directed Lightweight Checkpointing for Fine-Grained Guaranteed Soft Error Recovery

Authors: *Qingrui Liu (Virginia Polytechnic Institute and State University), Changhee Jung (Virginia Polytechnic Institute and State University), Dongyoon Lee (Virginia Polytechnic Institute and State University), Devesh Tiwari (Oak Ridge National Laboratory)*

This paper presents Bolt, a compiler-directed soft error recovery scheme, that provides fine-grained and guaranteed recovery without excessive performance and hardware overhead. To get rid of expensive hardware support, the compiler protects the architectural inputs during their entire liveness period by safely checkpointing the last updated value in idempotent regions. To minimize the performance overhead, Bolt leverages a novel compiler analysis that eliminates those checkpoints whose value can be reconstructed by other checkpointed values without compromising the recovery

guarantee. As a result, Bolt incurs only 4.7% performance overhead on average which is up to 57% reduction compared with the state-of-the-art schemes that require expensive hardware support for the same recovery guarantee as Bolt.
Award: Best Student Paper Finalists

Understanding Error Propagation in GPGPU Applications

Authors: *Guanpeng Li (University of British Columbia), Karthik Pattabiraman (University of British Columbia), Chen-Yong Cher (IBM), Pradip Bose (IBM)*

GPUs have emerged as general-purpose accelerators in high-performance computing (HPC) and scientific applications. However, the reliability characteristics of GPU applications have not been investigated in depth. While error propagation has been extensively investigated for non-GPU applications, GPU applications have a very different programming model which can have a significant effect on error propagation in them. We perform an empirical study to understand and characterize error propagation in GPU applications. We build a compiler-based fault-injection tool for GPU applications to track error propagation, and define metrics to characterize propagation in GPU applications. We find GPU applications exhibit significant error propagation for some kinds of errors, but not others, and the behavior is highly application specific. We observe the GPU-CPU interaction boundary naturally limits error propagation in these applications compared to traditional non-GPU applications. We also formulate various guidelines for the design of fault-tolerance mechanisms in GPU applications based on our results.

Scientific Data Management and Visualization

1:30pm-3pm
Room: 355-D

Simulation and Performance Analysis of the ECMWF Tape Library System

Authors: *Markus Mäsker (Johannes Gutenberg University of Mainz), Lars Nagel (Johannes Gutenberg University of Mainz), Tim Süß (Johannes Gutenberg University of Mainz), André Brinkmann (Johannes Gutenberg University of Mainz), Lennart Sorth (European Centre for Medium-Range Weather Forecasts)*

Improvements in hardware and software have enabled magnetic disks to become an alternative to tape in backup environments. Nevertheless, even considering its slow access times, tape is still part of most hierarchical storage management strategies. The main reasons are cost-effectiveness, long lifetimes, and that tape, continually improved, keeps up with magnetic disks in terms of capacity. The performance of tape libraries, however, has been scarcely analyzed in the literature.

In this paper, we present a tape library simulator and analyze workload traces of the European Centre for Medium-Range Weather Forecasts. Optimizing the cartridge management, we show that the load latency can be reduced by a factor of 2.1 and the number of load operations by 2.5, compared to standard library settings. Furthermore, we present eviction and placement strategies, which can additionally lower the load latency by 20%, achieving the same performance level with considerably fewer drives.

Real-Time Synthesis of Compression Algorithms for Scientific Data

Authors: *Martin Burtscher (Texas State University), Hari Mukka (Texas State University), Annie Yang (Texas State University), Farbod Hesaaraki (Texas State University)*

Many scientific programs produce large amounts of floating-point data that are saved for later use. To minimize the storage requirement, it is worthwhile to compress such data as much as possible. However, existing algorithms tend to compress floating-point data relatively poorly. As a remedy, we have developed FPcrush, a tool that automatically synthesizes an optimized compressor for each given input. The synthesized algorithms are lossless and parallelized using OpenMP. This paper describes how FPcrush is able to perform this synthesis in real-time, i.e., even when accounting for the synthesis overhead, it compresses the 16 tested real-world single- and double-precision data files more quickly than parallel bzip2. Decompression is an order of magnitude faster and exceeds the throughput of multicore implementations of bzip2, gzip, and FPC. On all but two of the tested files, as well as on average, the customized algorithms deliver higher compression ratios than the other three tools.

Performance Modeling of In Situ Rendering

Authors: *Matthew Larsen (University of Oregon), Cyrus Harrison (Lawrence Livermore National Laboratory), James Kress (University of Oregon), Dave Pugmire (Oak Ridge National Laboratory), Jeremy Meredith (Oak Ridge National Laboratory), Hank Childs (University of Oregon)*

With the push to exascale, in situ visualization and analysis will continue to play an important role in HPC. Tightly coupling in situ visualization with simulations constrains resources for both, and these constraints force a complex balance of trade-offs. A performance model that provides an a priori answer for the cost of using an in situ approach for a given

task would assist in managing the trade-offs between simulation and visualization resources. In this work, we present new statistical performance models, based on algorithmic complexity, that accurately predict the run-time cost of a set of representative rendering algorithms, an essential in situ visualization task. To train and validate the models, we conduct a performance study of an MPI+X rendering infrastructure used in situ with three HPC simulation applications. We then explore feasibility issues using the model for selected in situ rendering questions.
Award: Best Paper Finalists

Resilience

Session Chair: *Christian Engelmann (Oak Ridge National Laboratory)*
3:30pm-5pm
Room: 355-D

Failure Detection and Propagation in HPC systems

Authors: *George Bosilca (University of Tennessee), Aurelien Bouteiller (University of Tennessee), Amina Guermouche (University of Tennessee), Thomas Herault (University of Tennessee), Yves Robert (ENS Lyon), Pierre Sens (LIP6 Paris), Jack Dongarra (University of Tennessee)*

Building an infrastructure for exascale applications requires, in addition to many other key components, a stable and efficient failure detector. This paper describes the design and evaluation of a robust failure detector, able to maintain and distribute the correct list of alive resources within proven and scalable bounds. The detection and distribution of the fault information follow different overlay topologies that together guarantee minimal disturbance to the applications. A virtual observation ring minimizes the overhead by allowing each node to be observed by another single node, providing an unobtrusive behavior. The propagation stage is using a non-uniform variant of a reliable broadcast over a circulant graph overlay network, and guarantees a logarithmic fault propagation. Extensive simulations, together with experiments on the ORNL Titan supercomputer, show that the algorithm performs extremely well and exhibits all the desired properties of an exascale-ready algorithm.
Award: Best Paper Finalists

Improving Application Resilience to Memory Errors with Lightweight Compression

Authors: *Scott Levy (University of New Mexico), Kurt B. Ferreira (Sandia National Laboratories), Patrick G. Bridges (University of New Mexico)*

In next-generation extreme-scale systems, application performance will be limited by memory performance characteristics. The first exascale system is projected to contain many petabytes of memory. In addition to the sheer volume of the memory required, device trends, such as shrinking feature sizes and reduced supply voltages, have the potential to increase the frequency of memory errors. As a result, resilience to memory errors is a key challenge. In this paper, we evaluate the viability of using memory compression to repair detectable uncorrectable errors (DUEs) in memory. We develop a software library, evaluate its performance, and demonstrate that it is able to significantly compress memory of HPC applications. Further, we show that exploiting compressed memory pages to correct memory errors can significantly improve application performance on next-generation systems.

FlipBack: Automatic Targeted Protection Against Silent Data Corruption

Authors: *Xiang Ni (University of Illinois), Laxmikant Kale (University of Illinois)*

The decreasing size of transistors has been critical to the increase in capacity of supercomputers. It is predicted that transistors will likely be one third of their current size by the time exascale computers are available. The smaller the transistors are, less energy is required to flip a bit, and thus silent data corruptions (SDCs) become more common. Traditional approaches to protect applications from SDCs come at the cost of either doubling the hardware resource or elongating application execution time by two times. In this paper, we present FlipBack, an automatic software based approach that protects applications from SDCs. FlipBack provides targeted protection for different types of data and calculations based on their characteristics. We evaluate FlipBack with various HPC mini-applications that capture the behavior of real scientific applications and show that FlipBack is able to fully protect applications from silent data corruptions with only 10% performance degradation.

Tensor and Graph Algorithms

3:30pm-5pm
Room: 355-BC

Graph Coloring as a Challenge Problem for Dynamic Graph Processing on Distributed Systems

Authors: *Scott Sallinen (University of British Columbia), Keita Iwabuchi (Tokyo Institute of Technology), Suraj Poudel (University of Alabama), Maya Gokhale (Lawrence Livermore National Laboratory), Matei Ripeanu (University of British Columbia), Roger Pearce (Lawrence Livermore National Laboratory)*

An unprecedented growth in data generation is taking place. Data about larger dynamic systems is being accumulated, capturing finer granularity events, and thus processing requirements are increasingly approaching real-time. To keep up, data-analytics pipelines need to be viable at massive scale and switch away from static, offline scenarios to support fully online analysis of dynamic systems.

This paper uses a challenge problem, graph coloring, to explore massive-scale analytics for dynamic graph processing. We present an event-based infrastructure, and a novel, online, distributed graph coloring algorithm. Our implementation for coloring static graphs, used as a performance baseline, is up to an order of magnitude faster than previous results and handles massive graphs with over 257 billion edges. Our framework supports dynamic graph coloring with performance at large scale better than GraphLab's static analysis. Our experience indicates that online solutions are feasible, and can be more efficient than those based on snapshotting.

An Exploration of Optimization Algorithms for High Performance Tensor Completion

Authors: *Shaden Smith (University of Minnesota), Jongsoo Park (Intel Corporation), George Karypis (University of Minnesota)*

Tensor completion is a powerful tool used to estimate or recover missing values in multi-way data. It has seen great success in domains such as product recommendation and healthcare. Tensor completion is most often accomplished via low-rank sparse tensor factorization, a non-convex optimization problem which has only recently been studied in the context of parallel computing. In this work, we study three optimization algorithms that have been successfully applied to tensor completion: alternating least squares (ALS), stochastic gradient descent (SGD), and coordinate descent (CCD++). We explore opportunities for parallelism on shared- and distributed-memory systems and address challenges such as memory- and operation-efficiency, load balance, cache locality, and communication. Among our advancements are an SGD algorithm which combines stratification with asynchronous communication, an ALS algorithm rich in level-3 BLAS routines, and a communication-efficient CCD++ algorithm. We evaluate our optimizations on a variety of real datasets and demonstrate speedups through 1024 cores.
Award: Best Student Paper Finalists

An Efficient and Scalable Algorithmic Method for Generating Large-Scale Random Graphs

Authors: *Maksudul Alam (Virginia Polytechnic Institute and State University), Maleq Khan (Virginia Polytechnic Institute and State University), Anil Vullikanti (Virginia Polytechnic Institute and State University), Madhav Marathe (Virginia Polytechnic Institute and State University)*

Many real-world systems are modeled and analyzed using various random network models. For realistic analysis, models must incorporate relevant properties such as degree distribution and clustering coefficient. Many models, such as the Chung-Lu, stochastic Kronecker, stochastic blockmodels (SBM), and block two-level Erdos-Renyi (BTER) models have been devised to capture those properties. However, the generative algorithms for these models are mostly sequential and take a prohibitively long time. In this paper, we present a novel time and space efficient algorithm for the Chung-Lu model requiring O(m) time and O(Λ) space, where m and Λ are the number of edges and distinct degrees. We also present a distributed-memory parallel algorithm with P processors requiring O(m/P+Λ+P) time and O(Λ) space. Finally, we extend our algorithms for two other popular models: SBM and BTER. These algorithms are highly scalable. Generating a power-law network with 250 billion edges takes only 12 seconds using 1024 processors.
Award: Best Paper Finalists

Topics in Distributed Computing

3:30pm-4:30pm
Room: 355-E

HARP: Predictive Transfer Optimization Based on Historical Analysis and Real-time Probing

Authors: *Engin Arslan (University at Buffalo), Kemal Guner (University at Buffalo), Tevfik Kosar (University at Buffalo)*

Increasingly data-intensive applications require frequent movement of large datasets from one site to the other. Despite the growing capacity of the networking capacity, these data movements rarely achieve the promised data transfer rates due to poorly tuned data transfer protocols. In this paper, we present predictive end-to-end data transfer optimiza-

tion algorithms based on historical data analysis and real-time background traffic probing, dubbed HARP. Most of the existing work in this area is solely based on real time network probing, which either cause too much sampling overhead or fail to accurately predict the correct transfer parameters. Combining historical data analysis with real time sampling enables our algorithms to tune the application level data transfer parameters accurately to achieve close-to-optimal end-to-end data transfer throughput with very low overhead. Our experimental analysis over a variety of network settings shows that HARP outperforms existing solutions by up to 50% in terms of achieved throughput.

SERF: Efficient Scheduling for Fast Deep Neural Network Serving via Judicious Parallelism

Authors: *Feng Yan (University of Nevada, Reno), Yuxiong He (Microsoft Research), Olatunji Ruwase (Microsoft Research), Evgenia Smirni (College of William and Mary)*

Deep neural networks (DNNs) has enabled a variety of artificial intelligence applications. These applications are backed by large DNN models running in serving mode on a cloud computing infrastructure. Given the compute-intensive nature of large DNN models, a key challenge for DNN serving systems is to minimize the request response latencies. This paper characterizes the behavior of different parallelism techniques for supporting scalable and responsive serving systems for large DNNs. We identify and model two important properties of DNN workloads: homogeneous request service demand, and interference among requests running concurrently due to cache/memory contention. These properties motivate the design of SERF, a dynamic scheduling framework that is powered by an interference-aware queueing-based analytical model. We evaluate SERF in the context of an image classification service using several well known benchmarks. The results demonstrates its accurate latency prediction and its ability to adapt to changing load conditions.

Wednesday, November 16th

Performance Measurement and Analysis

Session Chair: *Peter Beckman (Argonne National Laboratory)*
10:30am-12pm
Room: 355-BC

Understanding Performance Interference in Next-Generation HPC Systems

Authors: *Oscar H. Mondragon (University of New Mexico), Patrick G. Bridges (University of New Mexico), Kurt B. Ferreira (Sandia National Laboratories), Scott Levy (University of New Mexico), Patrick Widener (Sandia National Laboratories)*

Next-generation systems face a wide range of new potential sources of application interference, including resilience actions, system software adaptation, and in situ analytics programs. In this paper, we present a new model for analyzing the performance of bulk-synchronous HPC applications based on the use of extreme value theory. After validating this model against both synthetic and real applications, the paper then uses both simulation and modeling techniques to profile next-generation interference sources and characterize their behavior and performance impact on a selection of HPC benchmarks, mini-applications, and applications. Lastly, this work shows how the model can be used to understand how current interference mitigation techniques in multi-processors work.
Award: Best Student Paper Finalists

Reliable and Efficient Performance Monitoring in Linux

Authors: *Maria Dimakopoulou (Stanford University), Stephane Eranian (Google), Nectarios Koziris (National Technical University of Athens), Nicholas Bambos (Stanford University)*

Processor hardware performance counters have advanced in quality and features in recent years. At the same time, performance monitoring support in Linux has been significantly revamped with the development of the perf_events subsystem. These factors result in making performance monitoring a more common practice among developers. However, no performance analysis is possible without reliable hardware counter data. In this paper, we focus on a published correctness erratum in the performance counters of recent prevalent processors with Simultaneous Multithreading (SMT). This erratum causes cross-thread hardware counter corruption and may produce unreliable performance data. We propose a cache-coherence style protocol to circumvent the issue by

introducing cross-thread dynamic event scheduling. We also introduce a better scheduling algorithm that achieves the optimal scheduling of events onto hardware counters at all times. The proposed optimizations do not require any user level changes. The improvements have been contributed to the upstream Linux kernel 4.1.

Evaluating and Optimizing OpenCL Kernels for High Performance Computing with FPGAs

Authors: *Hamid Reza Zohouri (Tokyo Institute of Technology), Naoya Maruyama (RIKEN), Aaron Smith (Microsoft Corporation), Motohiko Matsuda (RIKEN), Satoshi Matsuoka (Tokyo Institute of Technology)*

We evaluate the power and performance of the Rodinia benchmark suite using the Altera SDK for OpenCL targeting a Stratix V FPGA against a modern CPU and GPU. We study multiple OpenCL kernels per benchmark, ranging from direct ports of the original GPU implementations to loop-pipelined kernels specifically optimized for FPGAs. Based on our results, we find that even though OpenCL is functionally portable across devices, direct ports of GPU-optimized code do not perform well compared to kernels optimized with FPGA-specific techniques such as sliding windows. However, by exploiting FPGA-specific optimizations, it is possible to achieve up to 3.4x better power efficiency using an Altera Stratix V FPGA in comparison to an NVIDIA K20c GPU, and better run time and power efficiency in comparison to CPU. We also present preliminary results for Arria 10, which, due to hardened FPUs, exhibits noticeably better performance compared to Stratix V in floating-point-intensive benchmarks.

Systems and Networks II

10:30am-12pm
Room: 355-D

Enhancing InfiniBand with OpenFlow-Style SDN Capability

Authors: *Jason Lee (Florida State University), Zhou Tong (Florida State University), Karthik Achalkar (Florida State University), Xin Yuan (Florida State University), Michael Lang (Los Alamos National Laboratory)*

InfiniBand is the de facto networking technology for commodity HPC clusters and has been widely deployed. However, most production large-scale InfiniBand clusters use simple routing schemes such as the destination-mod-k routing to route traffic, which may result in degraded communication performance. In this work, we investigate using the OpenFlow-style Software-Defined Networking (SDN) technology to overcome the routing deficiency in InfiniBand. We design an enhanced InfiniBand with OpenFlow-style SDN capability and demonstrate a use case that illustrates how the SDN capability can be exploited in HPC clusters to improve the system and application performance. Finally, we quantify the potential benefits of InfiniBand with OpenFlow-style SDN capability in balancing the network load by simulating job traces from production HPC clusters. The results indicate that InfiniBand with SDN capability can achieve much better network load balancing than traditional InfiniBand for HPC clusters.

Designing MPI Library with On-Demand Paging (ODP) of InfiniBand: Challenges and Benefits

Authors: *Mingzhe Li (Ohio State University), Khaled Hamidouche (Ohio State University), Xiaoyi Lu (Ohio State University), Hari Subramoni (Ohio State University), Jie Zhang (Ohio State University), Dhabaleswar K. Panda (Ohio State University)*

Existing InfiniBand drivers require the communication buffers to be “pinned” in physical memory during communication. Most runtimes leave these buffers “pinned” until the end of the run. Such situation limits the swappable memory space for applications. To address these concerns, Mellanox has recently introduced the On-Demand Paging (ODP) feature for InfiniBand. With ODP, communication buffers are paged in when they are needed by the HCA and paged out when the OS needs to swap them. This paper presents a thorough analysis of ODP and studies its performance characteristics. With these studies, we propose novel designs of ODP-aware MPI communication protocols. To the best of our knowledge, this is the first work to study and analyze the ODP feature and design an ODP-aware MPI library. Performance evaluations

with applications show that ODP-aware designs can reduce the size of pin-down buffers by 11X without performance degradation compared with the pin-down scheme.

The Mont-Blanc Prototype: An Alternative Approach for HPC Systems

Authors: *Nikola Rajovic (Barcelona Supercomputing Center), Alejandro Rico (ARM), Filippo Mantovani (Barcelona Supercomputing Center), Daniel Ruiz (Barcelona Supercomputing Center), Josep Oriol Vilarrubi (Barcelona Supercomputing Center), Constantino Gomez (Barcelona Supercomputing Center), Luna Backes (Barcelona Supercomputing Center), Diego Nieto (Barcelona Supercomputing Center), Harald Servat (Barcelona Supercomputing Center), Xavier Martorell (Barcelona Supercomputing Center), Jesus Labarta (Barcelona Supercomputing Center), Eduard Ayguade (Barcelona Supercomputing Center), Chris Adeniyi-Jones (ARM), Said Derradji (Bull), Herve Gloaguen (Bull), Piero Lanucara (CINECA), Nico Sanna (CINECA), Jean-François Méhaut (Grenoble Alpes University), Kevin Pouget (Grenoble Alpes University), Brice Videau (Grenoble Alpes University), Eric Boyer (GENCI), Momme Allalen (Leibniz Supercomputing Centre), Axel Auweter (Leibniz Supercomputing Centre), David Brayford (Leibniz Supercomputing Centre), Daniele Tafani (Leibniz Supercomputing Centre), Volker Weinberg (Leibniz Supercomputing Centre), Dirk Brömmel (Forschungszentrum Juelich), Rene Halver (Forschungszentrum Juelich), Jan H. Meinke (Forschungszentrum Juelich), Ramon Beivide (University of Cantabria), Mariano Benito (University of Cantabria), Enrique Vallejo (University of Cantabria), Mateo Valero (Barcelona Supercomputing Center), Alex Ramirez (NVIDIA Corporation)*

HPC systems are usually designed using the state-of-the-art devices. On the other side, the much larger embedded and mobile market allows for rapid development of IP blocks and provide more flexibility in designing an application-specific SoC, in turn, providing the possibility in balancing performance, energy-efficiency and cost. We advocate for alternative HPC systems to be built from such commodity IP blocks currently used in embedded and mobile SoCs. As a first demonstration of such an approach, we present the Mont-Blanc prototype; the first HPC system built with commodity SoCs, memories, and NICs from the embedded and mobile domain, and off-the-shelf HPC networking, storage, cooling, using standard integration solutions. In this paper, we present the system's architecture and evaluate both performance and energy-efficiency. Further, we compare the system's abilities against a production-level supercomputer. Finally, we discuss parallel scalability and estimate the maximum parallel scalability point of this approach.
Award: Best Paper Finalists

Compilation for Enhanced Parallelism

1:30pm-3pm
Room: 355-D

PIPES: A Language and Compiler for Task-Based Programming on Distributed-Memory Clusters

Authors: *Martin Kong (Rice University), Louis-Noel Pouchet (Ohio State University), P. Sadayappan (Ohio State University), Vivek Sarkar (Rice University)*

Applications running on clusters of shared-memory computers are often implemented using OpenMP+MPI. Productivity can be vastly improved using task-based programming, a paradigm where the user expresses the data and control-flow relations between tasks, offering the runtime maximal freedom to place and schedule tasks. While productivity is increased, high-performance execution remains challenging: the implementation of parallel algorithms typically requires specific task placement and communication strategies to reduce inter-node communications and exploit data locality.

In this work, we present a new macro-dataflow programming environment for distributed-memory clusters, based on the Intel Concurrent Collections (CnC) runtime. Our language extensions let the user define virtual topologies, task mappings, task-centric data placement, task and communication scheduling, etc. We introduce a compiler to automatically generate Intel CnC C++ run-time, with key automatic optimizations including task coarsening and coalescing. We experimentally validate our approach on a variety of scientific computations, demonstrating both productivity and performance.

A Domain-Specific Compiler for a Parallel Multiresolution Adaptive Numerical Simulation Environment

Authors: *Samyam Rajbhandari (Ohio State University), Jinsung Kim (Ohio State University), Sriram Krishnamoorthy (Pacific Northwest National Laboratory), Louis-Noel Pouchet (Ohio State University), Fabrice Rastello (French Institute for Research in Computer Science and Automation), Robert Harrison (Brookhaven National Laboratory), P. Sadayappan (Ohio State University)*

This paper describes the design and implementation of a layered domain-specific compiler to support MADNESS – Multiresolution Adaptive Numerical Environment for Scientific Simulation. MADNESS uses k-d trees to represent spatial functions and operators like addition, multiplication, differentiation, integration on the numerical representation of functions. The MADNESS runtime system provides global namespace support and a task-based execution model includ-

ing futures. MADNESS is deployed on large scale supercomputers and is used for numerical solution of differential and integral equations arising in many scientific domains. The current distribution of MADNESS is deployed on massively parallel supercomputers and has enabled many science advances. However, since all optimization is based on runtime decisions, scalability is less than desirable. This paper describes a layered domain-specific compiler that enables significant performance improvement of the MADNESS framework.

Automating Wavefront Parallelization for Sparse Matrix Codes

Authors: *Anand Venkat (University of Utah), Mahdi Soltan Mohammadi (University of Arizona), Jongsoo Park (Intel Corporation), Hongbo Rong (Intel Corporation), Rajkishore Barik (Intel Corporation), Michelle Mills Strout (University of Arizona), Mary Hall (University of Utah)*

This paper presents a compiler and runtime framework for parallelizing sparse matrix computations that have loop-carried dependences. Our approach automatically generates a runtime inspector to collect data dependence information and achieves wavefront parallelization of the computation, where iterations within a wavefront execute in parallel, and synchronization is required across wavefronts. A key contribution of this paper involves dependence simplification, which reduces the time and space overhead of the inspector. This is implemented within a polyhedral compiler framework, extended for sparse matrix codes. Results demonstrate the feasibility of using automatically-generated inspectors and executors to optimize ILU factorization and symmetric Gauss-Seidel relaxations, which are part of the Preconditioned Conjugate Gradient (PCG) computation. Our implementation achieves a median speedup of 2.97x and 2.82x over the reference sequential PCG implementation and PCG parallelized with the Intel Math Kernel Library (MKL) respectively and are within 7% of the median performance of manually tuned code.

Award: Best Paper Finalists

Fluid Dynamics

Session Chair: *Carol Woodward (Lawrence Livermore National Laboratory)*

1:30pm-3pm

Room: 255-EF

Granularity and the Cost of Error Recovery in Resilient AMR Scientific Applications

Authors: *Anshu Dubey (Argonne National Laboratory), Hajime Fujita (Intel Corporation), Daniel Graves (Lawrence Berkeley National Laboratory), Andrew Chien (University of Chicago), Devesh Tiwari (Oak Ridge National Laboratory)*

Supercomputing platforms are expected to have larger failure rates in the future because of scaling and power concerns. The memory and performance impact may vary with error types and failure modes. Therefore, localized recovery schemes will be important for scientific computations, including failure modes where application intervention is suitable for recovery. We present a resiliency methodology for applications using structured adaptive mesh refinement, where failure modes map to granularities within the application for detection and correction. This approach also enables parameterization of cost for differentiated recovery. The cost model is built with tuning parameters that can be used to customize the strategy for different failure rates in different computing environments. We also show that this approach can make recovery cost proportional to the failure rate.

Extreme Scale Plasma Turbulence Simulations on Top Supercomputers Worldwide

Authors: *William Tang (Princeton University), Bei Wang (Princeton University), Stephane Ethier (Princeton University), Grzegorz Kwasniewski (ETH Zurich), Torsten Hoefler (ETH Zurich), Khaled Ibrahim (Lawrence Berkeley National Laboratory), Kamesh Madduri (Pennsylvania State University), Samuel Williams (Lawrence Berkeley National Laboratory), Leonid Oliker (Lawrence Berkeley National Laboratory), Carlos Rosales-Fernandez (University of Texas at Austin), Timothy Williams (Argonne National Laboratory)*

The goal of the extreme scale plasma turbulence studies described in this paper is to expedite the delivery of reliable predictions on confinement physics in large magnetic fusion systems by using world-class supercomputers to carry out simulations with unprecedented resolution and temporal duration. This has involved architecture-dependent optimizations of performance scaling and addressing code portability and energy issues, with the metrics for multi-platform comparisons being “time-to-solution” and “energy-to-solution”. Realistic results addressing how confinement losses caused

by plasma turbulence scale from present-day devices to the much larger $25 billion international ITER fusion facility have been enabled by innovative advances in the GTC-P code including (i) implementation of one-sided communication from MPI 3.0 standard; (ii) optimization techniques on Xeon Phi processors; and (iii) development of a novel performance model for the key kernels of the PIC code. Results show that modeling data movement is sufficient to predict performance on modern supercomputer platforms.

A Parallel Arbitrary-Order Accurate AMR Algorithm for the Scalar Advection-Diffusion Equation

Authors: *Arash Bakhtiari (Technical University Munich), Dhairya Malhotra (University of Texas at Austin), Amir Raoofy (Technical University Munich), Miriam Mehl (University of Stuttgart), Hans-Joachim Bungartz (Technical University Munich), George Biros (University of Texas at Austin)*

We present a novel numerical method to solve the Advection-Diffusion problem. In our approach, we solve the advection equation by using an efficient semi-Lagrangian scheme with multi-algorithmic interpolation. The diffusion equation is solved by applying an integral transform: a convolution with the fundamental solution of the modified Laplace PDE. For time-marching, the Stiffly-Stable method is used. We use a dynamic spatially-adaptive distributed-memory parallelized Chebyshev octree as our data structure. The scheme is arbitrary-order accurate in space and second-order in time. To address load imbalance and communication overhead in distributed-memory Lagrangian schemes, we introduce a novel and robust partitioning scheme. With our scheme, we achieved 50% parallel efficiency for our strong scalability test up to 16k cores on the Stampede at the Texas Advanced Computing Center. Finally, as an application example, we simulate the transport of a substance in a Stokes flow in a porous medium with highly complex pore structure.

Performance Tools

Session Chair: *Lauren L. Smith (National Security Agency)*
1:30pm-3pm
Room: 355-E

MUSA: A Multi-Level Simulation Approach for Next-Generation HPC Machines

Authors: *Thomas Grass (Barcelona Supercomputing Center), César Allande (Barcelona Supercomputing Center), Adrià Armejach (Barcelona Supercomputing Center), Miquel Moretó (Barcelona Supercomputing Center), Marc Casas (Barcelona Supercomputing Center), Alejandro Rico (ARM), Eduard Ayguade (Barcelona Supercomputing Center), Jesus Labarta (Barcelona Supercomputing Center), Mateo Valero (Barcelona Supercomputing Center)*

The complexity of HPC infrastructures is growing due to a larger number of components and increasing heterogeneity. Interactions between software and hardware are not transparent to programmers and system architects. Therefore, predicting the behavior of applications on future systems is a challenging task.

In this paper, we present MUSA, an end-to-end methodology for multi-level simulation. Combining different levels of abstraction, MUSA models the communication network, microarchitectural details, and system software interactions, trading off simulation cost and accuracy. We compare detailed MUSA simulations with native executions of up to 2,048 cores and find errors to be less than 10% in the common case. Furthermore, we use MUSA to simulate up to 16,384 cores and identify scalability bottlenecks due to different factors - e.g. memory contention and load imbalance. We also compare different system configurations, showing how MUSA can help to assess the usefulness of future technologies in next-generation HPC machines.

A Machine Learning Framework for Performance Coverage Analysis of Proxy Applications

Authors: *Tanzima Z. Islam (Lawrence Livermore National Laboratory), Jayaraman J. Thiagarajan (Lawrence Livermore National Laboratory), Abhinav Bhatele (Lawrence Livermore National Laboratory), Martin Schulz (Lawrence Livermore National Laboratory), Todd Gamblin (Lawrence Livermore National Laboratory)*

Proxy applications are written to represent subsets of performance behaviors of larger, more complex, and often restricted access applications. They enable easy evaluation of these behaviors across systems, e.g., for procurement or co-design

purposes. However, the intended correlation between the performance behaviors of proxy applications and their parent codes is often based solely on the developer's intuition.

In this paper, we present novel machine-learning techniques to methodically quantify the coverage of performance behaviors of parent codes by their proxy applications. We have developed a framework, Veritas, to answer these questions in the context of on-node performance: a) which hardware resources are covered by proxy applications and how well, and b) which resources are important, but not covered. We present our techniques in the context of two benchmarks, STREAM and DGEMM, and two production applications, OpenMC and CMTnek, and their respective proxy applications.

Caliper: Performance Introspection for HPC Software Stacks

Authors: *David Boehme (Lawrence Livermore National Laboratory), Todd Gamblin (Lawrence Livermore National Laboratory), David Beckingsale (Lawrence Livermore National Laboratory), Peer-Timo Bremer (Lawrence Livermore National Laboratory), Alfredo Gimenez (University of California, Davis), Matthew LeGendre (Lawrence Livermore National Laboratory), Olga Pearce (Lawrence Livermore National Laboratory), Martin Schulz (Lawrence Livermore National Laboratory)*

Many performance engineering tasks, from long-term performance monitoring to post-mortem analysis and on-line tuning, require efficient runtime methods for introspection and performance data collection. To understand interactions between components in increasingly modular HPC software, performance introspection hooks must be integrated into runtime systems, libraries, and application codes across the software stack. This requires an interoperable, cross-stack, general-purpose approach to performance data collection, which neither application-specific performance measurement nor traditional profile or trace analysis tools provide. With Caliper, we have developed a general abstraction layer to provide performance data collection as a service to applications, runtime systems, libraries, and tools. Individual software components connect to Caliper in independent data producer, data consumer, and measurement control roles, which allows them to share performance data across software stack boundaries. We demonstrate two case studies using Caliper for performance analysis in different production scenarios.

Storage Systems

1:30pm-3pm
Room: 355-BC

Exploring the Potentials of Parallel Garbage Collection in SSDs for Enterprise Storage Systems

Authors: *Narges Shahidi (Pennsylvania State University), Mohammad Arjomand (Pennsylvania State University), Myoungsoo Jung (Yonsei University), Mahmut Kandemir (Pennsylvania State University), Chita Das (Pennsylvania State University), Anand Sivasubramaniam (Pennsylvania State University)*

SSDs have been widely adopted for high-end enterprise systems in an attempt to provide a high-performance storage. However, inferior performance is frequently attained mainly due to the need for Garbage Collection (GC). GC is a high-latency operation, and once it is scheduled for service to a block of a plane in a flash chip, it can increase latency for later arriving requests to the same plane. Apart from that, the consequent high-latency also keeps other planes of the same chip idle for a long time. This paper proposes a novel GC strategy, called Parallel GC (PaGC), whose goal is to proactively run GC on the remaining planes of a flash chip whenever any of its planes needs to execute on-demand GC. The resulting PaGC system boosts the response time of I/O requests by up to 45% (32% on average) for different GC settings and I/O workloads.

Týr: Blob Storage Meets Built-In Transactions

Authors: *Pierre Matri (Technical University of Madrid), Alexandru Costan (IRISA-INSA), Gabriel Antoniu (Inria), Jesús Montes (Technical University of Madrid), María S. Pérez (Technical University of Madrid)*

Concurrent Big Data applications often require high-performance storage, as well as ACID (Atomicity, Consistency, Isolation, Durability) transaction support. Although blobs (binary large objects) are an increasingly popular model for addressing the storage needs of such applications, state-of-the-art blob storage systems typically offer no transaction semantics. This requires users to coordinate access to data carefully in order to avoid race conditions, inconsistent writes, overwrites and other problems that cause erratic behavior. We argue that there is a gap between existing storage solutions and application requirements, which limits the design of transaction-oriented applications. We introduce Týr, the first blob storage system to provide built-in, multiblob transactions, while retaining sequential consistency and high throughput under heavy access concurrency. Týr offers fine-grained random write access to data and in-place atomic operations. Large-scale experiments on Microsoft Azure with a production application from CERN LHC show Týr throughput outperforming

state-of-the-art solutions by more than 75%.
Award: Best Student Paper Finalists

DAOS and Friends: A Proposal for an Exascale Storage System

Authors: *Jay Lofstead (Sandia National Laboratories), Ivo Jimenez (University of California, Santa Cruz), Carlos Maltzahn (University of California, Santa Cruz), Quincey Koziol (Lawrence Berkeley National Laboratory), John Bent (Seagate Technology LLC), Eric Barton (Intel Corporation)*

The DOE Extreme-Scale Technology Acceleration Fast Forward Storage and IO Stack project is going to have significant impact on storage systems design within and beyond the HPC community. With phase two of the project starting, it is an excellent opportunity to explore the complete design and how it will address the needs of extreme scale platforms. This paper examines each layer of the proposed stack in some detail along with cross-cutting topics, such as transactions and metadata management. This paper not only provides a timely summary of important aspects of the design specifications but also captures the underlying reasoning that is not available elsewhere. We encourage the broader community to understand the design, intent, and future directions to foster discussion guiding phase two and the ultimate production storage stack based on this work. An initial performance evaluation of the early prototype implementation is also provided to validate the presented design.

Accelerator Programming Tools

3:30pm-5pm
Room: 355-D

Translating OpenMP Device Constructs to OpenCL Using Unnecessary Data Transfer Elimination

Authors: *Junghyun Kim (Seoul National University), Yong-Jun Lee (Seoul National University), Jungho Park (Seoul National University), Jaejin Lee (Seoul National University)*

In this paper, we propose a framework that translates OpenMP 4.0 directives for accelerators to OpenCL. By translating an OpenMP program to an OpenCL program, the program can be executed on any hardware platform that supports OpenCL. We also propose a run-time optimization technique that automatically eliminates unnecessary data transfers between the host and the target accelerator. It exploits the page-fault mechanism to detect if a copy of the memory object already resides in the accelerator and it has not been modified by the host. To evaluate the framework, we develop 17 OpenMP 4.0 benchmark applications in two versions: basic and hand-tuned. By evaluating them on three different GPUs with the original OpenCL and OpenMP programs, we show the effectiveness of the framework. To show the practicality of the framework, we compare the performance of generated OpenCL programs with that of equivalent OpenACC programs compiled by the commercial PGI compiler.

dCUDA: Hardware Supported Overlap of Computation and Communication

Authors: *Tobias Gysi (ETH Zurich), Jeremia Bär (ETH Zurich), Torsten Hoefler (ETH Zurich)*

Over the last decade, CUDA and the underlying GPU hardware architecture have continuously gained popularity in various HPC application domains such as climate modeling, computational chemistry, and machine learning. Despite this popularity, we lack a single coherent programming model for GPU clusters. We therefore introduce the dCUDA programming model, which implements device-side remote memory access with target notification. To hide instruction pipeline latencies, CUDA programs over-decompose the problem and over-subscribe the device by running many more threads than there are hardware execution units. Whenever a thread stalls, the hardware scheduler immediately proceeds with the execution of another thread ready for execution. This latency hiding technique is key to making best use of the available hardware resources. With dCUDA, we apply latency hiding at cluster scale to automatically overlap computation and communication. Our benchmarks demonstrate perfect overlap for memory bandwidth-bound tasks and good overlap for compute-bound tasks.

Daino: A High-Level Framework for Parallel and Efficient AMR on GPUs

Authors: *Mohamed Wahib Attia (RIKEN), Naoya Maruyama (RIKEN), Takayuki Aoki (Tokyo Institute of Technology)*

Adaptive Mesh Refinement methods reduce computational requirements of problems by increasing resolution for only areas of interest. However, in practice, efficient AMR implementations are difficult considering that the mesh hierarchy management must be optimized for the underlying hardware. Architecture complexity of GPUs can render efficient AMR to be particularity challenging in GPU-accelerated supercomputers. This paper presents a compiler-based high-level framework that can automatically transform serial uniform mesh code annotated by the user into parallel adaptive mesh code optimized for GPU-accelerated supercomputers. We also present a method for empirical analysis of a uniform mesh to project an upper-bound on achievable speedup of a GPU-optimized AMR code. We show experimental results on three production applications. The speedups of code generated by our framework are comparable to hand-written AMR code

while achieving good and weak scaling up to 1000 GPUs.
Award: Best Paper Finalists

Memory and Power

Session Chair: *Rolf Riesen (Intel Corporation)*
3:30pm-5pm
Room: 355-BC

Optimizing Memory Efficiency for Deep Convolutional Neural Networks on GPUs

Authors: *Chao Li (North Carolina State University), Yang Yi (NEC Laboratories), Min Feng (NEC Laboratories), Srimat Chakradhar (NEC Laboratories), Huiyang Zhou (North Carolina State University)*

Leveraging large data sets, deep Convolutional Neural Networks (CNNs) achieve state-of-the-art recognition accuracy. Due to the substantial compute and memory operations, however, they require significant execution time. The massive parallel computing capability of GPUs make them ideal platforms to accelerate CNNs, and a number of GPU-based CNN libraries have been developed. While existing work mainly focuses on the computational efficiency of CNNs, the memory efficiency of CNNs have been largely overlooked. Yet CNNs have intricate data structures and their memory behavior can have significant impact on the performance. In this work, we study the memory efficiency of various CNN layers and reveal the performance implication from both data layouts and memory access patterns. Experiments show the universal effect of our proposed optimizations on both single layers and various networks with up to 27.9x for a single layer and up to 5.6x on the whole networks.
Award: Best Student Paper Finalists

Unprotected Computing: A Large-Scale Study of DRAM Raw Error Rate on a Supercomputer

Authors: *Leonardo Bautista-Gomez (Barcelona Supercomputing Center), Ferad Zyulkyarov (Barcelona Supercomputing Center), Simon McIntosh-Smith (University of Bristol), Osman S. Unsal (Barcelona Supercomputing Center)*

Resilience is a pressing issue to solve for extreme scale computing. DRAM errors have been analyzed in the past but little is known about errors escaping hardware checks, which lead to silent data corruption. This work attempts to fill that gap by analyzing memory errors for over a year on a cluster with about 1000 nodes featuring low-power memory without error correction. The study gathered millions of events recording thousands of errors. Temporal and spatial correlation are analyzed, but also temperature and the time of day. The study showed that most multi-bit errors corrupted non-adjacent bits in the memory word and that most errors switched from 1 to 0. In addition, we observed thousands of cases of multiple single-bit errors occurring simultaneously in different regions of the memory. We propose several directions in which the findings of this study can help the design of more reliable systems in the future.

A Data Driven Scheduling Approach for Power Management on HPC Systems

Authors: *Sean Wallace (Illinois Institute of Technology), Xu Yang (Illinois Institute of Technology), Venkatram Vishwanath (Argonne National Laboratory), William E. Allcock (Argonne National Laboratory), Susan Coghlan (Argonne National Laboratory), Michael E. Papka (Argonne National Laboratory), Zhiling Lan (Illinois Institute of Technology)*

Schedulers running on HPC systems traditionally consider the number of resources and the time requested for each job for scheduling. However as systems get larger, other metrics like power become necessary to ensure system stability.

In this paper, we propose a data driven scheduling approach for controlling the power consumption of the entire system under any user defined budget. Our approach actively observes, analyzes, and assesses power behaviors of the system and user jobs to guide scheduling decisions for power management. This design is based on the key observation that HPC jobs have distinct power profiles. This work contains an empirical analysis of workload power characteristics on a production system, dynamic learner to estimate job power profiles for scheduling, and an online power-aware scheduler for managing overall system power. Using real workload traces, we demonstrate that our design effectively controls system power consumption while minimizing the impact on system utilization.

Numerical Algorithms, Part II

3:30pm-5pm
Room: 355-E

GreenLA: Green Linear Algebra Software for GPU-Accelerated Heterogeneous Computing

Authors: *Jieyang Chen (University of California, Riverside), Li Tan (University of California, Riverside), Panruo Wu (University of California, Riverside), Dingwen Tao (University of California, Riverside), Hongbo Li (University of California, Riverside), Xin Liang (University of California, Riverside), Sihuan Li (University of California, Riverside), Rong Ge (Clemson University), Laxmi Bhuyan (University of California, Riverside), Zizhong Chen (University of California, Riverside)*

While many linear algebra libraries have been developed to optimize their performance, no linear algebra library considers their energy efficiency at the library design time. In this paper, we present GreenLA - an energy efficient linear algebra software package that leverages linear algebra algorithmic characteristics to maximize energy savings with negligible overhead. GreenLA is (1)energy efficient: it saves up to several times more energy than the best existing energy saving approaches that do not modify library source codes; (2)high performance: its performance is comparable to the highly optimized linear algebra library MAGMA; and (3)transparent to applications: with the same programming interface, existing MAGMA users do not need to modify their source codes to benefit from GreenLA. Experimental results demonstrate that GreenLA is able to save up to three times more energy than the best existing energy saving approaches while delivering similar performance compared to the state-of-the-art linear algebra library MAGMA.

Merge-Based Parallel Sparse Matrix-Vector Multiplication (SpMV)

Authors: *Duane Merrill (NVIDIA Corporation), Michael Garland (NVIDIA Corporation)*

We present a strictly balanced method for the parallel computation of sparse matrix-vector products (SpMV). Our algorithm operates directly upon the Compressed Sparse Row (CSR) sparse matrix format without preprocessing, inspection, reformatting, or supplemental encoding. Regardless of nonzero structure, our equitable 2D merge-based decomposition tightly bounds the workload assigned to each processing element. Furthermore, our technique is suitable for recursively partitioning CSR datasets themselves into multi-scale, distributed, NUMA, and GPU environments that are constrained by fixed-size local memories.

We evaluate our method on both CPU and GPU microarchitectures across a very large corpus of diverse sparse matrix datasets. We show that traditional CsrMV methods are inconsistent performers, often subject to order-of-magnitude performance variation across similarly-sized datasets. In comparison, our method provides predictable performance that is substantially uncorrelated to the distribution of nonzeros among rows and broadly improves upon that of current CsrMV methods.

Strassen's Algorithm Reloaded

Authors: *Jianyu Huang (University of Texas at Austin), Tyler M. Smith (University of Texas at Austin), Greg M. Henry (Intel Corporation), Robert A. van de Geijn (University of Texas at Austin)*

We dispel with "street wisdom" regarding the practical implementation of Strassen's algorithm for matrix-matrix multiplication (DGEMM). Conventional wisdom: it is only practical for very large matrices. Our implementation is practical for small matrices. Conventional wisdom: the matrices being multiplied should be relatively square. Our implementation is practical for rank-k updates, where k is relatively small (a shape of importance for libraries like LAPACK). Conventional wisdom: it inherently requires substantial workspace. Our implementation requires no workspace beyond buffers already incorporated into conventional high-performance DGEMM implementations. Conventional wisdom: a Strassen DGEMM interface must pass in workspace. Our implementation requires no such workspace and can be plug-compatible with the standard DGEMM interface. Conventional wisdom: it is hard to demonstrate speedup on multi-core architectures. Our implementation demonstrates speedup over conventional DGEMM even on an Intel(R) Xeon Phi(TM) coprocessor utilizing 240 threads. We show how a distributed memory matrix-matrix multiplication also benefits from these advances.

Thursday, November 17th

Data Analytics

10:30am-12pm
Room: 355-BC

Optimal Execution of Co-Analysis for Large-Scale Molecular Dynamics Simulations

Authors: *Preeti Malakar (Argonne National Laboratory), Venkatram Vishwanath (Argonne National Laboratory), Christopher Knight (Argonne National Laboratory), Todd Munson (Argonne National Laboratory), Michael E. Papka (Argonne National Laboratory)*

The analysis of scientific simulation data enables scientists to derive insights from their simulations. This analysis of the simulation output can be performed at the same execution site as the simulation using the same resources or can be done at a different site. The optimal output frequency is challenging to decide and is often chosen empirically. We propose a mathematical formulation for choosing the optimal frequency of data transfer for analysis and the feasibility of performing the analysis, under the given resource constraints such as network bandwidth, disk space, available memory, and computation time. We propose formulations for two cases of co-analysis -- local and remote. We consider various analysis features such as computation time, input data and memory requirement, importance of the analysis and minimum frequency required for performing the analysis. We demonstrate the effectiveness of our approach using molecular dynamics applications on the Mira and Edison supercomputers.

ScaleMine: Scalable Parallel Frequent Subgraph Mining in a Single Large Graph

Authors: *Ehab Abdelhamid (King Abdullah University of Science and Technology), Ibrahim Abdelaziz (King Abdullah University of Science and Technology), Panos Kalnis (King Abdullah University of Science and Technology), Zuhair Khayyat (King Abdullah University of Science and Technology), Fuad Jamour (King Abdullah University of Science and Technology)*

Frequent Subgraph Mining is an essential operation for graph analytics and knowledge extraction. Due to its high computational cost, parallel solutions are necessary. Existing approaches either suffer from load imbalance or high communication and synchronization overheads. In this paper we propose ScaleMine; a novel parallel frequent subgraph mining system for a single large graph. ScaleMine introduces a novel two-phases approach. The first phase is approximate; it quickly identifies subgraphs that are frequent with high probability, while collecting various statistics. The second phase computes the exact solution by employing the results of the approximation to achieve good load balance; prune the search space; generate efficient execution plans; and guide intra-task parallelism. Our experiments show that ScaleMine scales to 8,192 cores on a Cray XC40 (12x more than competitors); supports graphs with one billion edges (10x larger than competitors), and is at least an order of magnitude faster than existing solutions.

Efficient Delaunay Tessellation through K-D Tree Decomposition

Authors: *Dmitriy Morozov (Lawrence Berkeley National Laboratory), Tom Peterka (Argonne National Laboratory)*

Delaunay tessellations are fundamental data structures in computational geometry. They are important in data analysis, where they can represent the geometry of a point set or estimate its density. The algorithms for computing these tessellations at scale perform poorly when the input data is unbalanced. We investigate the use of k-d trees to evenly distribute points among processes, comparing two strategies for picking split points between domain regions. Because resulting point distributions no longer satisfy the assumptions of the existing algorithms, we develop a new parallel algorithm that adapts to its input and prove its correctness. We evaluate the new algorithm using two late-stage cosmology datasets. The new running times are up to 50 times faster using k-d tree compared to regular grid decomposition. In the unbalanced data sets, decomposing the domain into a k-d tree is up to five times faster than decomposing it into a regular grid.

Performance Analysis of Network Systems

Session Chair: *David Lowenthal (University of Arizona)*
10:30am-12pm
Room: 355-D

A PCIe Congestion-Aware Performance Model for Densely Populated Accelerator Servers

Authors: *Maxime Martinasso (Swiss National Supercomputing Center), Grzegorz Kwasniewski (ETH Zurich), Sadaf R. Alam (Swiss National Supercomputing Center), Thomas C. Schulthess (Swiss National Supercomputing Center), Torsten Hoefler (ETH Zurich)*

MeteoSwiss, the Swiss national weather forecast institute, has selected densely populated accelerator servers as their primary system to compute weather forecast simulation. Servers with multiple accelerator devices that are primarily connected by a PCI-Express (PCIe) network achieve a significantly higher energy efficiency. Memory transfers between accelerators in such a system are subjected to PCIe arbitration policies. In this paper, we study the impact of PCIe topology and develop a congestion-aware performance model for PCIe

communication. We present an algorithm for computing congestion factors of every communication in a congestion graph that characterizes the dynamic usage of network resources by an application. Our validation results on two different topologies of 8 GPU devices demonstrate that our model achieves an accuracy of over 97% within the PCIe network. We demonstrate the model on a weather forecast application to identify the best algorithms for its communication patterns among GPUs.

Watch Out for the Bully! Job Interference Study on Dragonfly Network

Authors: *Xu Yang (Illinois Institute of Technology), John Jenkins (Argonne National Laboratory), Misbah Mubarak (Argonne National Laboratory), Robert B. Ross (Argonne National Laboratory), Zhiling Lan (Illinois Institute of Technology)*

High-radix, low-diameter dragonfly networks will be a common choice in next-generation supercomputers. Preliminary studies show that random job placement with adaptive routing should be the rule of thumb to utilize such networks, since it uniformly distributes traffic and alleviates congestion. Nevertheless, in this work we find that while random job placement coupled with adaptive routing is good at load balancing network traffic, it cannot guarantee the best performance for every job. The performance improvement of communication-intensive applications comes at the expense of performance degradation of less intensive ones. We identify this "bully" behavior and validate its underlying causes with the help of detailed network simulation and real application traces. We further investigate a hybrid contiguous-noncontiguous job placement policy as an alternative. Initial experimentation shows that hybrid job placement aids in reducing the worst-case performance degradation for less communication-intensive applications while retaining the performance of communication-intensive ones.

Measuring and Understanding Throughput of Network Topologies

Authors: *Sangeetha Abdu Jyothi (University of Illinois), Ankit Singla (ETH Zurich), P. Brighten Godfrey (University of Illinois), Alexandra Kolla (University of Illinois)*

High throughput is of particular interest in data center and HPC networks. Although myriad network topologies have been proposed, a broad head-to-head comparison across topologies and across traffic patterns is absent, and the right way to compare worst-case throughput performance is a subtle problem. In this paper, we develop a framework to benchmark the throughput of network topologies, using a two-pronged approach. First, we study performance on a variety of synthetic and experimentally-measured traffic matrices (TMs). Second, we show how to measure worst-case throughput by generating a near-worst-case TM for any given topology. We apply the framework to study the performance of these TMs in a wide range of network topologies, revealing insights into the performance of topologies with scaling, robustness of performance across TMs, and the effect of scattered workload placement. Our evaluation code is freely available.

Combinatorial and Multigrid Algorithms

1:30pm-3pm
Room: 355-E

Designing Scalable b-Matching Algorithms on Distributed Memory Multiprocessors by Approximation

Authors: *Arif Khan (Purdue University), Alex Pothen (Purdue University), Md. Mostofa Ali Patwary (Intel Corporation), Mahantesh Halappanavar (Pacific Northwest National Laboratory), Nadathur Rajagopalan Satish (Intel Corporation), Narayanan Sundaram (Intel Corporation), Pradeep Dubey (Intel Corporation)*

A b-Matching is a subset of edges M such that at most b(v) edges in M are incident on each vertex v, where b(v) is specified. We present a distributed-memory parallel algorithm, b-Suitor, that computes a b-Matching with more than half the maximum weight in a graph with weights on the edges. The approximation algorithm is designed to have high concurrency and low time complexity. We organize the implementation of the algorithm in terms of asynchronous supersteps that combine computation and communication, and balance the computational work and frequency of communication to obtain high performance. We present several strategies to reduce the communication volume. We implement the algorithm using a hybrid strategy using MPI and OpenMP. We demonstrate strong and weak scaling of b-Suitor up to 16K processors on two supercomputers. We compute a b-Matching in a graph with 2 billion edges in under 4 seconds using 16K processors.

A Parallel Algorithm for Finding All Pairs k-Mismatch Maximal Common Substrings

Authors: *Sriram P. Chockalingam (Indian Institute of Technology Bombay), Sharma V. Thankachan (Georgia Institute of Technology), Srinivas Aluru (Georgia Institute of Technology)*

We present an efficient parallel algorithm for the following problem: Given an input collection D of n strings of total length N, a length threshold H and a mismatch threshold k, report all k-mismatch maximal common substrings of length at least H over all pairs of strings in D. This problem is moti-

vated by various applications in computational biology, where D is a collection of millions of short DNA sequences. Sequencing errors and massive size of these datasets necessitate efficient parallel approximate sequence matching algorithms. We present a novel distributed memory parallel algorithm that solves this approximate sequence matching problem in O((N/p + occ) log^{k} N) expected time and takes O(log^{k+1} N) expected rounds of global communications, under some realistic assumptions, where p is the number of processors and occ is the output size. We demonstrate the performance and scalability of our algorithm using large high throughput sequencing data sets.

Accelerating Lattice QCD Multigrid on GPUs Using Fine-Grained Parallelization

Authors: *M. A. Clark (NVIDIA Corporation), Bálint Joó (Thomas Jefferson National Accelerator Facility), Alexei Strelchenko (Fermi National Laboratory), Michael Cheng (Boston University), Arjun S. Gambhir (College of William and Mary), Richard C. Brower (Boston University)*

The past decade has witnessed a dramatic acceleration of lattice quantum chromodynamics calculations in nuclear and particle physics. This has been due to both significant progress in accelerating the iterative linear solvers using multigrid algorithms, and due to the throughput improvements brought by GPUs. Deploying hierarchical algorithms optimally on GPUs is non-trivial owing to the lack of parallelism on the coarse grids, and as such, these advances have not proved multiplicative. Using the QUDA library, we demonstrate that by exposing all sources of parallelism that the underlying stencil problem possesses, and through appropriate mapping of this parallelism to the GPU architecture, we can achieve high efficiency even for the coarsest of grids. Results are presented for the Wilson-clover discretization, where we demonstrate up to 10x speedup over present state-of-the-art GPU-accelerated methods on Titan. Finally, we look to the future and consider the software implications of our findings.

File Systems and I/O

1:30pm-3pm
Room: 355-D

An Ephemeral Burst-Buffer File System for Scientific Applications

Authors: *Teng Wang (Florida State University), Kathryn Mohror (Lawrence Livermore National Laboratory), Adam Moody (Lawrence Livermore National Laboratory), Kento Sato (Lawrence Livermore National Laboratory), Weikuan Yu (Florida State University)*

Burst buffers are becoming an indispensable hardware resource on large-scale supercomputers to buffer the bursty I/Os from scientific applications. However, there is a lack of software solutions for burst buffers to be efficiently shared by applications within a batch-submitted job and recycled across different batch jobs. In addition, burst buffers need to cope with a variety of challenging I/O patterns from data-intensive scientific applications. In this study, we have designed an ephemeral Burst-buffer based File System (BurstFS) that supports scalable and efficient aggregation of I/O bandwidth from burst buffers while having the same life cycle as a batch-submitted job. BurstFS features several techniques including scalable metadata indexing, I/O delegation, and service-side read clustering and pipelining. Through extensive tuning and analysis, we have validated that BurstFS has accomplished our design objectives, with linear scalability in terms of aggregated I/O bandwidth for parallel writes and reads.

Server-Side Log Data Analytics for I/O Workload Characterization and Coordination on Large Shared Storage Systems

Authors: *Yang Liu (North Carolina State University), Raghul Gunasekaran (Oak Ridge National Laboratory), Xiaosong Ma (Qatar Computing Research Institute), Sudharshan S. Vazhkudai (Oak Ridge National Laboratory)*

Inter-application I/O contention and performance interference has been recognized as a severe problem. In this work, we demonstrate, through measurement from the world's No. 2 supercomputer, that high I/O variance co-exists with the fact that individual storage units remain under-utilized for the majority of the time. This motivates us to propose AID, a system that performs automatic application I/O characterization and I/O-aware job scheduling. AID analyzes existing I/O traffic and batch job history logs, without any application-related prior knowledge or user/developer's involvement. It identifies the small set of I/O-intensive candidates among parallel applications running on a supercomputer, and subsequently mines their I/O patterns, using more detailed per-I/O-node traffic logs. Based on such auto-extracted information, AID provides

online I/O-aware scheduling recommendations to steer I/O-intensive applications away from heavy ongoing I/O activities. We evaluate AID on the same supercomputer, using both real applications and our own pseudo-applications.

G-Store: High-Performance Graph Store for Trillion-Edge Processing

Authors: *Pradeep Kumar (George Washington University), H. Howie Huang (George Washington University)*

High-performance graph processing brings great benefits to a wide range of scientific applications, e.g., biology networks, recommendation systems, and social networks, where the size of such graphs has grown to terabytes of data with billions of vertices and trillions of edges. Subsequently, storage and I/O performance plays a critical role in designing a high-performance computer system for graph analytics. In this paper, we present G-Store, a new graph store that incorporates three techniques to accelerate the I/O and computation of graph algorithms. We evaluate G-Store on a number of graphs against two state-of-the-art graph engines and show that G-Store achieves 2 to 8× saving in storage and outperforms X-Stream by up to 32× and Flash-Graph by up to 2.4×. G-Store is able to run different algorithms on trillion-edge graphs within tens of minutes, setting a new milestone in semi-external graph processing system.

Inverse Problems and Quantum Circuits

1:30pm-3pm
Room: 355-BC

Distributed-Memory Large Deformation Diffeomorphic 3D Image Registration

Authors: *Andreas Mang (University of Texas at Austin), Amir Gholami (University of Texas at Austin), George Biros (University of Texas at Austin)*

We present a parallel distributed-memory algorithm for large deformation diffeomorphic registration of volumetric images that produce large isochoric deformations. Finding the optimal deformation mapping requires the solution of a highly nonlinear problem that involves pseudodifferential operators, biharmonic operators, and pure advection operators. We use a preconditioned, inexact, Gauss-Newton-Krylov solver. Our algorithm integrates several components: a spectral discretization in space, a semi-Lagrangian formulation in time, analytic adjoints, different regularization functionals including volume-preserving ones, a spectral preconditioner, a highly optimized distributed FFT, and a cubic interpolation scheme for the semi-Lagrangian time-stepping. We demonstrate the scalability of our algorithm on images with resolution of up to 1024^3 on Maverick and Stampede systems. The challenging problem in the medical field is to solve the registration problem for moderate size of 256^3. We are able to compute registration for the size of 256^3 in less than five seconds on 64 x86 nodes of Maverick.

ZNNi - Maximizing the Inference Throughput of 3D Convolutional Networks on CPUs and GPUs

Authors: *Aleksandar Zlateski (Massachusetts Institute of Technology), Kisuk Lee (Massachusetts Institute of Technology), H. Sebastian Seung (Princeton University)*

Sliding window convolutional networks (ConvNets) have become a popular approach to computer vision problems such as image segmentation and object detection and localization. Here we consider the parallelization of inference, i.e., the application of a previously trained ConvNet, with emphasis on 3D images. Our goal is to maximize throughput, defined as the number of output voxels computed per unit time. We propose CPU and GPU primitives for convolutional and pooling layers, which are combined to create CPU, GPU, and CPU-GPU inference algorithms. The primitives include convolution based on highly efficient padded and pruned FFTs. Our theoretical analyses and empirical tests reveal a number of interesting findings. For example, adding host RAM can be a more efficient way of increasing throughput than adding another GPU or more CPUs. Furthermore, our CPU-GPU algorithm can achieve greater throughput than the sum of CPU-only and GPU-only throughputs.

High Performance Emulation of Quantum Circuits

Authors: *Thomas Haener (ETH Zurich), Damian S. Steiger (ETH Zurich), Mikhail Smelyanskiy (Intel Corporation), Matthias Troyer (ETH Zurich)*

As quantum computers of non-trivial size become available in the near future, it is imperative to develop tools to emulate small quantum computers. This allows for validation and debugging of algorithms as well as exploring hardware-software co-design to guide the development of quantum hardware and architectures. The simulation of quantum computers entails multiplications of sparse matrices with very large dense vectors of dimension 2^n, where n denotes the number of qubits, making this a memory-bound and network bandwidth-limited application. We introduce the concept of a quantum computer emulator as a component of a software framework for quantum computing, enabling a significant performance advantage over simulators by emulating quantum algorithms at a high level rather than simulating individual gate operations. We describe various optimization approaches and present benchmarking results, establishing the superiority of quantum computer emulators in terms of performance.

Manycore Architectures

1:30pm-3pm
Room: 255-EF

Elastic Multi-Resource Fairness: Balancing Fairness and Efficiency in Coupled CPU-GPU Architectures

Authors: *Shanjiang Tang (Tianjin University), Bingsheng He (National University of Singapore), Shuhao Zhang (National University of Singapore), Zhaojie Niu (Nanyang Technological University)*

Fairness and efficiency are two important concerns for users in a shared system, and there tends to be a tradeoff between them. Heterogeneous computing poses new challenging issues on the fair allocation of throughput among users. Prior work either considers the fair resource allocation separately for each computing device or is unable to balance the tradeoff between the fairness and system utilization. We consider Coupled CPU-GPU Architectures. We first show that it is essential to have a new fair policy that can consider both CPU and GPU as a whole. We then propose Elastic Multi-Resource Fairness(EMRF). It extends DRF by adding a knob that allows users to tune and balance fairness and performance flexibly. We show that EMRF satisfies fairness properties of sharing incentive, envy-freeness, and pareto efficiency. Finally, the experimental results show that EMRF can achieve good performance and fairness.

DCA: a DRAM-Cache-Aware DRAM Controller

Authors: *Cheng-Chieh Huang (University of Edinburgh), Vijay Nagarajan (University of Edinburgh), Arpit Joshi (University of Edinburgh)*

3D-stacking technology has enabled the option of embedding a large DRAM onto the processor. Prior studies have proposed using this as a DRAM cache. Since the DRAM cache can be orders of magnitude larger than a conventional SRAM cache, the size of the associated cache tags can also be large. Recent works have proposed storing these tags in the stacked DRAM array itself. However, this increases the complexity, with each DRAM cache request translating into multiple DRAM accesses (tag/data).

In this work, we address how to schedule these DRAM cache accesses. We start by exploring whether or not a conventional DRAM controller will work well. We introduce two potential baseline designs and study their limitations. We then derive a set of design principles that a DRAM cache controller must ideally satisfy. Our DRAM-cache-aware (DCA) DRAM controller, that is based on these principles, consistently improves performance over various DRAM cache organizations.

Enabling Efficient Preemption for SIMT Architectures with Lightweight Context Switching

Authors: *Zhen Lin (North Carolina State University), Lars Nyland (NVIDIA Corporation), Huiyang Zhou (North Carolina State University)*

Context switching is a key technique enabling preemption and time-multiplexing for CPUs. However, for single-instruction multiple-thread (SIMT) processors such as high-end graphics processing units (GPUs), it is challenging to support context switching due to the massive number of threads, which leads to a huge amount of architectural states to be swapped during context switching. Recent works present thread-block-level preemption on SIMT processors to avoid context switching overhead. However, because the execution time of a thread block (TB) is highly dependent on the kernel program. The response time of preemption cannot be guaranteed. In this paper, we propose three complementary ways to reduce and compress the architectural states to achieve lightweight context switching on SIMT processors. Based on lightweight context switching, we enable instruction-level preemption on SIMT processors with compiler and hardware co-design. With our proposed schemes, the preemption latency is reduced by 59.7% on average compared to the naive approach.

Accelerating Science

Session Chair: *Michael Bader (Technical University Munich)*
3:30pm-5pm
Room: 255-EF

High-Frequency Nonlinear Earthquake Simulations on Petascale Heterogeneous Supercomputers

Authors: *Daniel Roten (San Diego State University), Yifeng Cui (San Diego Supercomputer Center), Kim B. Olsen (San Diego State University), Steven M. Day (San Diego State University), Kyle Withers (San Diego State University), William Savran (San Diego State University), Peng Wang (NVIDIA Corporation), Dawei Mu (San Diego Supercomputer Center)*

The omission of nonlinear effects in large-scale 3D ground motion estimation, which are particularly challenging due to memory and scalability issues, can result in costly misguidance for structural design in earthquake-prone regions. We have implemented nonlinearity using a Drucker-Prager yield condition in AWP-ODC and further optimized the CUDA kernels to more efficiently utilize the GPU's memory bandwidth. The application has resulted in a significant increase in the model region and accuracy for state-of-the-art earthquake simulations in a realistic Earth structure, which are now able to resolve the wavefield at frequencies relevant for the most

vulnerable buildings (> 1 Hz) while maintaining the scalability and efficiency of the method. We successfully ran the code on 4,200 Kepler K20X GPUs on NCSA Blue Waters and OLCF Titan to simulate a M 7.7 earthquake on the southern San Andreas fault with a spatial resolution of 25 m for frequencies up to 4 Hz.

Refactoring and Optimizing the Community Atmosphere Model (CAM) on the New Sunway Many-Core Supercomputer

Authors: *Haohuan Fu (Tsinghua University), Junfeng Liao (Tsinghua University), Wei Xue (Tsinghua University), Lanning Wang (Beijing Normal University), Dexun Chen (Tsinghua University), Long Gu (National Research Center of Parallel Computer Engineering and Technology), Jinxiu Xu (National Research Center of Parallel Computer Engineering and Technology), Nan Ding (Tsinghua University), Xinliang Wang (Tsinghua University), Conghui He (Tsinghua University), Shizhen Xu (Tsinghua University), Yishuang Liang (Beijing Normal University), Jiarui Fang (Tsinghua University), Yuanchao Xu (Tsinghua University), Weijie Zheng (Tsinghua University), Jingheng Xu (Tsinghua University), Zhen Zheng (Tsinghua University), Wanjing Wei (Tsinghua University), Xu Ji (Tsinghua University), He Zhang (Tsinghua University), Bingwei Chen (Tsinghua University), Kaiwei Li (Tsinghua University), Xiaomeng Huang (Tsinghua University), Wenguang Chen (Tsinghua University), Guangwen Yang (Tsinghua University)*

This paper reports our efforts in refactoring and optimizing the Community Atmosphere Model(CAM) on the Sunway TaihuLight supercomputer, which uses many-core processors that consist of management processing elements (MPEs) and clusters of computing processing elements (CPEs). To map the large code base of CAM to the millions of cores on the Sunway system, we take OpenACC-based refactoring as the major approach, and apply source-to-source translator tools to exploit the most suitable parallelism for the CPE cluster, and to fit the intermediate variable into the limited on-chip fast buffer. For individual kernels, when comparing the original ported version using only MPEs and the refactored version using both the MPE and CPE clusters, we achieve up to 22x speedup for the compute-intensive kernels. For the 25km resolution CAM global model, we manage to scale to 24,000 MPEs and 1,536,000 CPEs, and we achieve a simulation speed of 2.81 model years per day.

LIBXSMM: Accelerating Small Matrix Multiplications by Runtime Code Generation

Authors: *Alexander Heinecke (Intel Corporation), Greg Henry (Intel Corporation), Maxwell Hutchinson (University of Chicago), Hans Pabst (Intel Corporation)*

Many modern highly scalable scientific simulations packages rely on small matrix multiplications as their main computational engine. Math libraries or compilers are unlikely to provide the best possible kernel performance. To address this issue, we present a library which provides high performance small matrix multiplications targeting all recent x86 vector instruction set extensions up to Intel AVX-512. Our evaluation proves that speed-ups of more than 10X are possible depending on the CPU and application. These speed-ups are achieved by a combination of several novel technologies. We use a code generator which has a built-in architectural model to create code which runs best without requiring an auto-tuning phase. Since such code is very specialized, we leverage just-in-time compilation to only build the required kernel variant at runtime. To keep ease-of-use, overhead, and kernel management under control we accompany our library with a BLAS-compliant frontend which features a multi-level code-cache hierarchy.

Clouds & Job Scheduling

3:30pm-5pm
Room: 355-BC

Transient Guarantees: Maximizing the Value of Idle Cloud Capacity

Authors: *Supreeth Shastri (University of Massachusetts), Amr Rizk (University of Massachusetts), David Irwin (University of Massachusetts)*

To reduce waste, platforms have begun to offer idle capacity in the form of transient servers, which they may unilaterally revoke, for much lower prices—~70-90% less—than on-demand servers. However, transient servers' revocation characteristics—their volatility and predictability—influence their performance, since they affect the overhead of fault-tolerance mechanisms applications use to handle revocations. Unfortunately, current cloud platforms offer no guarantees on revocation characteristics, which makes it difficult for users to optimally configure (and value) transient servers.

To address the problem, we propose the abstraction of a transient guarantee, which offers probabilistic assurances on revocation characteristics. We present policies for partitioning a variable amount of idle capacity into classes with different transient guarantees to maximize performance and value. We then implement and evaluate these policies on job traces

from a production Google cluster. We show that our approach can increase the aggregate revenue from idle server capacity by ~6.5× compared to existing approaches.

Multi-Resource Fair Sharing for Datacenter Jobs with Placement Constraints

Authors: *Wei Wang (Hong Kong University of Science and Technology), Baochun Li (University of Toronto), Ben Liang (University of Toronto), Jun Li (University of Toronto)*

Providing quality-of-service guarantees by means of fair sharing has never been more challenging in datacenters. Due to the heterogeneity of physical machines, datacenter jobs frequently specify placement constraints, restricting them to run on a particular class of machines meeting specific hardware/software requirements. In addition, jobs have diverse demands across multiple resource types, and may saturate any of CPU, memory, or network. Despite the rich body of recent work on datacenter scheduling, it remains unclear how multi-resource fair sharing is defined and achieved for jobs with placement constraints. We propose a new sharing policy, called Task Share Fairness (TSF). With TSF, jobs are better off sharing the datacenter, and are better off reporting demands and constraints truthfully. We have prototyped TSF on Apache Mesos and confirmed its service guarantees in a 50-node EC2 cluster. Trace-driven simulations have further revealed that TSF speeds up 60% of tasks over existing fair schedulers.

A Multi-Faceted Approach to Job Placement for Improved Performance on Extreme-Scale Systems

Authors: *Christopher Zimmer (Oak Ridge National Laboratory), Saurabh Gupta (Oak Ridge National Laboratory), Scott Atchley (Oak Ridge National Laboratory), Sudharshan Vazhkudai (Oak Ridge National Laboratory), Carl Albing (US Naval Academy)*

Job placement plays a pivotal role in application performance on supercomputers. We present a multi-faceted exploration to influence placement in extreme-scale systems, to improve network performance and decrease variability. In our first exploration, Scores, we developed a machine learning model that extracts features from a job's node-allocation and grades performance. This identified several important node-metrics that led to Dual-Ended scheduling, a means of reducing network contention without impacting utilization. In evaluations on the Titan supercomputer, we observed reductions in average hop-count by up to 50%. We also developed an improved node-layout strategy that targets a better balance between network latency and bandwidth, replacing the default ALPS layout on Titan that resulted in an average of 10% runtime improvement. Both of these efforts underscore the importance of a job placement strategy that is cognizant of workload mixture and network topology.

State-of-the-Practice: System Characterization and Design

3:30pm-4:30pm
Room: 355-D

Characterizing Parallel Scientific Applications on Commodity Clusters: An Empirical Study of a Tapered Fat-Tree

Authors: *Edgar A. Leon (Lawrence Livermore National Laboratory), Ian Karlin (Lawrence Livermore National Laboratory), Abhinav Bhatele (Lawrence Livermore National Laboratory), Steven H. Langer (Lawrence Livermore National Laboratory), Chris Chambreau (Lawrence Livermore National Laboratory), Louis H. Howell (Lawrence Livermore National Laboratory), Trent D'Hooge (Lawrence Livermore National Laboratory), Matthew L. Leininger (Lawrence Livermore National Laboratory)*

Understanding the characteristics and requirements of applications that run on commodity clusters is key to properly configuring current machines and, more importantly, procuring future systems effectively. There are only a few studies, however, that are current and characterize realistic workloads. For HPC practitioners and researchers, this limits our ability to design solutions that will have an impact on real systems.

We present a systematic study that characterizes applications with an emphasis on communication requirements. It includes cluster utilization data, identifying a representative set of applications from a U.S. Department of Energy laboratory, and characterizing their communication requirements. The driver for this work is understanding application sensitivity to a tapered fat-tree network. These results provided key insights into the procurement of our next generation commodity systems. We believe this investigation can provide valuable input to the HPC community in terms of workload characterization and requirements from a large supercomputing center.

Performance Analysis, Design Considerations, and Applications of Extreme-Scale In Situ Infrastructures

Authors: *Utkarsh Ayachit (Kitware Inc), Andy Bauer (Kitware Inc), Earl P. N. Duque (Intelligent Light), Greg Eisenhauer (Georgia Institute of Technology), Nicola Ferrier (Argonne National Laboratory), Junmin Gu (Lawrence Berkeley National Laboratory), Kenneth Jansen (University of Colorado, Boulder), Burlen Loring (Lawrence Berkeley National Laboratory), Zarija Lukic (Lawrence Berkeley National Laboratory), Suresh Menon (Georgia Institute of Technology), Dmitriy Morozov (Lawrence Berkeley National Laboratory), Patrick O'Leary (Kitware Inc), Reetesh Ranjan (Georgia Institute of Technology), Mirchel Rasquin (Cenaero), Chrisopher P. Stone (Computational Science and Engineering LLC), Venkat Vishwanath (Argonne National Laboratory), Gunther Weber (Lawrence Berkeley National Laboratory), Brad J. Whitlock (Intelligent Light), Matthew Wolf (Georgia Institute of Technology), Kesheng Wu (Lawrence Berkeley National Laboratory), E. Wes Bethel (Lawrence Berkeley National Laboratory)*

A key trend facing extreme-scale computational science is the widening gap between computational and I/O rates, and the challenge that follows is how to best gain insight from simulation data when it is increasingly impractical to save it to persistent storage for subsequent visual exploration and analysis. One approach to this challenge is centered around the idea of in situ processing, where visualization and analysis processing is performed while data is still resident in memory. We examine key design and performance issues of in situ processing at extreme scale on modern platforms: scalability, overhead, performance measurement and analysis, comparison and contrast with a traditional post hoc approach, and interfacing with simulation codes. We illustrate these principles in practice with studies, conducted on large-scale HPC platforms, that include a miniapplication and multiple science application codes, one of which demonstrates in situ methods in use at greater than 1M-way concurrency.

Task-Oriented Runtimes

3:30pm-4:30pm
Room: 355-E

Extended Task Queuing: Active Messages for Heterogeneous Systems

Authors: *Michael LeBeane (University of Texas at Austin), Brandon Potter (Advanced Micro Devices Inc), Abhisek Pan (Advanced Micro Devices Inc), Alexandru Dutu (Advanced Micro Devices Inc), Vinay Agarwala (Advanced Micro Devices Inc), Wonchan Lee (Stanford University), Deepak Majeti (Hewlett Packard Enterprise), Bibek Ghimire (Louisiana State University), Eric Van Tassell (Advanced Micro Devices Inc), Samuel Wasmundt (University of California, San Diego), Brad Benton (Advanced Micro Devices Inc), Mauricio Breternitz (Advanced Micro Devices Inc), Michael L. Chu (Advanced Micro Devices Inc), Mithuna Thottethodi (Purdue University), Lizy K. John (University of Texas at Austin), Steven K. Reinhardt (Advanced Micro Devices Inc)*

Accelerators have emerged as an important component of modern cloud, datacenter, and HPC computing environments. However, launching tasks on remote accelerators across a network remains unwieldy, forcing programmers to send data in large chunks to amortize the transfer and launch overhead. By combining advances in intra-node accelerator unification with one-sided Remote Direct Memory Access (RDMA) communication primitives, it is possible to efficiently implement lightweight tasking across distributed-memory systems.

This paper introduces Extended Task Queuing (XTQ), an RDMA-based active messaging mechanism for accelerators in distributed-memory systems. XTQ's direct NIC-to-accelerator communication decreases inter-node GPU task launch latency by 10-15% for small-to-medium sized messages and ameliorates CPU message servicing overheads. These benefits are shown in the context of MPI accumulate, reduce, and allreduce operations with up to 64 nodes. Finally, we illustrate how XTQ can improve the performance of popular deep learning workloads implemented in the Computational Network Toolkit (CNTK).

Perilla: Metadata-based Optimizations of an Asynchronous Runtime for Adaptive Mesh Refinement

Authors: *Tan Nguyen (Lawrence Berkeley National Laboratory), Didem Unat (Koc University), Weiqun Zhang (Lawrence Berkeley National Laboratory), Ann Almgren (Lawrence Berkeley National Laboratory), NUFAIL FAROOQI (Koc University), John Shalf (Lawrence Berkeley National Laboratory)*

Hardware architecture is increasingly complex, urging the development of asynchronous runtime systems with advance resource and locality management supports. However, these supports may come at the cost of complicating the user interface while programming remains one of the major constraints to wide adoption of asynchronous runtimes in practice. In this paper, we propose a solution that leverages application metadata to enable challenging optimizations as well as to facilitate the task of transforming legacy code to an asynchronous representation. We develop Perilla, a task graph-based runtime system that requires only modest programming effort. Perilla utilizes metadata of an AMR software framework to enable various optimizations at the communication layer without complicating its API. Experimental results with different applications on up to 24K processor cores show that Perilla can realize up to 1.44x speedup over the synchronous code variant. The metadata-enabled optimizations account for 25% to 100% of the performance improvement.

SC16 Salt Lake City, Utah | hpc matters.

Posters

Tuesday, November 15

Research Poster Exhibition
8:30am-5:15pm
Room: Lower Lobby Concourse

Research Poster Reception
5:15pm-7pm
Room: Lower Lobby Concourse
The SC16 Poster program will kick-off with a Poster Reception where attendees will have the opportunity to engage in open discussion with poster authors while enjoying light foods and drinks.

ACM Student Research Competition (SRC) Poster Exhibition
10:00am-5:15pm
Room: Exhibit Hall E, Booth #104

ACM SRC Poster Reception
5:15pm-7pm
Room: Exhibit Hall E, Booth #104
The SC16 Poster program will kick-off with a Poster Reception where attendees will have the opportunity to engage in open discussion with poster authors while enjoying light foods and drinks.

Doctoral Showcase Posters
For details on the Doctoral Showcase poster exhibition and reception, please see the Doctoral Showcase section on page 53.

Wednesday, November 16

Research Poster Exhibition
8:30am-5pm
Room: Lower Lobby Concourse

ACM SRC Poster Exhibition
10:00am-6pm
Room: Exhibit Hall E, Booth #104

Thursday, November 17

Research Poster Exhibition
8:30am-5pm
Room: Lower Lobby Concourse

ACM SRC Poster Exhibition
10:00am-3pm
Room: Exhibit Hall E, Booth #104

ACM Student Research Competition Posters

Poster #SRC1

Scheduling Strategies and Bounds for Cholesky Factorization on Heterogeneous Platforms

Suraj Kumar (Inria)

We provide an analysis and comparison of different dynamic strategies for task graph scheduling on platforms consisting of heterogeneous and unrelated resources, such as GPUs and CPUs. Parallelism inside processing nodes makes static scheduling strategies difficult to precisely predict the performance of both communications and computations, due to shared resources and co-scheduling effects. Recently, to cope with this limitation, many dynamic task-graph based runtime schedulers (StarPU, StarSs, QUARK, PaRSEC) have been proposed. Dynamic schedulers decisions

depend on information such as the set of available tasks, the location of data, the state of the resources and task priorities. Our analysis is deep, but we concentrate on a single kernel, namely Cholesky factorization of dense matrices on heterogeneous platforms. We analyze different dynamic strategies and propose a set of intermediate strategies, by adding more static features into dynamic strategies. We also compute theoretical upper bounds on task graphs performance.

Poster #SRC2

Performance Modeling and Engineering with Kerncraft

Julian Hammer (University of Erlangen-Nuremberg)

Achieving optimal program performance requires deep insight into the interaction of hardware and software. For software developers without an in-depth background in computer architecture, understanding and fully utilizing modern architectures is close to impossible. Therefore, we must support them with easy to understand models and tools.

We present two simple to apply approaches and a tool for regular stencil and streaming codes: Execution-Cache-Memory (ECM) and layer condition (LC) modeling with Kerncraft. The ECM performance model gives precise overall performance predictions of computational kernels, and the LC analysis gives analytically derived hints for well-suited spatial blocking factors. By utilizing Kerncraft, we are able to take the pain out of the process using automation. We show its applicability to stencils and the insights gained through analysis, performance improvements achieved by following Kerncraft's suggestions derived from generalized LCs, and multi-core saturation point predictions gained through ECM modeling.

Poster #SRC3

Discovering Optimal Execution Policies in KRIPKE using RAJA

William Killian (University of Delaware)

Newer architectures create application porting problems for legacy physics applications. With architectures changing frequently (multicore, many-core, GPU), applications need to be adaptable to many different architectures. "KRIPKE is a proxy application whose primary purpose is to research how data layout, programming paradigms, and architectures affect the implementation and performance of Sn transport". We create policy generation tools that emit RAJA policies to generate over 850,000 versions of KRIPKE. RAJA provides a performance portability layer with high-level C++ constructs to allow its users to quickly change execution policies for CPUs and GPUs. We use hill-climbing and subspace search strategies to explore this space without the need for exhaustive search. With hill-climbing, we explore 10% of the total versions while achieving 95.6% of optimal performance. The subspace search strategy yields 98.8% of optimal while only exploring 20% of all versions. Overall, we improved the best existing version of KRIPKE by 19.5%.

Poster #SRC4

Leveraging Neural Networks to Predict Job I/O in HPC Systems

Michael R. Wyatt (University of Delaware)

The Lustre File System is a shared resource for many HPC systems. The metadata server associated with the Lustre File System is liable to become unresponsive and crash if there are too many metadata operations from the connected cluster. In order to avoid this and maintain high throughput on a cluster, jobs must be scheduled with consideration of metadata operations associated with job I/O. Predicting runtime and metadata operations is an essential step towards an I/O aware scheduler. We use a neural network to accurately predict runtime and I/O for clusters located at Lawrence Livermore National Laboratory. Our method is novel in that entire user-submitted job scripts are analyzed by our neural network. In our poster, we present our methods to achieve high predictive accuracy for runtime and I/O.

Poster #SRC5

Transactional Storage Class Memory

Ellis Giles (Rice University)

Emerging Storage Class Memory, or SCM, technologies are promising both byte-addressability and persistence near DRAM speeds operating on the main memory bus. This high-speed, byte-addressable persistence will give rise to new applications, but programmers are faced with a dual edged problem of how to catch spurious cache evictions while atomically grouping stores to manage consistency guarantees in case of failure.

Consistency for byte-addressable persistent data coupled with highly concurrent applications for high performance introduces new challenges. Both Hardware and Software Transactional Memory attempt to solve concurrency challenges but do not address durability. On the other extreme, logging techniques that address durability of data structures in view of failure are not conducive for concurrent high performance applications.

This poster identifies and evaluates areas in both HTM and STM that can be extended to provide durability on Storage Class Memories with little burden on performance. Several approaches are implemented and discussed.

Poster #SRC6

Toward Exascale Computing: An Automatic Hardware-Software Co-Design Framework for Aspen

Mariam Umar (Virginia Polytechnic Institute and State University)

Hardware-software co-design is becoming a constraint as we approach the exascale era, particularly because of increases in complexity and scale. Existing solutions to maximize performance within power constraints, like memory throttling and DVFS, are becoming less efficient with generational increases in type and density of cores and accelerators. This problem is exacerbated when implementing application specific hardware. We, therefore propose automated generation of performance models for hardware and software to simplify the co-design process. Our hardware-software co-design framework is developed for Aspen: a domain specific language designed for testing and analyzing exascale systems. The framework uses application and hardware information, extracts their characteristics, generates intermediate representations, verifies them and produces Aspen models. Our approach has numerous applications, e.g., sensitivity analysis, estimating and comparing performance models on current and future architectures. Here we present sensitivity analysis of six applications, and scaling and energy profiling for Titan and future exascale system.

Poster #SRC7

Job Startup at Exascale: Challenges and Solutions

Sourav Chakraborty (Ohio State University)

We identify the major performance bottlenecks and memory constraints in bootstrapping MPI and PGAS applications and propose scalable solutions for them. We introduce on-demand connection management to OpenSHMEM to reduce startup time. We also propose extensions to the PMI (Process Management Interface) standard to reduce the data transferred over the network as well as overlap the communication with other initialization tasks. Finally, we introduce a shared memory based design for PMI to reduce its memory footprint by a factor of processes per node (PPN). Our evaluation shows that with sufficient overlap, near-constant initialization time can be achieved at any process count for MPI and hybrid MPI+PGAS applications. Time taken for MPI_Init is reduced by 2.88 times at 16,384 processes. Initialization time of OpenSHMEM is improved by 29.6 times at 8,192 processes. Estimated memory footprint for PMI is reduced by nearly 1GB with 1 million processes and 16 PPN.

Poster #SRC8

Application of Numerical Accuracy to the Selection of Lossy Compression Error Tolerances

Jon Calhoun (University of Illinois)

Long running HPC applications depend on checkpoint restart to recover from failures and utilize multiple time allocations. Memory bandwidth and in particular file system bandwidth continues to be limiters to application performance and scalability. Compression techniques can be used to reduce data size limiting its impact. Lossless compression fails to generate high compression factors for floating-point data. Lossy compression generates noticeably higher compression factors at the expense of adding a small but controllable amount of error into the simulation. In this poster, a methodology for selecting lossy compression error tolerances is presented and evaluated that interprets compression error as numerical error. This interpretation allows defining error tolerances that hide compression error in the high-order truncation error already present in the simulation. This error tolerance selection methodology is shown to add error that does not effect the result of the simulation.

Poster #SRC9

Enabling a Data-Centric Model on the Open Community Runtime

Sri Raj Paul (Rice University)

Achieving exascale performance requires addressing challenges arising from the complexity of exascale architecture, its power constraints, as well as application complexity. Exascale architectures, expected to consist of millions of heterogeneous cores and extremely non-uniform hierarchical memory, have to optimally distribute applications, data and parallelize computation. Dynamic task-based execution models hold promise in achieving this, as they help express fine-grained parallelism, along with decoupling computation and data from underlying resources. Open Community Runtime(OCR) is a community-led effort to explore various asynchronous task-parallel runtime principles that can support a broad range of higher-level programming constructs. Legion is a data-centric programming model in which the runtime extracts task-based parallelism from programs, freeing the developer from having to explicitly express them. In this poster we describe our efforts to run Legion programs on top of OCR as their underlying runtime, to combine the parallelism extracted by Legion with the flexibility provided by OCR.

Poster #SRC10

Design and Evaluation of Topology-Aware Scatter and AllGather Algorithms for Dragonfly Networks

Nathanael Cheriere (ENS Rennes)

High-radix direct network topologies such as Dragonfly have been proposed for petascale and exascale supercomputers because they ensure fast interconnections and reduce the cost of the network compared with traditional network topologies. The design of new machines such as Theta with a Dragonfly network present an opportunity to further improve the performance of distributed applications by creating communication algorithms that are aware of this particular topology. In this way, one can exploit the strength of the Dragonfly network while mitigating its bottlenecks.

In this poster, we study existing and new algorithms for the AllGather and Scatter operations. We evaluate them using an event-driven simulator. The simulations show expected as well as unintuitive results and demonstrate that the network's hardware must also be taken into consideration when creating topology-aware algorithms to reach the best performance.

Poster #SRC11

Electron Dynamics Simulation with Time-Dependent Density Functional Theory on Large Scale Many-Core Systems

Yuta Hirokawa (University of Tsukuba)

Many-core processors such as the Intel Xeon Phi (Knights Corner, KNC) and GPUs provide new solutions for HPC systems. However, it is difficult to achieve high computation efficiency on them due to the different characteristics from traditional processors such as Intel Xeon CPU. In this study, I implement an electron dynamics simulation as real scientific code to the three types of many-core processors. A kernel of stencil computation that dominates the total computation time is optimized in a single-thread level with explicit vectorization with 256/512-bit SIMD instructions. As a result, the stencil computation performance achieves maximum 591.2 GFLOPS with two KNC and 215.1 GFLOPS with a SPARC64 XIfx while single KNL (Knights Landing) chip achieves 707.2 GFLOPS. For the entire code, the cooperative computation with KNC and Xeon achieves 1.97 times better performance compared with the CPU-only utilization at strong scaling case.

Poster #SRC12

Parallel Algorithms for Updating Large Dynamic Networks Using Graph Sparsification

Sriram Srinivasan (University of Nebraska Omaha)

Algorithms for analyzing dynamic networks help study how properties of complex systems, such as those in bioinformatics and social sciences, evolve with time. Since the networks are very large, it is essential to design parallel algorithms for updating their properties. However, because network data is highly unstructured and exhibit poor locality of access, designing fast and scalable dynamic algorithms is challenging.

We present a framework for converting sequential algorithms on static networks to efficient parallel algorithms on dynamic networks using a technique called graph sparsification that expresses network algorithms in a reduction-like format. We describe how our framework can be used to update connected components and the minimum weighted spanning tree (MST), present a new graph ordering method for reducing computation time and the scalability results for these methods on shared memory systems. To the best of our knowledge, this is the first parallel implementation for updating MST.

Poster #SRC13

Mapping Applications on Irregular Allocations

Seonmyeong Bak (University of Illinois)

Mapping applications on clusters becomes more difficult as the number of nodes become larger. Supercomputers assign allocations with irregular shapes to users to maximize the utilization of resources, and it is much more difficult to map applications on these irregular allocations.

We extended Rubik, a python based framework to map applications on irregular allocations with a few lines of python code. Rubik was originally designed for regular allocations, so we added features to handle allocations with irregular structure and unavailable nodes and two mapping algorithms such as row-ordering and recursive splitting.

We evaluate our work with two widely used HPC applications on Blue Waters: MILC and Qbox. We reduced execution time by 32.5% in MILC, and by 36.3% in Qbox, and communication time by 60% in MILC and 56% in Qbox.

Poster #SRC14

The Search for Missing Parallel I/O Performance on the Cori Supercomputer

Matt Bryson (Lawrence Berkeley National Laboratory)

With scientific computing approaching exascale rapidly, I/O is becoming the main cause of bottlenecks in High Performance Computing (HPC). To solve this, next generation systems such as the NERSC Cori Supercomputer are equipped with I/O nodes equipped with NVRAM, otherwise known as burst buffers (BB). BB systems are designed to provide more I/O throughput than traditional Parallel File Systems (PFS) that rely on magnetic storage. This has not been the case with the Cray Burst Buffer (CBB) which is performing at 11.07% of peak performance when confronted with the popular HDF5 file library. We use the Vector Particle in Cell (VPIC) I/O kernel to benchmark the system and find potential optimization strategies. By changing the I/O access pattern of VPIC I/O kernel we are able to improve performance up to 4.6 times in some configurations.

Poster #SRC15

Touring Dataland? Automated Recommendations for the Big Data Traveler

William C. Agnew (University of Chicago)

We explore how recommendation techniques can be adapted and applied to big data science. Using features specific to big data science, we develop a set of data location prediction heuristics. We combine these heuristics into a single recommendation engine using a deep recurrent neural network. We show, via analysis of historical Globus operations, that our approaches can predict the storage locations accessed by users with 78.2% and 95.5% accuracy for top-1 and top-3 recommendations, respectively.

Poster #SRC16

Revisiting Popcount Operations in CPUs/GPUs

Chenfan Sun (University of Washington)

Popcount is a binary operation where the input is a binary word and the output is the number of set bits. Popcount is a common building block for many applications such as the Hamming distance calculation. Recently, popcount is used to approximate multiplications in convolutional neural networks. Considering renewed interest in popcount, this work asks the question: can programmers lazily use the builtin popcount intrinsic or is further tuning necessary to achieve peak performance?

In this work, we benchmark the efficacy of several popcount implementations on both the CPU and GPU to analyze their behaviors under different working set sizes. On the CPU, results show that for memory bound workloads, the builtin popcount compiler intrinsic is within 0.01% of the fastest hand-tuned implementations suggesting that no hand-tuning is required, while this gap is up to 60% in the compute bound scenario where hand-tuned implementations of popcount matter.

Poster #SRC17

FemtoGraph: A Pregel Based Shared-Memory Graph Processing Library

Alexander Ballmer (Illinois Institute of Technology)

The emerging applications for large graphs in big data science and social networks has led to the development of numerous parallel or distributed graph processing applications. The need for faster manipulation of graphs has driven the need to scale across large core counts and many parallel machines. While distributed memory parallel systems continue to be used for high performance computing, some smaller systems make use of shared memory (SMP) and larger core counts. We have implemented a graph processing framework for shared memory systems capable of scaling past 48 parallel cores. This system leverages and scale to large core counts and provide a framework for later incorporating distributed processing across multiple nodes.

Poster #SRC18

Analysis of Variable Selection Methods on Scientific Cluster Measurement Data

Jonathan Wang (University of California, Berkeley)

The goal of the project was to use parallelized variable selection methods to improve the performance of machine learning models on the PTF astrophysics dataset by reducing model training time and removing disruptive variables. Several methods were implemented in Spark to utilize high performance computing and tested on the PTF data. The results from the PTF data tests showed that Sequential Backward Selection was able to approximate the optimal subset relatively quickly. This subset took significantly less time to train on and had higher accuracy than the full feature set. We also experimented with correlation-based grouping to take advantage of feature correlations in the PTF data. This method allows large correlated datasets to be handled more efficiently. We were able to further improve the performance of Sequential Backward Selection on this dataset without significant loss in accuracy.

Poster #SRC19

GNI Provider Lovector Support for Libfabric

Evan Harvey (Cray Inc.)

Portable HPC middleware such as MPICH has typically targeted low-level network APIs that are vendor specific. Libfabric is a new portable, low-level API that aims to alleviate the burden of porting to new network APIs without sacrificing performance. The Libfabric GNI "provider" was developed to enable testing Libfabric-based middle-ware at scale using Cray XC(TM) systems with the Aries interconnect. This poster describes the implementation of scatter-gather lists for point-to-point data transfers using the Aries Chip.

Poster #SRC20

Narrowing the Gap: Effects of Latency with Docker in IP Networks

Corbin L. Higgs (Pendleton High School)

Recent work has shown that Docker containers can useful for packaging HPC applications with their runtime environments. Applications in containers perform similarly to native; however, some work has shown that containers can adversely affect the latency variation of network traffic. This variation may have an impact on the performance of some HPC workloads, especially those that are dependent on synchronization.

In this work, we report on the latency characteristics of MPI communications using Docker containers and the Linux bridge. Although we observe lower mean and variation in latency in simple microbenchmarks, synchronization time of applications using the bridge is higher and more variable. We then report on application benchmarks that suffer from low performance with the Linux bridge, and correlate these results to high communications frequency. In conclusion, we find that Docker containers add little to no performance cost to HPC applications, but the Linux bridge can cause degradation.

Poster #SRC21

Quantization for Energy Efficient Convolutional Neural Networks

Joao Vitor Mascarenhas (Federal University of Ouro Preto)

A traditional Convolutional Neural Network (CNN) is parameterized by floating point weights and biases and takes floating point data as input. In many cases, the floating point representation of these parameters and input is more than necessary. The use of a more compact representation of the parameters and input allows CNNs to be deployed on energy efficient architectures that operate with a few bits and much lower memory footprint. This work focuses on data reduction and quantization schemes that can be applied to a trained CNN for classifying scientific simulation data. We show that each neuron and synapse can be encoded with only one byte to maintain accuracy above 98%.

Poster #SRC22

Accelerated Signed Distance Queries For Performance Portable Physics Codes

Evan, Jordan DeSantola, Backes (Lawrence Livermore National Laboratory)

Signed distance is commonly employed to numerically represent material interfaces with complex boundaries in multi-material numerical simulations. However, the performance of computing the signed distance field is hindered by the complexity and size of the input. Recent trends in HPC architecture consist of multi-core CPUs and accelerators that collectively expose tens to thousands of cores to the application. Harnessing this massive parallelism for computing the signed distance field presents significant challenges. Chief among them, the design and implementation of a performance portable solution that can work across architectures. Addressing these challenges to accelerate signed distance queries is the primary merit of this work. Herein, we employ the RAJA programming model, which provides a loop-level abstraction that decouples the loop-body from the parallel execution and insulates application developers from non-portable compiler and platform-specific directives. Implementation and performance results are discussed in more detail.

Poster #SRC23

Energetically Efficient Acceleration EEA-Aware For Scientific Applications of Large-Scale On Heterogeneous Architectures

John Anderson Garcia Henao (Industrial University of Santander)

Heterogeneous parallel programming has two main problems on large computation systems: the first is the increase of power consumption on supercomputers in proportion to the amount of computational resources used to obtain high performance, the second one is the underuse of these resources by scientific applications with improper distribution of tasks. Select the optimal computational resources and make a good mapping of task granularity is the fundamental challenge for the next generation of Exascale Systems. This research proposes an integrated energy-aware scheme called efficiently energetic acceleration (EEA) for large-scale scientific applications running on heterogeneous architectures. The EEA scheme uses statistical techniques to get GPU power levels to create a GPU power cost function and obtains the computational resource set that maximizes energy efficiency

for a provided workload. The programmer or load balancing framework can use the computational resources obtained to schedule the map parallel task granularity in static time.

Poster #SRC24

Discovering Energy Resource Usage Patterns on Scientific Clusters

Matthew Bae (Harvey Mudd College)

With the growth of scientific clusters, there has been an increase in volumes of data, number of machines, and exploited parallelism. We are now seeing increasing interactions of hardware components within clusters. As a result, system resource usage patterns are becoming increasingly harder to detect. The National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL) uses Simple Linux Utility for Resource Management (SLURM) in its clusters, which outputs logs about the characteristics of jobs. On Cori, LBNL's Cray XC40 supercomputer, we are able to read energy counters for jobs. This allows us to analyze patterns related to energy consumption. We show that energy consumption patterns arise based on different variables such as CPU load and CPU utilization.

Poster #SRC25

Parallel Provenance Databases for High Performance Workflows

Jennifer A. Steffens (Drake University)

In scientific computing, understanding the origins and derivation of data is crucial. Provenance models aim to provide a means of capturing this in an efficient and effective manner. For the Swift/T language, the current provenance handling system requires improvement. In this poster, we discuss the development of a new Swift/T provenance model, the Multiple Parallel Databases Model (MPDM), which parallelizes the real-time storage of provenance data in a user-accessible database system. Utilizing multiple databases in high performance, parallel workflows can increase the practicality of lightweight, relational databases engines such as SQLite, as we show MPDM to be more efficient and have better scalability than the previous, single database model.

Poster #SRC26

GPU-Accelerated Jacobi-Like Algorithm for Eigendecomposition of General Complex Matrices

Basileal Y. Imana (Trinity College)

Jacobi-like methods, though introduced in 1846, were not popular as they were computationally expensive for sequential computing. With the advent of parallel computing, however, it has become feasible to efficiently implement such algorithms in parallel. In addition, the Jacobi method has been shown to compute very small eigenvalues with high accuracy compared to the conventional methods. In this research, we present a novel parallel implementation of Jacobi method for eigendecomposition of general complex matrices on the GPU. Our preliminary results show a significant improvement over those on the CPU, running up to 94 times faster for general dense complex matrices of moderate size.

Research Posters

Algorithms

Poster #28

A Scalable Evolutionary Algorithm with Intensification and Diversification Protocols Designed for Statistical Models

Wendy K. Cho (University of Illinois), Yan Y. Liu (University of Illinois)

Important insights into many problems that are traditionally analyzed via statistical models can be obtained by re-formulating within a large-scale optimization framework. The theoretical underpinnings of statistical models may shift the goal of the solution space traversal from the traditional search for an optimal solution to a traversal with the purpose of yielding a set of high quality, independent solutions. We examine statistical frameworks with astronomical solution spaces where the independence requirement constitutes a significant additional challenge for standard optimization methodologies. We design a hybrid metaheuristic with intensification and diversification protocols in the base search algorithm. Via our grant on the Blue Waters supercomputer, we extend our algorithm to the high-performance-computing realm. We experimentally demonstrate the effectiveness of our algorithm to utilize multiple processors to collaboratively hill climb, broadcast messages to one another about the landscape characteristics, diversify across the landscape, and request aid in climbing particularly difficult peaks.

Poster #29

GPU Accelerated Graph Analytics Using Abstract Sparse Linear Algebra

Stephen T. Kozacik (EM Photonics Inc), Aaron L. Paolini (EM Photonics Inc), Paul Fox (EM Photonics Inc), James L. Bonnett (EM Photonics Inc), Evenie M. Chao (EM Photonics Inc), Eric Kelmelis (EM Photonics Inc), Dennis W. Prather (University of Delaware)

Large-scale graph analytics using conventional processing systems often involves prohibitively excessive runtimes. Phrasing graph operations as abstract linear algebra operations offers potential for leveraging backend processing engines that scale better on modern heterogeneous computing platforms. To achieve this goal, we present a language for graph analytics that allows users to write in a familiar, vertex-centric API while leveraging the computational power of many-core accelerators. Our prototype toolchain automatically employs abstract sparse linear algebra (ASLA) operations using custom semi-rings in order to maximize performance. We use a C++ EDSL to map the user's algorithm to efficient, fused ASLA operations at compile time with no runtime overhead. Using this technique, we have implemented several algorithms, including single-source shortest path, PageRank and single-source widest path. With a GPU-accelerated linear algebra backend, we can achieve better than 500% speedup over a multi-core ASLA library, using real and synthetic datasets.

Poster #30

Accelerating DMFT-MatDeLab GW Calculation with GPU

Zhihua Dong (Brookhaven National Laboratory), Kwangmin Yu (Brookhaven National Laboratory), Sangkook Choi (Brookhaven National Laboratory), Nicholas D'Imperio (Brookhaven National Laboratory)

GW approximation and DMFT are theories to calculate electron structure in material science. The electron self-energy Σ and polarizability χ matrix elements construction in GW calculation iterations are the most time consuming parts, consuming about 40% of the total running time when calculating the Cerium α band on Titan at the ORNL Leadership Computing Facility. GPU acceleration was implemented on those two parts. To optimize the code, we overcome the difficulty of heavy non-sequential memory access caused by updating random elements of arrays at the inner most loop by sorting the index pattern and rearranging the calculation so that each GPU thread calculates one element, avoiding memory writing during updates. Test runs were conducted on Titan with 768 MPI processes. Results show a 26X speed up of the χ matrix calculation and 31.6% performance improvement of a complete α -Ce run.

Poster #31

Cache-Oblivious Wavefront Algorithms for Dynamic Programming Problems: Efficient Scheduling with Optimal Cache Performance and High Parallelism

Jesmin Jahan Tithi (Stony Brook University), Pramod Ganapathi (Stony Brook University), Rezaul Chowdhury (Stony Brook University), Yuan Tang (Fudan University)

Wavefront algorithms are algorithms on grids where execution proceeds in a wavefront manner from the start-to-the-end of the execution. Iterative-wavefront algorithms for evaluating dynamic programming (DP) recurrences exploit optimal-parallelism but show poor cache performance. Tiled-iterative-wavefront algorithms achieve optimal cache-complexity and high-parallelism but are cache-aware and not portable, neither cache-adaptive. In contrast, standard cache-oblivious recursive divide-and-conquer (CORDAC) algorithms have optimal serial-cache-complexity but often have low-parallelism due to artificial-dependencies among subtasks.

The cache-oblivious wavefront algorithms for DP problems are variants of CORDAC algorithms with reduced artificial-dependencies and, hence, have better parallelism.

We show how to transform a CORDAC algorithm to a recursive-wavefront algorithm to achieve optimal parallel-cache-complexity and high-parallelism under state-of-the-art schedulers for fork-join programs. These cache-oblivious wavefront algorithms use closed-form formulas to compute at what time each divide-and-conquer function must be launched in-order-to achieve high-parallelism without losing cache-performance. We present experimental performance and scalability results showing superiority of these new algorithms.

Poster #32

Domain Decomposition Techniques for Contour Integration Eigenvalue Solvers

Vassilis Kalantzis (University of Minnesota), Yousef Saad (University of Minnesota), James Kestyn (University of Massachusetts), Eric Polizzi (University of Massachusetts)

This poster discusses techniques for computing a few selected eigenvalue-eigenvector pairs of large and sparse symmetric matrices. A recently developed powerful class of techniques to solve this type of problems is based on integrating the matrix resolvent operator along a complex contour that encloses the interval containing the eigenvalues of interest. This poster considers such contour integration techniques from a domain decomposition viewpoint, and extends the concept of domain decomposition linear system solvers in the framework of contour integration methods for eigenvalue problems, such as FEAST. We describe a multi-parallel implementation of FEAST using domain decomposition, discuss how the different levels of parallelism can be exploited, and report results on distributed computing environments reported. These results show that domain decomposition approaches can lead to reduced runtimes and improved scalability.

Poster #33

Massively Parallel Simulation of Plasma Turbulence with the Sparse Grid Combination Technique

Mario Heene (University of Stuttgart), Alfredo Parra Hinjosa (Technical University Munich), Dirk Pflüger (University of Stuttgart)

The problem sizes for the solution of higher-dimensional PDEs, such as for the simulation of plasma turbulence in a fusion device, often are very limited due to the exponential increase in the degrees of freedom with the dimensionality. With the Sparse Grid Combination Technique (SGCT) we can mitigate this so-called curse of dimensionality and push the computational limits of high-dimensional simulations significantly. Offering a second level of parallelism, the SGCT minimizes the need for global communication and ensures scalability on future (exascale) HPC systems. First experimental results demonstrate the scalability of our distributed algorithm up to 180,225 cores on the supercomputer Hazel Hen. Furthermore, the inherent data redundancy of this technique enables fault-tolerance on the algorithmic level at low cost without the need of checkpointing or process replication. These are the key issues we investigate within the project EXAHD of Germany's priority program "Software for Exascale Computing" (SPPEXA).

Poster #34

Sparse Grid Algorithms to Recover from Hard and Soft Faults

Alfredo Parra Hinojosa (Technical University Munich), Hans-Joachim Bungartz (Technical University Munich), Mario Heene (University of Stuttgart), Dirk Pflüger (University of Stuttgart)

High dimensional PDEs present a challenge in computing due to the exponential growth of discretization points with increasing dimension. An algorithm to tackle such problems is the Sparse Grid Combination Technique (SGCT), which is an extrapolation scheme. The SGCT has some inherent data redundancy that can be exploited to make it tolerant to both hard and soft faults: it can recover from process failures as well as from data corruption without the need for checkpointing, process replication or any of the typical system-level approaches. We describe two main results: first, that our parallel implementation of the SGCT scales well with simulated hard faults on a large parallel system (Hazel Hen). And second, that the SGCT can be extended to deal with Silent Data Corruption, a type of soft fault that is becoming more common as supercomputers grow in size. This makes the SGCT a promising algorithm for future exascale systems.

Poster #35

High-Performance Tensor Contraction without BLAS

Devin A. Matthews (University of Texas at Austin)

Tensor computations – in particular tensor contraction (TC) – are important kernels in many scientific computing applications (SCAs). Due to the fundamental similarity of TC to matrix multiplication (MM) and to the availability of optimized implementations such as the BLAS, tensor operations have traditionally been implemented in terms of BLAS operations, incurring both a performance and a storage overhead. Instead, we implement TC using the flexible BLIS framework, which allows for reshaping of the tensor to be fused with partitioning and packing operations, requiring no reshaping operations or additional workspace. This implementation, TBLIS, achieves performance approaching that of MM, and in some cases considerably higher than that of traditional TC.

Our implementation also supports multithreading using an approach identical to that used for MM in BLIS, with similar performance characteristics. The complexity of managing tensor-to-matrix transformations is also handled automatically in our approach, greatly simplifying use in SCAs.

Poster #36

Minimizing Communication for Tensor Decompositions

Travis Bartley (University of California, Irvine), Mona Elkoussy (University of California, Irvine), Anima Anandkumar (University of California, Irvine), Aparna Chandramowlishwaran (University of California, Irvine)

Previous research in numerical linear algebra has proven that the same communication lower bound is attained by variety of algorithms, including matrix multiplication, LU factorization, Cholesky factorization, and algorithms for eigenvalues and singular values. This article extends this result to three tensor decomposition methods. For computing the CP decomposition, the alternating least squares method (CP-ALS) is studied. For the Tucker decomposition, the higher-order singular value decomposition (HOSVD) and higher-order orthogonal iteration (HOOI) methods are studied. For each method, this article contributes an algorithm that attains the previously mentioned communication lower bound.

Poster #37

Enabling K-Nearest Neighbor Algorithm Using a Heterogeneous Streaming Library: hStreams

Jesmin Jahan Tithi (Intel Corporation)

hStreams is a recently proposed (IPDPSW 2016) task-based target-agnostic heterogeneous streaming library that supports task concurrency on heterogeneous platforms. In this poster, we share our experience of enabling a non-trivial machine learning (ML) algorithm: K-nearest neighbor using hStreams. The K-nearest neighbor (KNN) is a popular algorithm with numerous applications in machine learning, data-mining, computer vision, text processing, scientific computing such as in computational biology, astronomy, physics, and in other areas. This is the first example of showcasing hStreams' ability to enable an ML algorithm. hStreams enabled KNN achieves the best performance achievable by either Xeon R or Xeon PhiTM 1 by utilizing both platforms simultaneously and selectively.

Poster #38

Parallelized Dimensional Decomposition for Dynamic Stochastic Economic Models

Aryan Eftekhari (University of Lugano), Olaf Schenk (University of Lugano), Simon Scheidegger (University of Zurich)

This project explores a technique called Dimensional Decomposition, which allows for the separation of a function into a finite number of lower-dimensional component functions. The method leverages the lack of input-output correlation to effectively reduce the dimensionality of the problem. This program has been integrated with a sparse grid approximation to form an efficient approximation method referred to DDSG (dimensional decomposition with sparse grid). Due to the intrinsic separability and hierarchical construction, in both dimensional decomposition and sparse grid, a highly parallelizable framework has been developed. This framework has been applied in the context of computational economics, in which we provide an efficient solution method for high-dimensional dynamic stochastic models. Our findings show that DDSG can effectively capture model dynamics with relatively low-dimensional component functions, thus mitigating the so-called "curse of dimensionality".

Poster #75

Toward Portable Machine Learning Kernels for Deep Neural Networks with Autotuning on Top of OpenCL and High Bandwidth Memory GPUs

Yaohung Tsai (University of Tennessee), Piotr Luszczek (University of Tennessee), Jakub Kurzak (University of Tennessee), Jack Dongarra (University of Tennessee)

We present a portable, highly-optimized Deep Neural Network (DNN) algorithm and its implementation techniques. Our approach combines in novel ways existing HPC techniques such as autotuning, data layout, and low-level optimizations that, when applied simultaneously, achieve performance that matches and exceeds what is possible with either reverse engineering and manual assembly coding or proprietary vendor libraries. The former was done inside the maxDNN implementation, and the latter is represented by cuDNN. Our work may be directly applied to the most time consuming part DNN workflow, namely the training process which often needs restart when it stagnates due to, among other reasons, diminishing gradients and getting stuck in local minima. We used for our performance tests a consumer-grade GPU with the latest High Bandwidth Memory (HBM) stack which can match a number of server grade hardware at fraction of the price which attests to the portability of our approach and implementation.

Poster #76

Acceleration of the Boundary Element Library BEM4I on the Knights Corner and Knights Landing Architectures

Michal Merta (Technical University of Ostrava), Jan Zapletal (Technical University of Ostrava)

The aim of the poster is to present acceleration of the boundary element method (BEM) by the Intel Xeon Phi technology. Since the classical BEM produces dense matrices, it is particularly suitable for acceleration on manycore architectures due to data locality and the possibility to employ vectorization and optimized dense BLAS routines. Efficient implementation of BEM is necessary to enable its application to large engineering problems. The poster provides brief overview of BEM followed by the discretization approach and efficient numerical assembly of the BEM matrices (full and sparsified). We discuss the parallelization by OpenMP in shared memory and the SIMD vectorization necessary to exploit the full potential of the Xeon and Xeon Phi architectures. An interface to the domain decomposition library Espreso enables solution of large scale problems in distributed memory. We present numerical experiments performed both on the Knights Corner coprocessor and the Knights Landing standalone processor.

Poster #77

Fast Sparse General Matrix-Matrix Multiplication on GPU with Low Memory Usage

Yusuke Nagasaka (Tokyo Institute of Technology), Akira Nukada (Tokyo Institute of Technology), Satoshi Matsuoka (Tokyo Institute of Technology)

Sparse general matrix-matrix multiplication (SpGEMM) is one of the key kernel of preconditioner such as algebraic multigrid method or graph algorithms. The performance of SpGEMM is quite low because of its random memory access to both input and output matrices. Moreover, the pattern of non-zero elements of resulting matrix is not known beforehand, which makes it hard to manage the memory usage. There are several GPU implementations of fast SpGEMM computation while consuming large temporal memory. We devise new SpGEMM algorithm requiring small amount of memory so that we can compute larger matrices using limited device memory of GPU. Accesses to input matrices are optimized for coalesced memory access. We devise efficient hash table on shared memory to calculate output matrix with appropriate case analysis for better load-balancing. Our algorithm achieves speedups of up to x4.0 in single precision and x3.3 in double precision compared to existing fast SpGEMM libraries.

Poster #78

A Task-Based Directed Acyclic Graph Implementation of Additive AMG

Amani Alonazi (King Abdullah University of Science and Technology), George Markomanolis (King Abdullah University of Science and Technology), David Keyes (King Abdullah University of Science and Technology)

Algebraic multigrid is the solver of choice in many and growing applications in today's petascale environments and scales well in a weak, distributed-memory sense. Exascale architectures exacerbate the challenges of increased strong SIMT concurrency, reduced memory capacity and bandwidth per core, and vulnerability to global synchronization. The recent Vassilevski-Yang (2014) "mult-additive" AMG nearly preserves the convergence rate of the multiplicative form of AMG, while reducing communication and synchronization frequency and controlling memory growth. However, it remains bulk-synchronous. We extend the algorithm further toward exascale environments with a task-based directed acyclic graph implementation. The algorithm can then be represented as a task-based DAG where vertices represent tasks, and edges are dependencies among them. We implement a tiling approach for decomposing the grid hierarchy into parallel units within task containers. The distribution of tiles is left to the compiler and OmpSs runtime system to expose task-level parallelism.

Poster #79

Parallel FETI Solver for Modern Architectures

Lubomir Riha (IT4Innovations National Supercomputing Center, Ostrava, Czech Republic), Tomas Brzobohaty (IT4Innovations National Supercomputing Center, Ostrava, Czech Republic), Michal Merta (IT4Innovations National Supercomputing Center, Ostrava, Czech Republic), Alexandros Markopoulos (IT4Innovations National Supercomputing Center, Ostrava, Czech Republic), Ondrej Meca (IT4Innovations National Supercomputing Center, Ostrava, Czech Republic), Tomas Kozubek (IT4Innovations National Supercomputing Center, Ostrava, Czech Republic)

This poster presents a parallel implementation of the Hybrid Total FETI (HTFETI) method. HTFETI is a multilevel domain decomposition method in which a small number of neighboring subdomains is aggregated into clusters. The presented HTFETI solver is able to solve 3D heat transfer problems of size up to 124 billion unknowns on 17,576 nodes of the Titan supercomputer. A superlinear strong scalability of the iterative solver is presented for 20 billion unknown heat transfer problems and 11 billion DOF structural mechanics problem executed on 2,744 to 17,576 nodes. In addition, a super-linear strong scaling is presented for 300 million DOF real world problem.

The acceleration of the HTFETI method using Local Schur Complement method delivers 7.8 speedup (heat transfer) and 2.7 speedup (structural mechanics) of the iterative solver when executed on the Intel Xeon Phi and 3.4 speedup for structural mechanics delivered by the GPUs on the Titan machine.

Poster #80

symPACK: a Solver for Sparse Symmetric Matrices

Mathias Jacquelin (Lawrence Berkeley National Laboratory), Yili Zheng (Google), Esmond Ng (Lawrence Berkeley National Laboratory), Katherine Yelick (Lawrence Berkeley National Laboratory)

Systems of linear equations arise at the heart of many scientific applications. Most of these systems are sparse. Direct methods are sometimes needed to ensure accurate solutions (e.g. in shift-invert Lanczos). Performance and resource usage of sparse matrix factorizations are critical to time-to-solution and maximum solvable problem size. In many applications, matrices are symmetric. Exploiting the symmetry reduces both the amount of work and storage cost. On large-scale distributed memory platforms, communication cost can become critical. In addition, network topologies have become more complex. Modern platforms thus exhibit a higher level of performance variability, which makes scheduling of computations an intricate task. We investigate the novel use of an asynchronous task paradigm coupled with dynamic scheduling, in implementing sparse Cholesky factorization. Our solver symPACK, shows good strong scaling. It relies on efficient communication primitives provided by the UPC++ library. Performance evaluation shows that symPACK outperforms state-of-the-art packages, validating our approach.

Poster #81

Exploring Performance of Domain Decomposition Strategies for Monte Carlo Radiation Transport

Platon Karpov (University of California, Santa Cruz), David Huff (New Mexico Institute of Mining and Technology), Xinyu Chen (University of New Mexico), Ryan Wollaeger (Los Alamos National Laboratory), Gabriel Rockefeller (Los Alamos National Laboratory), Brendan Krueger (Los Alamos National Laboratory)

SuperNu is a Monte Carlo (MC) radiation transport code to simulate light curves of explosive outflows from supernovae. The MC transport step is domain replicated. To enable scaling on next-generation HPC systems, we implemented the recursive coordinate bisection approach of domain decomposition for the opacity calculation. Then, we propagated the decomposition to other steps in the simulation and constructed a communication infrastructure to support the decomposition. In this poster, we demonstrate the results of two communication schemes: the Improved KULL and Improved Milagro algorithms. We present results from a range of processors: Ivybridge, Haswell on local clusters at LANL, and the AMD Bulldozer Opteron 6200 series processors at Blue Waters' Cray XE6 nodes. Finally, we tested the scalability of SuperNu on the latest Intel Xeon Phi architecture, Knights Landing.

Poster #82

Distributed Graph-Based Clustering for Network Intrusion Detection

Corbin A. McNeill (Wheaton College), Enyue Lu (Salisbury University), Matthias Gobbert (University of Maryland, Baltimore County)

In order to process large volumes of network traffic data and quickly detect intrusions, we use parallel computation frameworks for Network Intrusion Detection Systems (NIDS). Additionally, we model network data as graphs to highlight the interconnected nature of network traffic, and apply unsupervised graph-based clustering to flag anomalies as network intrusions. We particularly examine the effectiveness of barycentric clustering on graph network traffic models for intrusion detection. The clustering algorithms have been implemented in Hadoop MapReduce for the purpose of rapid intrusion detection. Furthermore, k-nearest neighbor graphs (kNNGs) are used to optimize the clustering process. We find that across various data sets, unsupervised graph based clustering is able to exceed a 92% intrusion detection accuracy and that kNNGs can effectively optimize the network traffic graphs for larger data sets.

Poster #83

A Novel Variable-Blocking Representation for Efficient Sparse Matrix-Vector Multiply on GPUs

Tuowen Zhao (University of Utah), Tharindu Rusira (University of Utah), Khalid Ahmad (University of Utah), Mary Hall (University of Utah)

Fillrate-guided block compressed sparse row (FBCSR) is a novel approach to improve the performance of sparse matrix-vector multiply (SpMV) on GPUs. Motivated by the observation that in the finite element method, many of the matrices consist of dense block of different sizes and unaligned starting positions, FBCSR can identify and extract those local nonzero patterns that could improve the memory access and reduce intra-warp divergence of the corresponding SpMV kernels. As compared to other variable blocking methods such as unaligned block compressed sparse row, it relies on local patterns and can generate larger blocks well-suited for single instruction, multiple threads processing, while also tolerating a bounded number of generated zero-fillins. We present the

SpMV performance results for FBCSR on two generations of Nvidia GPUs, which shows that FBCSR outperforms available alternatives for the matrices to which it is applicable.

Poster #84

Energy and Communication Efficient Partitioning for Large-Scale Finite Element Computations

Milinda Fernando (University of Utah), Dmitry Duplyakin (University of Colorado, Boulder), Hari Sundar (University of Utah)

Load balancing and partitioning are critical when it comes to parallel computations. Generally partitioning involves equally dividing the work and data among the processors, reducing processor idle time and communication costs. As we march towards exascale machines, the cost of data movement and load-imbalances therein are a major bottleneck for achieving scalability. We propose an alternative Space Filling Curve (SFC)-based partitioning scheme where we allow some (user-specified) flexibility in the work assignment, so as to minimize the data-dependencies across partitions. Effectively, we show that the flexibility in SFC based partitioning schemes leads towards minimizing the communication load-imbalance (Hilbert- 4.9% and Morton- 12.18%) and overall energy consumption (Hilbert- 22.0% and Morton-5.0%) at the cost of a marginal increase in workload-imbalance. The traditional SFC-based partitioning can be recovered by setting the flexibility to zero.

Poster #85

Identifying Malicious Entities through Massive Graph Analysis of DNS Resolution Patterns

Trevor J. Goodyear (Georgia Institute of Technology), Evan B. Stuart (Georgia Institute of Technology), Michael S. Fields (Georgia Institute of Technology), Eric R. Hein (Georgia Institute of Technology), Mark Wisneski (Georgia Institute of Technology)

In this work, we apply graph metrics to an Active DNS dataset, provided by the Georgia Tech Astrolavos Lab, to characterize and identify malicious actors within the DNS infrastructure of the Internet. Active DNS uses a set of data sources - including TLD Zone Files, public blacklists, and many others - to gather approximately 250 million DNS records per day. We use A records from this dataset to form a graph. From this graph we compute various metrics for each vertex using the following algorithms: PageRank, betweenness centrality, k-core decomposition, and in- and out-degree. Using Random Forest, we identify malicious nodes within Active DNS. This work uses the high-performance streaming graph analysis platform, STINGER for rapid generation of graph-based features for hundreds of millions of data domain names and IP addresses.

Poster #86

STRUMPACK: Scalable Preconditioning Using Low-Rank Approximations and Random Sampling

Pieter Ghysels (Lawrence Berkeley National Laboratory), Xiaoye S. Li (Lawrence Berkeley National Laboratory), Christopher Gorman (University of California, Santa Barbara), Francois-Henry Rouet (Lawrence Berkeley National Laboratory)

We present a parallel and fully algebraic preconditioner based on an approximate sparse factorization using rank-structured matrix compression. In sparse multifrontal LU factorization, the fill-in occurs in dense frontal matrices. These are approximated as Hierarchically Semi-Separable (HSS) rank-structured matrices using an efficient randomized sampling technique. The resulting fast solver or preconditioner has optimal or close to optimal complexity – in terms of floating point operations and memory usage – for matrices from several types of discretized partial differential equations. Our STRUMPACK approximate solver is a viable alternative to other state-of-the-art preconditioners like ILU and AMG, and we demonstrate that our new method is more robust and scalable than existing ones. Our poster presents results with the distributed memory MPI+OpenMP code for DOE applications on supercomputers from NERSC.

Applications

Poster #7

Optimizing Turbomachinery CFD applications for Modern Multi-Core and Accelerator HPC Systems

Christopher P. Stone (Computational Science and Engineering LLC), Daryl Y. Lee (University of California, Davis), Roger L. Davis (University of California, Davis)

MBFLO3 is a general-purpose, multi-disciplinary simulation code focused on turbomachinery applications including turbulent conjugate heat transfer. The multi-block, structured mesh algorithm has been refactored to perform efficiently on multicore CPUs, NVIDIA Kepler GPUs, and Intel Xeon Phi accelerators by extending the existing distributed memory (MPI) parallelism with finer-grained thread and data parallelism. Thread and data parallelism was implemented using OpenACC and OpenMP compiler directives. Significant memory and loop restructuring was necessary to achieve high performance on all three platforms. Refactoring lead to improvements of over 2x on the host CPU and more than 4x on the Xeon Phi accelerator. OpenMP threading on the multi-core host CPU increased the parallelism by an order of magnitude with over 60% parallel efficiency.

Poster #8

Spectral Domain Decomposition Using Local Fourier Basis: Application to Ultrasound Simulation on a Cluster of GPUs

Jiri Jaros (Brno University of Technology), Filip Vaverka (Brno University of Technology), Bradley E. Treeby (University College London)

The simulation of ultrasound wave propagation through biological tissue has a wide range of practical applications. However, large grid sizes are generally needed to capture the phenomena of interest. Here, a novel approach to reduce the computational complexity is presented. The model uses an accelerated k-space pseudospectral method which enables more than one hundred GPUs to be exploited to solve problems with more than 3 x 10^9 grid points. The classic communication bottleneck of Fourier spectral methods, all-to-all global data exchange, is overcome by the application of domain decomposition using local Fourier basis. Compared to global domain decomposition, for a grid size of 1536 x 1024 x 2048, this reduces the simulation time by a factor of 7.5 and the simulation cost by a factor of 3.8.

Poster #9

Tsunami Run-Up and Inundation Simulations Using LexADV_EMPS Solver Framework on Fujitsu FX100

Masao Ogino (Nagoya University), Hongjie Zheng (Toyo University), Kohei Murotani (University of Tokyo), Seiichi Koshizuka (University of Tokyo), Ryuji Shioya (Toyo University)

In this research, we describe applications of the LexADV_EMPS solver framework to tsunami run-up and inundation simulations on the Fujitsu FX100. To achieve high performance of particle-based simulations, the LexADV_EMPS framework supports hierarchical domain decomposition, halo exchange pattern of communication, and dynamic load balancing on the distributed-memory parallel computers. By using our framework, we have been successfully performed large-scale numerical simulations of tsunami run-up and inundation with hundreds of millions of particles for the purpose of the tsunami impact assessment.

Poster #10

Optimizing Application I/O by Leveraging the Storage Hierarchy Using the Scalable Checkpoint Restart Library with a Monte Carlo Particle Transport Application on the Trinity Advanced Computing System

Michael M. Pozulp (Lawrence Livermore National Laboratory), Gregory B. Becker (Lawrence Livermore National Laboratory), Patrick S. Brantley (Lawrence Livermore National Laboratory), Shawn A. Dawson (Lawrence Livermore National Laboratory), Kathryn Mohror (Lawrence Livermore National Laboratory), Adam T. Moody (Lawrence Livermore National Laboratory), Matthew J. O'Brien (Lawrence Livermore National Laboratory)

The poster accompanying this summary exhibits our experience using the Scalable Checkpoint Restart library (SCR) to achieve I/O speedups during checkpoint and restart. We ran Lawrence Livermore National Laboratory's (LLNL) Monte Carlo particle transport code, Mercury, on Trinity at Los Alamos National Laboratory (LANL). We performed a weak scaling study and observed speedups at 16 nodes and above, including a 30x maximum speedup at 4096 nodes. We benchmarked read performance by restarting from the checkpoints we wrote and observed speedups for 11 out of 12 counts, including a 9x maximum speedup at 2048 nodes. Finally, we ran a user problem in which using SCR reduced median time-to-checkpoint by 20x. Our results show that leveraging the storage hierarchy is necessary for optimizing application I/O.

Poster #11

GPU Acceleration of a Non-Hydrostatic Ocean Model with Lagrangian Particle Tracking

Takateru Yamagishi (Research Organization for Information Science and Technology), Yoshimasa Matsumura (Hokkaido University)

To achieve detailed simulations of several different types of particles in numerical ocean simulations, we have implemented and optimized a non-hydrostatic ocean model with Lagrangian particle tracking on an NVIDIA GPU. We have revised the algorithm for particle tracking and optimized the kernels for ocean dynamics calculation. The sorting of all particles at every time step was introduced to coalesce the access to the GPU global memory, and a texture cache was assigned to the ocean current velocity array to accelerate particle tracking. Thread-level parallelism was exploited with the help of additional calculations, and registers were effectively used for the ocean dynamics calculation. When comparing its execution on a Fujitsu SPARC64 IXfx to that on the NVIDIA K20C, the GPU-implemented model was three times faster. This model successfully reproduced the nonlinear distribution of the particles in the ocean.

Poster #12

Performance of Popular HPC Applications on the Intel Knights Landing Platform

Antonio Gomez-Iglesias (University of Texas at Austin), Feng Chen (University of Texas at Austin), Lei Huang (University of Texas at Austin), Hang Liu (University of Texas at Austin), Si Liu (University of Texas at Austin), Antia Lamas-Linares (University of Texas at Austin), John Cazes (University of Texas at Austin), Carlos Rosales (University of Texas at Austin)

We present direct performance measurements for six popular HPC applications in the Knights Landing (KNL) platform. Performance numbers for Sandy Bridge and Haswell processors are provided for contrast. The applications (NAMD, Gromacs, FLASH4, WRF, Quantum Espresso, and NCBI BLAST) were selected from among the ten most used in the Stampede supercomputer at the Texas Advanced Computing Center and, given their diversity, should be representative of typical HPC workloads. All runs were performed with publicly available codes without modification -- except a single line added to FLASH4 to enable threading in a given code section -- and so results should be expected to improve as developers gain access to KNL. Current results are promising, with execution on a single KNL processor showing speedups up to 2.7X with respect to a dual socket Sandy Bridge and up to 1.7x with respect to a dual socket Haswell.

Poster #13

Accelerating PETSc-Based CFD Codes with Multi-GPU Computing

Pi-Yueh Chuang (George Washington University), Lorena A. Barba (George Washington University)

We wrote a wrapper code to bridge PETSc and AmgX libraries to use AmgX's multi-GPU linear solvers in existing PETSc-based CFD codes. With the wrapper, those codes are able to exploit all available CPU and GPU resources without heavy coding efforts. The wrapper features a simple usage: the two functions for setting and solving a linear system in the wrapper can directly replace the same functions in PETSc. Data conversion, transfer, scatters and gathers, and MPI communications are all taken care of. Benchmarks with real CFD applications show that, with multi-GPU computing, we can save 1) run times and hardware cost (e.g. a 6-CPU-core workstation with 1 NVIDIA K40c can compete with a 16-node CPU cluster); and 2) cloud HPC cost (e.g. a benchmark on Amazon EC2 shows a 16x cost saving between CPU and GPU clusters).

Poster #14

Scaling a High Energy Laser Physics Application (VBL) using MPI and the RAJA Portability Layer

Kathleen McCandless (Lawrence Livermore National Laboratory), Tom Epperly (Lawrence Livermore National Laboratory), Jean Michel Di Nicola (Lawrence Livermore National Laboratory), Katie Lewis (Lawrence Livermore National Laboratory), Gabriel Mennerat (Lawrence Livermore National Laboratory), Jarom Nelson (Lawrence Livermore National Laboratory), Rick Sacks (Lawrence Livermore National Laboratory), Samuel Schrauth (Lawrence Livermore National Laboratory), Paul Wegner (Lawrence Livermore National Laboratory)

LLNL is a world leader in designing and maintaining high energy lasers, built upon decades of leadership in the modeling of high energy laser systems. Here we present initial results for a parallel mini-app based on the National Ignition Facility's (NIF) Virtual Beamline (VBL) code, a single-node laser physics modeling engine. Recent advances in ultra-intense short-pulse laser systems are driving us to develop massively parallel laser physics capabilities similar to the laser physics code Miró (an MPI-only implementation) to support the multi-order increase in time/space resolution needed for these types of broadband, chirped-pulse amplification lasers. Here we present a demonstration of our new scalable simulation code architecture using MPI and the RAJA Portability Layer. This hybrid parallelization approach promises to bridge the gap in resolution allowing us to deliver future simulations with the requisite physics fidelity at an unprecedented scale.

Poster #15

Laminar Unsteady Navier-Stokes Flow on Multicore Architectures

Bahareh Mostafazadeh Davani (University of California, Irvine), Ferran Marti (University of California, Irvine), Feng Liu (University of California, Irvine), Aparna Chandramowlishwaran (University of California, Irvine)

Stencil computations are at the core of many scientific applications, and given the fact that these computations are often memory-bound, they eventually become the limiting portion of these applications. Such limitations give rise to the increasing focus on optimizing stencil computations. A vast number of improvements have been applied to stencil kernels, but very few of these attempts have applied the optimized stencils in real applications. Our work is unique in the sense that we target an entire solver that captures the interaction between multiple stencil patterns and one that is capable of simulating real applications. Our implementation solves the fluid motion of compressible viscous flow at transonic speeds. In this work, we take advantage of the full potential of modern multicore architectures to achieve an efficient and scalable Computational Fluid Dynamics solver. Our parallel solver achieves more than 40x speedup over a baseline C++ code.

Poster #16

NEMO5, Xeon Phi, and hStreams: Physics of Ultrascaled 2D Nanotransistors

Xinchen Guo (Purdue University), Kuang-Chung Wang (Purdue University), James Charles (Purdue University), Junzhe Geng (Purdue University), Daniel Mejia (Purdue University), Daniel Valencia (Purdue University), Daniel Lemus (Purdue University), James E. Fonseca (Purdue University), Gerhard Klimeck (Purdue University), Tillmann Kubis (Purdue University)

The detailed electrical and thermal properties of future ultrascaled transistors are critically dependent on quantum effects. This is particularly true for the promising class of 2D-material based transistors. NEMO5 – a multipurpose, multiscale semiconductor device simulation tool is applied on these devices in many academic and industrial groups (including Intel, Samsung, TSMC etc.). Harnessing the power of modern HPC hardware in world's largest supercomputing centers, NEMO5 is able to unveil more fine details of nanotransistors. This work discusses how a physical Xeon Phi is partitioned into several virtual Xeon Phis using the hStreams library and put into a MPI parallelization scheme.

Poster #97

Developing A Scalable Platform For Next-Generation Sequencing Data Analytics Over Heterogeneos Clouds and HPCs : A Case for Transcriptomes and Metagenomes

Shayan Shams (Louisiana State University), Nayong Kim (Louisiana State University), Ming-Tai Ha (Rutgers University), Shantenu Jha (Rutgers University), Jian Tao (Louisiana State University), Ramesh Subramanian (Louisiana State University), Vladmir Chouljenko (Louisiana State University), K. Gus Kousoulas (Louisiana State University), Ram J. Ramanujam (Louisiana State University), Seung-Jong ParK (Louisiana State University), Joohyun Kim (Louisiana State University)

A novel scalable pipeline for metagenome/transcriptome is presented. Thanks to the underlying distributed computing platform, a significant roadblock in Next-Generation Sequencing data analytics, associated with ever-growing and noisy data sets, can be effectively resolved.

On top of the core feature for accessing and utilizing heterogeneous distributed computing resources including HPCs and Clouds (EC2, OpenStack-based, and IBM Bluemix), the distributed application runtime environment is built for efficient management of massive workloads and data processing tasks by leveraging high-end HPC technologies, emerging Hadoop-based software models, and DOCKER. The consequently available repertoire of options for flexible and scalable runtime scenarios constitutes the pipeline for dealing with any size of data sets. In order to maximize benefits from the scalable platform, a novel method was developed for de novo genome sequence reconstruction with Multiple Assembly Multiple Parameter (MAMP) and available with the pipeline. Preliminary results indicate great potentials of MAMP.

Poster #98

First Experiences with ab initio Molecular Dynamics on OpenPOWER: The Case of CPMD

Valery Weber (IBM), A. Cristiano I. Malossi (IBM), Ivano Tavernelli (IBM), Teodoro Laino (IBM), Costas Bekas (IBM), Manish Modani (IBM), Nina Wilner (IBM), Tom Heller (IBM), Alessandro Curioni (IBM)

The algorithmic adaptation and code re-engineering to port CPMD code to next-generation heterogeneous Open-POWER architectures (CPU+GPU) is presented here. The construction of the electronic density, the application of the potential to the wavefunctions, and the orthogonalization procedure are offloaded to the GPU. The different GPU kernels consist mainly of fast Fourier transforms (FFT) and basic linear algebra operations (BLAS). The performance of the new implementation obtained on Firestone (POWER8/Tesla) is discussed. The communication between the host and the GPU contributes a large fraction of the total run time. We expect a strong attenuation of the communication bottleneck when the NVLink high-speed inter-connect, between CPU to GPU will be available. These results will be added in the final version of the poster for SC16.

Poster #99

A Fast Implicit Solver with Low Memory Footprint and High Scalability for Comprehensive Earthquake Simulation System

Kohei Fujita (RIKEN), Tsuyoshi Ichimura (University of Tokyo), Kentaro Koyama (Fujitsu Ltd), Masashi Horikoshi (Intel Corporation), Hikaru Inoue (Fujitsu Ltd), Larry Meadows (Intel Corporation), Seizo Tanaka (University of Tsukuba), Muneo Hori (University of Tokyo), Lalith Maddegeddara (University of Tokyo), Takane Hori (Japan Agency for Marine-Earth Science and Technology)

We developed a comprehensive earthquake simulation system that improves the reliability of conventional earthquake disaster estimates by significant speedup of unstructured implicit finite-element simulations needed to solve nonlinear wave-propagation problems in complex-shaped domains. A fast implicit solver with a low-memory footprint was developed using an algorithm combining a multi-grid, mixed-precision, communication avoiding inexact LU precon-

ditioning, and element-by-element method with novel SIMD-buffering and multi-core coloring methods. Using this solver with 88.8% size-up scalability up to 663,552 CPU cores of the full K computer, a challenging practical 2.05 Tera degrees-of-freedom problem (205 times more degrees-of-freedom compared to the current state-of-the-art) was solved 30 times faster compared with a memory-efficient SC14 solver. This poster presents the algorithmic design, implementation, performance, and portability of this fast and memory-efficient implicit solver.

Poster #100

Civet: A Framework for Reproducible Bioinformatics Analysis

Glen Beane (Jackson Laboratory), Al Simons (Jackson Laboratory)

Bioinformatics analysis often combines multiple publicly available command line tools with custom tools and filters into an analysis pipeline. Analysts often script together tools into an ad hoc solution for the task at hand without designing a robust and easily maintainable pipeline. At The Jackson Laboratory we had a requirement to be able to quickly develop multiple easily maintainable, robust, and reproducible pipelines to support new clinical programs. To address this, we designed and implemented a framework, called Civet, for building bioinformatics pipelines. Civet ensures consistency and reliability for pipelines developed by different analysts throughout our organization, and because it was developed with clinical pipelines in mind, we included unique validation capabilities into Civet. Civet has the capability to verify that all the binaries, shared libraries, and reference files used in the execution of the pipeline are unchanged since the pipeline was validated.

Poster #101

Network-Optimized Distributed Memory Parallel Breadth-First Search

Praveen Sharma (University of Southern California)

Graphs are a very common modeling tool used for many applications including mapping relations in physical, social or information systems. In the worst case, the complexity of Breadth-first search (BFS) is linear in the number of edges and vertices, and the conventional top-down approach always takes as much time as the worst case. The bottom-up approach recently proposed manages to cut down the complexity all the way to the number of vertices in the best case, which is typically at least an order of magnitude less than the number of edges. However these algorithms can better exploit the network usage in terms of the time spent transmitting. In this paper, we propose an improved version of the direction-optimized distributed breadth-first search algorithm to further reduce the network traffic and improve performance in a multi-processor environment. Our implementation reduces network traffic by 11% and results in general speedups of 6-9%.

Poster #102

Improved Global Weather Prediction with GFDL's FV3 Dynamical Core

Shannon Rees (Engility Corporation)

Global weather models are extremely computationally intensive, especially when run at high resolutions. With increasing HPC resources and the desire for better forecasts, the resolutions these models are demanded to run at is always increasing. The current U.S. National Weather Service operational Global Spectral Model (GSM) is a hydrostatic model, and it has reached its end-of-life. The field of global weather modeling has reached a point where resolutions are so high that non-hydrostatic, or cloud-resolving, models are required. The Next Generation Global Prediction System (NGGPS) is being built to replace the GSM with a model that will improve forecasts and extend the predictability range out to 30 days, while keeping pace with changes in HPC resources over the next two decades. The Geophysical Fluid Dynamics Laboratory's Finite Volume dynamical core was chosen, through a two-year program of rigorous testing, to be part of this NGGPS.

Poster #103

Neuroscience Gateway – Understanding the Scaling Behavior of NEURON Application

Subhashini Sivagnanam (San Diego Supercomputer Center), Amit Majumdar (San Diego Supercomputer Center)

In this poster, we describe the Neuroscience Gateway (NSG) that has enabled HPC access for the computational neuroscience community since 2013. We will also discuss the scaling performance on various HPC architectures for NEURON application. A central challenge in neuroscience is to understand how brain function emerges from interactions of a large number of biological processes at multiple physical and temporal scales. Computational modeling is an essential tool for developing this understanding. Driven by a rapidly expanding body of empirical observations, models and simulation protocols are becoming increasingly complex. This has stimulated development of powerful, open source computational neuroscience simulators, which run efficiently on HPC systems. NEURON is one such simulator that is widely used by computational neuroscientists. The scaling performance of NEURON on KNL, FX10 will be discussed in this poster.

Poster #104

MPI-GIS : High Performance Computing and I/O for Spatial Overlay and Join

Satish Puri (Marquette University)

For certain GIS and Spatial Database applications, spatial overlay and join on two or more layers of geo-spatial data may be necessary. However, using a sequential paradigm to process them is time-consuming. For large datasets, I/O, spatial indexing, and geometric refinement phase are time consuming. These operations involve irregular I/O due to varying number of vertices in different shapes and irregular computations without well-defined communication pattern due to irregular spatial and/or temporal task or data distributions. These irregularities makes parallelization, partitioning, and load balancing more challenging.

We have undertaken parallelization of polygon clipping and overlay algorithms, and spatial join using GPU and MPI. We briefly describe our MPI-GIS system and highlight our work on parallel I/O for OpenStreetMap data. We also present a parallel algorithm for Intersection of polygons which is an elementary operations in polygon overlay. Its time complexity is O((n+k) logn) where k is the number of intersections.

Poster #105

A Scalable Approach for Topic Modeling with R

Tiffany A. Connors (Texas State University), Ritu Arora (University of Texas at Austin)

Topic Modeling (TM) helps in automatically classifying documents under different topics, and is especially useful for exploring a large corpus of documents to discover new relationships. The R programming language has a TM library that is easy to install and use. However, due to its interpreted nature, the performance of TM code written in R is poor as compared to the same code rewritten in C/C++/Fortran. Despite its poor performance, R is a high-productivity language that is commonly used by non-traditional High Performance Computing (HPC) users to do TM and other similar data analyses. Many such users do not have access to expertise for rewriting their R code in C/C++/Fortran but have large datasets to analyze. With TM as an example, we demonstrate that such end-users can reduce the time-to-results (sometimes, up to a factor of 23) by running their R code in High-Throughput Computing (HTC) mode on HPC resources.

Poster #106

Extreme Fidelity Computational Electromagnetic Analysis in the Supercomputer Era

Brian MacKie-Mason (University of New Mexico), Zhen Peng (University of New Mexico), Christopher Kung (Engility Corporation)

Ever-increasing fidelity and accuracy needs for advanced electromagnetic applications have pushed problem sizes toward extreme scales. It puts a high premium on parallel and scalable algorithms with optimal computational complexity. This poster displays the research into high-performance, geometry-aware domain decomposition (DD) methods for the solution of time-harmonic Maxwell's Equations. The technique ingredients include a volume-based optimized Schwarz finite element DD method, and a surface-based interior penalty boundary element DD method.

The work has three benefits: (i) it results in robust, cost-effective preconditioning techniques that reduce the condition number of very large systems of equations; (ii) it provides a flexible and natural way to set up the mathematical models, to create the problem geometries and to discretize the computational domain; (iii) it leads to parallel and scalable algorithms to reduce the time complexity of extreme-scale simulations. The capability of the algorithms is illustrated through real-world applications on high performance computing systems.

Architecture and Networks

Poster #1

Comparison of High Performance Network Options: EDR InfiniBand vs. 100Gb RDMA Capable Ethernet

Kari N. Erickson (Los Alamos National Laboratory), Faith V. Van Wig (Los Alamos National Laboratory), Luke A. Kachelmeier (Los Alamos National Laboratory)

InfiniBand (IB) has long been the network of choice for high performance computing (HPC). However, advancements in Ethernet and IB technology, as well as other high-performance networks, have made it necessary to analyze the performance of these network options in detail – specifically, we look at 100Gb Ethernet and IB. Advancements in Ethernet include upwards of 100Gb data rates and standardization of RDMA-over-Converged-Ethernet (Routable RoCE). Similarly, IB has introduced Enhanced Data Rate (EDR) hardware, which nearly doubles previous bandwidth, increasing it to 100Gb. The goal of this study is to compare and contrast these two options by looking at their respective bandwidth and latency performance, as well as message injection rates and deployment effort. This research will allow a clear definition of the solution space for the challenges and problems being faced by networks with HPC workloads. LA-UR-16-25359

Poster #2

Scalable Communication Architectures for GPU-Centric Systems

Benjamin Klenk (University of Heidelberg), Holger Fröning (University of Heidelberg)

Heterogeneity in computing has enabled higher performance and increased energy efficiency in the past years. Accelerators, such as GPUs, have been deployed to offload compute-intensive tasks. However, the CPU has been the main processor, delegating work to the GPU and handling communication with other nodes. This model is changing, as the latest GPUs bear more responsibilities with hardware-assisted unified memory and integrated NVLink, enhancing the GPU into a peer device.

Allowing GPUs to autonomously source and sink network traffic seems promising, as costly interactions with CPUs can be avoided. Communication can be offloaded to dedicated networking hardware or performed on the GPU on top of shared memory models like Nvidia's NVLink. This poster shows that for GPU-to-GPU traffic, offloading is always superior to CPU-controlled communication. We also present some early results and insights of communication being processed on the GPU without additional networking hardware.

Poster #3

Acceleration of All-to-All Communication on Multi-Layer Full Mesh, Low-Cost Connectable Network Topology

Toshihiro Shimizu (Fujitsu Ltd), Masahiro Miwa (Fujitsu Ltd), Kohta Nakashima (Fujitsu Ltd)

Recently, due to the increasing scale of computation, more and more servers are used to make calculations simultaneously and communicate with each other. These servers are connected by links and switches. Since the cost of the switch is relatively high, reduction of the number of switches is desirable to realize a larger cluster system cost-efficiently. We have already proposed the multilayer full mesh (MLFM) topology in this purpose and showed that MLFM topology can connect more servers than conventional fat tree topology and can achieve congestion-free all-to-all communication using all servers.

In this poster we propose a method of congestion-free all-to-all communication using part of the servers on MLFM. This is necessary because users of supercomputer systems typically use part of the servers rather than all of them. Our experimental results show 2.2 times higher throughput compared to the conventional communication pattern.

Poster #4

Exploring Randomized Multipath Routing on Multi-Dimensional Torus Networks

Prajakt Shastry (Illinois Institute of Technology), Daniel Parker (University of Chicago), Sanjiv Kapoor (Illinois Institute of Technology), Ioan Raicu (Illinois Institute of Technology)

Network performance is a critical aspect of HPC, and improving performance is a major goal in the design of future systems; this work proposes to improve network performance through new routing algorithms, leveraging the rich multi-path topologies of multi-dimensional torus networks commonly found in supercomputers built in the past fifteen years. Virtually all torus networks in production today utilize the dimension order routing algorithm, which is essentially a static and deterministic routing strategy to allow internode communication. This static routing strategy has significant load balancing implications, leading to sub-par performance. We propose a new Random Distance Routing algorithm, which randomly distributes packets to different neighboring nodes that are closer to the destination, leading to global load balanced network. Through the CODES/ROSS [4] simulator, we show that the proposed randomized multi-path routing algorithm can increases throughput of a 5D-Torus network by 1.6X, as well as reduce latency by 40%.

Poster #5

GPU Approximation Acceleration For Scientific Applications

Ang Li (Pacific Northwest National Laboratory), Shuaiwen Leon Song (Pacific Northwest National Laboratory)

Approximate computing, the technique that sacrifices certain amount of accuracy in exchange for substantial performance boost or power reduction, is one of the most promising solutions to enable power control and performance scaling towards exascale. In this poster, we introduce a transparent, tractable, and portable design framework for SFU-driven approximate acceleration on GPUs, providing fine-grained tuning for performance and accuracy trade-offs. Our design is software-based and requires no hardware or application modifications.

Poster #6

Training Restricted Boltzmann Machines Using a Quantum Annealer

Vaibhaw Kumar (Booz Allen Hamilton), Gideon P. Bass (Booz Allen Hamilton), Joseph S. Dulny (Booz Allen Hamilton)

Machine learning and the optimization involved therein is of critical importance for commercial and military applica-

tions. Due to the extremely complex nature of many-variable optimization, the conventional approach is to employ a meta-heuristic technique to find suboptimal solutions. Quantum Annealing (QA) hardware offers a completely novel approach for obtaining significantly better or optimal solutions with considerably large speed-ups when compared to traditional computing hardware. In this presentation, we describe our development of new machine learning algorithms tailored for QA hardware. We train a Restricted Boltzmann Machine (RBM) using QA hardware as a sampler. We present our initial results obtained by training RBMs on model data sets. We also discuss strategies for scaling up, including enhanced embedding and partitioned RBMs, to overcome the limitation imposed by current QA hardware.

Poster #107

Evaluating Best and Worst Case Scenarios on Two-Level Memory Systems

Ryan J. Huber (University of Minnesota), Edgar A. Leon (Lawrence Livermore National Laboratory)

To achieve the capacity and bandwidth requirements of an exascale memory system, vendors are employing multiple levels of memory: a small high-bandwidth memory close to the processor and a large but low-bandwidth memory. To leverage the high-bandwidth effectively, application developers could explicitly place data structures in fast memory, but this becomes impractical for large HPC codes.

Our long-term objective is to develop metrics to semi-automatically identify candidate data structures to place in fast memory. In this poster, we quantify upper and lower bounds on the potential performance gain/loss of a two-level memory system. We also show that a data placement policy guided by an experienced application developer may not lead to a significant performance improvement, motivating our future work in this area. Finally, we demonstrate how a single-level memory system can provide meaningful insight on the effect of data placement policies before porting a code to a two-level memory system.

Poster #108

Designing Accelerators for Data Analytics: A Dynamically Scheduled Architecture

Marco Minutoli (Pacific Northwest National Laboratory), Vito Giovanni Castellana (Pacific Northwest National Laboratory), Antonino Tumeo (Pacific Northwest National Laboratory), Marco Lattuada (Polytechnic University of Milan), Fabrizio Ferrandi (Polytechnic University of Milan)

Conventional High Level Synthesis (HLS) tools mainly target compute intensive kernels typical of digital signal processing applications. We are developing techniques and architectural templates to enable HLS of data analytics applications. These applications are memory intensive, present fine-grained, unpredictable data accesses, and irregular, dynamic task parallelism. We introduce a dynamic task scheduling approach to efficiently execute heavily unbalanced workloads, at the opposite of conventional HLS flows that employ execution paradigms based on static scheduling. Our approach is validated by analyzing and synthesizing queries from the Lehigh University Benchmark (LUBM), a well know SPARQL benchmark.

Poster #109

Concurrent Dynamic Memory Coalescing on GoblinCore-64 Architecture

Xi Wang (Texas Tech University), John Leidel (Texas Tech University), Yong Chen (Texas Tech University)

The majority of modern microprocessors are architected to utilize multi-level data caches as a primary optimization to reduce the latency and increase the perceived bandwidth from an application. However, applications that exhibit random or non-deterministic memory access patterns often induce a significant number of data cache misses, thus reducing the natural performance benefit from the data cache.

In response to the performance penalties inherently present with non-deterministic applications, we have constructed a unique memory hierarchy within the GoblinCore-64 (GC64) architecture explicitly designed to exploit memory performance from irregular memory access patterns with RISC-V-based core and Hybrid Memory Cube (HMC) devices.

In this work, we present two parallel methodologies and associated implementations for coalescing non-deterministic memory requests into the largest potential HMC request by constructing a binary tree-based memory coalescing model. Cogent test results are also presented to further convince the outstanding efficacy of this concurrent DMC design in GoblinCore-64 architecture.

Poster #110

A Comparative Power-Performance Analysis of Microarchitecture Effects on Heterogeneous CPU-GPU

Vijayalakshmi Saravanan (University at Buffalo), Sridhar Ramalingam (University at Buffalo)

High-Performance Computing (HPC) users are exploring heterogeneous computing through the integration of CPU-GPUs to maximize the computational throughput. The performance of the processors depends on many micro-architectural parameters such as issue-width, functional units, and pipeline

depth. The increasing number of processor cores and depth of instruction pipelines continues to add complexity to task/thread parallelism and the ratio of power/performance. To adequately address these issues and potential benefits and pitfalls that may arise from this heterogeneous processors, it is important to have a deep understanding of application-level and microarchitecture-level demands from CPU-GPU cores. Multi-core CPUs have already been studied thoroughly owing to their larger history in the field. This paper aims to evaluate the scalability of CPU-GPU per stream cores and micro-architectural behavior for parallel applications executing on each core type in heterogeneous CPU-GPU processor simulation. To obtain this, we conduct a set of detailed benchmarks for many-core systems from PARSEC.

Poster #111

A Comparison of x86 Computer Architecture Simulators

Ayaz Akram (Western Michigan University), Lina Sawalha (Western Michigan University)

Simulation is used as a primary performance evaluation methodology in most of the papers published in top computer architecture conferences. There is not much literature dealing with the evaluation of simulators by comparing them to each other and to the state-of-the-art processors. The absence of performance validation of simulators may cause experimental errors that can result in incorrect conclusions, if the errors are large. This work provides a detailed simulation accuracy comparison of four modern computer architecture simulators: gem5 , Sniper, Multi2sim and PTLsim. We configured these simulators to model Intel's high-performance processor, Core-i7 Haswell microarchitecture based CPU. Then we quantified the experimental errors and evaluated the accuracy of such an approach.

Poster #112

Accelerated Particle-Grid Mapping

Ahmed Sanaullah (Boston University)

Charge Mapping is critical to electrostatic computations for Molecular Dynamics. It reduces the complexity of evaluating long-range coulombic forces by diffusing discrete particle charges to a regular grid. Efficient charge mapping on accelerators (GPUs, FPGAs) is non-trivial, with the compute and memory intensive nature of the algorithm limiting performance benefits in naive implementations. On FPGAs, resource constraints have only allowed low order (bicubic) interpolations [5]. In our work, we explore methods for improving the performance on both platforms. These include application specific data structures and low complexity kernels for GPUs and deep pipelines with interleaved memory access for FPGAs. Our best case implementation shows > 62x speed-up over existing CPU codes and >30x speed-up over existing GPU codes. We also find that, when using the Altera Arria 10, high resource availability enables the building of a balanced accelerator for the entire long-range electrostatics computation on a single FPGA.

Clouds & Distributed Computing

Poster #90

Fault-Tolerant Scheduler for Shareable Virtualized GPU Resource

Daeyoun Kang (Korea Advanced Institute of Science and Technology), Tae Joon Jun (Korea Advanced Institute of Science and Technology), Daeyoung Kim (Korea Advanced Institute of Science and Technology)

Recently container-based virtualization is variously used to maximize utilization of computer resource, along with traditional Virtual Machine. However, different from traditional resources, GPU was hard to be shared by multiple containers. Lately, a GPU can be shared by multiple containers using volume share feature. In addition, high-end GPU like NVIDIA K20 supports Hyper-Q which allows multiple CPU processes to access a single GPU. However, problems still exist because of GPU's distinctive characteristics. Unlike system memory, GPU memory cannot be swappable. Also, GPU kernels in single Streaming Microprocessor cannot be switched during its running. These restrictions make it hard to share a GPU by multiple containers, and may result in deadlock situation. In this poster, we propose an interface for new fault-tolerant scheduler considering GPU memory usage. We have implemented this interface to restrict the usage of GPU memory for each container to prevent deadlock situation.

Poster #91

Lazer: A Memory-Efficient Framework for Large-Scale Genome Assembly

Sayan Goswami (Louisiana State University), Arghya Kusum Das (Louisiana State University), Richard Platania (Louisiana State University), Kisung Lee (Louisiana State University), Seung-Jong Park (Louisiana State University)

Genome sequencing technology has witnessed tremendous progress both in terms of throughput as well as the cost per base pair. However, when it comes to sequence assembly, there still exists a dilemma in the state-of-the-art technology. On one hand, we have a number of distributed assemblers that can utilize several nodes but require massive amounts of memory. On the other hand, there are a few assemblers that can assemble mammalian genomes on a single node but cannot scale up. In this paper, we present a distributed assembler

that is both scalable and memory efficient. Using partitioned De Bruijn graphs we enhance memory-to-disk swapping and reduce network communication in the cluster. Experimental results show that our framework can assemble a human genome dataset (452GB) in 14.5 hours using two nodes and 23 minutes using 128 nodes.

Poster #92

An I/O Load Balancing Framework for Large-Scale Applications (BPIO 2.0)

Sarah M. Neuwirth (University of Heidelberg), Feiyi Wang (Oak Ridge National Laboratory), Sarp Oral (Oak Ridge National Laboratory), Ulrich Bruening (University of Heidelberg)

Designed for capacity and capability, HPC storage systems are inherently complex and shared among multiple, concurrent jobs competing for resources. The lack of centralized coordination and control often render the end-to-end I/O paths vulnerable to load imbalance and contention. With the emergence of data-intensive HPC applications, storage systems are further contended for performance and scalability. BPIO is a topology-aware, load balancing library for mitigating resource contention. This work introduces BPIO 2.0, a dynamic, shared load balancing framework based on BPIO. It transparently intercepts file creation calls during runtime to balance the workload over all available storage targets. The utilization of BPIO 2.0 requires no source code modification and is independent from any I/O middleware. We demonstrate the effectiveness of our framework on the Titan system with a synthetic benchmark in a noisy production environment.

Poster #93

Prototype Implementation of Simulation Caching Framework for Multi-User Interaction

Jiachao Zhang (University of Fukui), Yu Yamamoto (University of Fukui), Shinji Fukuma (University of Fukui), Shin-ichiro Mori (University of Fukui)

In order to realize human-in-the-loop scientific computing in a cloud-like environment, we have to conquer the problem of network latency. For this purpose, we have proposed a simulation model which we referred to it as “Simulation Caching.” Simulation Caching is a sort of cooperative cloud technique where a high performance remote server somewhere on the cloud cooperates with a moderate scale local server to realize interactive steering and remote collaboration over on-going simulation. To hide the latency to the remote server, Simulation Caching lets the local server cache a part of the simulation from the remote server and performs the duplicated simulation concurrently with the remote server, while keeping the accuracy of the cached simulation by weakly cooperating with the original simulation running on the remote server. In this poster, the simulation framework based on simulation caching technique and its extension to multi-user environment are reported.

Poster #94

Toward Understanding HPC- Big Data Convergence Using Cloud Platforms

Shweta Salaria (Tokyo Institute of Technology)

The path to HPC- Big Data convergence has resulted in numerous research that demonstrates the performance-cost tradeoff between running applications on supercomputers and cloud platforms. Previous studies typically focus on evaluating scientific HPC benchmarks on old cloud configurations, failing to consider the new opportunities offered by cloud platforms. We present a comparative study of the performance of representative big data benchmarks, or “Big Data Ogres”, and HPC benchmarks running on supercomputer and cloud. Our work distinguishes itself from previous studies in a way that we explore the new generation of compute-optimized instances oriented for HPC applications by Amazon EC2. Our results indicate that the new cloud environment appears to be a promising HPC system alternative for a set of high performance computing applications.

Poster #95

Marrying HPC and Cloud for Long Term Happiness

Apoorve Mohan (Northeastern University), Ravi S. Gudimetla (Northeastern University), Ata Turk (Boston University), Sourabh Bollapragada (Northeastern University), Rajul Kumar (Northeastern University), Jason Hennessey (Boston University), Evan Weinberg (Boston University), Dimitri Makrigiorgos (Boston University), Christopher N. Hill (Massachusetts Institute of Technology), Gene Cooperman (Northeastern University), Peter Desnoyers (Northeastern University), Richard Brower (Boston University), Orran Krieger (Boston University)

Traditional HPC clusters, with deep job submission queues, by construction are always almost fully utilized. On the other hand, the cloud has time-varying workloads, and the cloud business model depends on being underutilized to instantly support all customer requests. A marriage of these two environments could provide additional resources to HPC users while offering increased utilization to the cloud. In this poster, we present the frameworks we built to enable a successful symbiotic co-existence of HPC and the cloud and showcase the benefits achievable with a prototype deployment.

Data Analytics, Visualization & Storage

Poster #17

GPU Accelerated Surface Reconstruction for Particle-Based Fluids

Wei Wu (Ocean University of China), Hongping Li (Ocean University of China), Tianyun Su (First Institute of Oceanography), Haixing Liu (First Institute of Oceanography)

Animating fluids with particle-based methods has attracted great attention in computer graphics. To render discrete particles, implicit surfaces need to be extracted. But this process is often seen as a bottleneck due to the expensive computation and high memory usage especially for constructing a detailed and artifact-free surface. In this work, a GPU accelerated fluids surface reconstruction method using 2-level grid structure is employed with a scheme of arranging fine surface vertices, which could preserve the spatial locality to facilitate the coalesced memory access on GPU. Meanwhile a parallel cuckoo hashing method is taken to help reduce the memory consumption. A parallel version of the optimized surface reconstruction was performed based on CUDA architecture. In the poster, the whole algorithm was outlined, subsequently with an explanation of the 2-level grid structure. The algorithm efficiency was verified by the performance comparison with other approach.

Poster #18

Experiences with a Burst Buffer at NERSC

Wahid Bhimji (Lawrence Berkeley National Laboratory), Debbie Bard (Lawrence Berkeley National Laboratory), David Paul (Lawrence Berkeley National Laboratory)

NVRAM-based Burst Buffers are an important part of the emerging HPC storage landscape. The National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory recently installed one of the first Burst Buffer systems as part of its new Cori supercomputer, collaborating with Cray on the development of the DataWarp software. NERSC has over 6500 users in 750 different projects spanning a wide variety of scientific applications, including climate modeling, combustion, fusion, astrophysics, computational biology, and many more. The applications of the Burst Buffer at NERSC are therefore also considerable and diverse.

We describe here experiences with the first year of the NERSC Burst Buffer. A number of research projects have had early access to the Burst Buffer and exercise its different capabilities to enable new scientific advancements. We present in-depth performance results and lessons-learnt from these real applications as well as benchmark results and system configuration experiences.

Poster #19

Utilizing In-Memory Storage for MPI-IO

Julian Kunkel (German Climate Computing Center), Eugen Betke (German Climate Computing Center)

In contrast to disk or flash based storage solutions, throughput and latency of in-memory storage promises to be close to the best performance. Kove's XPD offers pooled memory for cluster systems. However, the system does not expose access methods to treat the memory like a traditional parallel file system that offers POSIX or MPI-IO semantics.

Contributions of this poster are: (1) Implementation of a MPI-IO driver for the XPD. (2) Thorough performance evaluation of the XPD using IOR with MPI-IO mode.

This MPI independent file driver enables high-level libraries (HDF5, NetCDF) to utilize the XPD's pooled memory. We demonstrate that the MPI-IO driver is able to efficiently take benefit of the pooled memory by providing it as in-memory storage.

Poster #20

Simulating the Burst Buffer Storage Architecture on an IBM BlueGene/Q Supercomputer

Jian Peng (Illinois Institute of Technology), Ioan Raicu (Illinois Institute of Technology), Michael Lang (Los Alamos National Laboratory)

As the computing power of supercomputers keeps increasing, the I/O subsystem of these machines has become one of the major bottlenecks of the overall performance. Towards solving this problem, the burst buffer storage architecture has been adopted in the next generation supercomputers. In this paper, we simulated the burst buffer storage architecture on an IBM BlueGene/Q supercomputer using the CODES/ROSS simulation framework to study the potential I/O performance improvement with burst buffers. These results are a stepping stone toward studying this new storage architecture on future dragon-fly network based supercomputers, including resource management covering both storage and job scheduling.

Poster #21

In Situ Data Steering on Sedimentation Simulation with Provenance Data

Vítor Silva (Federal University of Rio de Janeiro), José Camata (Federal University of Rio de Janeiro), Daniel de Oliveira (Fluminense Federal University), Alvaro L.G.A. Coutinho (Federal University of Rio de Janeiro), Patrick Valduriez (Inria), Marta Mattoso (Federal University of Rio de Janeiro)

Parallel adaptive mesh refinement and coarsening (AMR) are optimal strategies for tackling large-scale simulations. libMesh is an open-source finite-element library that supports parallel AMR and is used in multiphysics applications. In complex simulation runs, users have to track quantities of interest (residuals, errors estimates, etc.) to control as much as possible the execution. However, this tracking is typically done only after the simulation ends. This paper presents DfAnalyzer, a solution based on provenance data to extract and relate strategic simulation data for online queries. We integrate DfAnalyzer to libMesh and ParaView Catalyst, so that queries on quantities of interest are enhanced by in situ visualization.

Poster #22

Optimizing Search in Un-Sharded Large-Scale Distributed Systems

Suraj Chafle (Illinois Institute of Technology), Jonathan Wu (Washington University in St. Louis), Kyle Chard (University of Chicago), Ioan Raicu (Illinois Institute of Technology)

Distributed file systems and storage networks are used to store large volumes of unstructured data. While these systems support large-scale storage, they create new challenges relating to efficiently discovering, accessing, managing, and analyzing distributed data. At the core of these challenges is the need to support scalable discovery of unstructured data. Traditional search methods leverage centralized and globally sharded indexes. We present a distributed search framework that does not rely on sharding and can be applied to a range of distributed storage models. Our approach is built on top of Lucene and utilizes search trees to distribute and parallelize queries. To further optimize query performance we explore methods to prioritize indexes based on size. We evaluate our search framework against alternatives, Grep and Solr, comparing our hierarchical query distribution with a centralized model. Our implementation proved to be faster and scale better.

Poster #23

Large Histogram Computation for Normalized Mutual Information on GPU

Sophie Voisin (Oak Ridge National Laboratory), Devin A. White (Oak Ridge National Laboratory), Jeremy S. Archuleta (Oak Ridge National Laboratory)

Our Photogrammetric Registration of Imagery from Manned and Unmanned Systems pipeline requires accurate computation of multiple Normalized Mutual Information (NMI) coefficients to perform a registration process in two steps, referred as global localization and registration refinement. Computing the NMI coefficients requires generating large joint-histograms to compute the images' joint-entropy, which allows performing an exhaustive search of all possible translations between two images for global localization, and to perform accurate keypoint matching for registration refinement.

Contrary to existing GPU implementation our kernels uniquely accommodate the specificity of each step. They both compute all the NMI coefficients at once using 65536-bins joint-histograms. However the first kernel has the particularity to use a mask of valid pixels for the computation while the second kernel massively compute a multitude of small registration problem corresponding to each keypoint matching with intensity values as descriptor.

The poster will provide implementation details and performance for both kernels.

Poster #24

HPC Enabled Data Analytics for High-Throughput High-Content Cellular Analysis

Ross A. Smith (Engility Corporation), Rhonda J. Vickery (Engility Corporation), Jack Harris (Engility Corporation), Thomas Wischgoll (Wright State University)

Biologists doing high-throughput high-content cellular analysis are generally not computer scientists or high performance computing (HPC) experts, and they want their workflow to support their science without having to be. We describe a new HPC enabled data analytics workflow with a web interface, HPC pipeline for analysis, and both traditional and new analytics tools to help them transition from a single workstation mode of operation to power HPC users. This allows the processing of multiple plates over a short period of time to ensure timely query and analysis to match potential countermeasures to individual responses.

Poster #25

Big Data Helps Particle Physicists to Concentrate on Science

Saba Sehrish (Fermi National Laboratory), Jim Kowalkowski (Fermi National Laboratory), Oliver Gutsche (Fermi National Laboratory), Matteo Cremonesi (Fermi National Laboratory), Bo Jayatilaka (Fermi National Laboratory), Cristina Mantilla (Fermi National Laboratory), Jim Pivarski (Princeton University), Alexy Svyatkovskiy (Princeton University)

In this poster, we evaluate Apache Spark for High Energy Physics (HEP) analyses using an example from the CMS experiment at the Large Hadron Collider (LHC) in Geneva, Switzerland. HEP deals with the understanding of fundamental particles and the interactions between them and is a very compute- and data-intensive statistical science. Our goal is to understand how well this technology performs for HEP-like analyses. Our use case focuses on searching for new types of elementary particles explaining Dark Matter in the universe. We provide different implementations of this analysis workflow; one using Spark on the Hadoop ecosystem, and the other using Spark on high performance computing platforms. The analysis workflow uses official experiment data formats as input and produces publication level physics plots. We compare the performance and productivity of the current analysis with the two above-mentioned approaches and discuss their respective advantages and disadvantages.

Poster #26

A Software-Defined Approach for QoS Control in High-Performance Computing Storage Systems

Neda Tavakoli (Texas Tech University), Dong Dai (Texas Tech University), John Jenkins (Argonne National Laboratory), Philip Carns (Argonne National Laboratory), Robert Ross (Argonne National Laboratory), Yong Chen (Texas Tech University)

HPC storage systems become increasingly critical to scientific applications given the data-driven discovery paradigm shift. As a storage solution for large-scale HPC systems, dozens of applications share the same storage system, and will compete and can interfere with each other. Application interference can dramatically degrade the overall storage system performance. Therefore, developing a flexible and effective storage solution to assure a certain level of resources per application, i.e. the Quality-of-Service (QoS) support, is critical. One of the common solutions to achieve QoS assurance for storage systems is using provisioning technique. However, provisioning has limitations such as requiring the detailed knowledge of the expected workloads. In addition, the storage workloads are transient hence expensive to be satisfied. Due to these limitations, providing QoS storage systems through provisioning is challenging. In this research, a software-defined approach is proposed as a flexible solution to achieve QoS guarantee for storage systems.

Poster #27

Node-local IO on Aurora - CPPR

Christopher Holguin (Intel Corporation), Kalyana Chadalavada (Intel Corporation), Jeffrey Olivier (Intel Corporation), John Carrier (Intel Corporation)

Common Persistent-memory POSIX Runtime (CPPR) enables applications to leverage the fast persistent memory available on Aurora compute nodes (CNs) for I/O which enables more efficient use of compute cores. CPPR comprises 3 components - 1) async API for developers to utilize CPPR services 2) Compute node session services (CNSS) daemon 3) file movement utilities. Using CPPR, applications can 1) take advantage of the node-local (NL) filesystem backed by persistent memory hardware to reduce interaction with the global filesystem, 2) move files into and out of the NL filesystem asynchronously, and 3) create fault-tolerant checkpoint data with SCR and FTI, which utilize CPPR services.

Poster #87

A Cross-Layer Solution in Scientific Workflow System for Tackling Data Movement Challenge

Dong Dai (Texas Tech University), Robert Ross (Argonne National Laboratory), Dounia Khaldi (Stony Brook University), Yonghong Yan (Oakland University), Dorier Matthieu (Argonne National Laboratory), Neda Tavakoli (Texas Tech University), Yong Chen (Texas Tech University)

Scientific applications running in HPC environment are more complex and more data-intensive nowadays. Workflow systems are typically used to manage such complexity. Traditionally, scientific workflow systems work with parallel file systems. As such, the data need to be transferred between compute nodes and storage systems, which introduces a significant performance bottleneck on I/O operations. One promising solution to tackle this challenge is to exploit the data locality in HPC storage hierarchy. Several recent studies have been done regarding building a shared storage system, utilizing compute node resources, to serve HPC workflows with locality, such as Hercules and WOSS etc. However, in this research, we argue that providing a compute-node side storage system is not sufficient to fully exploit data locality. A cross-layer solution together with storage, compiler, and runtime is necessary. We take Swift/T, a workflow system for data-intensive applications, as a prototype platform to demonstrate our solution.

Poster #88

Attributed Consistent Hashing for Heterogeneous Storage Systems

Jiang Zhou (Texas Tech University), Wei Xie (Texas Tech University), Yong Chen (Texas Tech University)

Cloud-scale storage system is an important building block of the cloud infrastructure. It demands the flexibility to distribute data and provide high I/O performance. Consistent hashing algorithm is widely used in large-scale parallel/distributed storage systems for the decentralized design, scalability and adaptability. It can evenly distribute data among nodes but lacks efficiency in a heterogeneous environment. In this research, we propose a novel data placement algorithm, which is based on consistent hashing while making it more efficient for heterogeneous storage systems. By considering both the capacity and bandwidth attributes of nodes, our algorithm can make better use of heterogeneous devices. Our current proof-of-concept evaluations with a distributed storage system, Sheepdog, show promising results.

Poster #89

Interactive and Offline Rendering in Blender Cycles Using MPI and Intel Xeon Phi Offload

Milan Jaros (Technical University of Ostrava), Lubomir Riha (Technical University of Ostrava)

In this paper, we describe interactive and offline rendering performed on multiple nodes of an HPC cluster. Compute nodes can be equipped with Intel Xeon processors and also Intel Xeon Phi coprocessors. Rendering is utilized in Blender (open source 3D creation suite). We have modified the kernel of the Blender Cycles rendering engine and then extended its capabilities to support the HPC cluster. We call it CyclesPhi. The CyclesPhi supports hybrid MPI/OpenMP/Offload concepts. The original Blender Cycles engine has limited network rendering capabilities, which cannot be used on supercomputers. Our paper presents a basic algorithm for image rendering, decomposition tasks for application of parallel strategies, and basic collective communication routine using MPI methods. All of the presented techniques improve strong-scalability of a cluster in both the offline and also interactive rendering.

Performance

Poster #43

MuMMI_R: Analyzing and Modeling Power and Time Under Different Resilience Strategies

Xingfu Wu (Texas A&M University), Valerie Taylor (Texas A&M University), Zhiling Lan (Illinois Institute of Technology)

While reducing execution time is still a major objective for high performance computing, future systems and applications will have additional power and resilience requirements that represent a multidimensional tuning challenge. In this poster we present MuMMI_R: analyzing and modeling power and time under different resilience strategies. We use FTI (Fault Tolerance Interface) library to conduct our experiments to evaluate how using FTI with checkpoints of different levels at different frequencies impacts the power consumptions at different node components (Node, CPU, Memory, Disk and Network) and energy consumptions of the MPI memory benchmark STREAM on three different architectures IBM BG/Q, Intel Haswell and AMD Kaveri. Our experimental results provide a better understanding the tradeoffs among runtime, power and resilience.

Poster #44

Modeling and Simulation of Tape Libraries for Hierarchical Storage Systems

Jakob Luettgau (German Climate Computing Center), Julian Kunkel (German Climate Computing Center)

The variety of storage technologies results in deep storage hierarchies to be the only feasible choice to meet performance and cost requirements when handling vast amounts of data. Long-term storage systems employed by scientific users are mainly reliant on tape storage, as it remained the most cost-efficient option. Archival systems are often loosely integrated into the HPC storage infrastructure. With the rise of exascale systems and in situ analysis also burst buffers will require integration with the archive. Exploring new strategies and developing open software for tape systems is a hurdle due to the lack of affordable storage silos and availability outside of large organizations and due to increased wariness requirements when dealing with ultra-durable data. Lessening these problems by providing virtual storage silos should enable community-driven innovation, and enable site operators to add features where they see fit while being able to verify strategies before deploying on production systems.

Poster #45

Performance Engineering FUN3D at Scale with TAU Commander

John Linford (Paratools), Srinath Vadlamani (ParaTools), Sameer Shende (ParaTools), Allen Malony (ParaTools), William Jones (NASA), William Kyle Anderson (NASA), Eric Nielsen (NASA)

FUN3D is an unstructured-grid computational fluid dynamics suite widely used to support major national research and engineering efforts. FUN3D is being applied to analysis and design problems across all the major service branches at the Department of Defense. These applications span the speed range from subsonic to hypersonic flows and include both fixed- and rotary-wing configurations. This poster presents performance profiles of a high Reynolds number simulation of the flow over a wing-body-pylon-nacelle geometry on 14,400 cores of a Cray XC30 at the Navy DSRC. Profiles are gathered via TAU~Commander, which implements a new performance engineering methodology to improve user productivity. TAU~Commander highlights source code regions that limit scalability through profiling, tracing, and aggregate summary statistics with respect to computational time, memory allocation, and memory access. This analysis approach is being carefully documented to assist other DoD groups in similar performance evaluation activities.

Poster #46

Advancing Parabolic Operators in Thermodynamic MHD models: Explicit Super Time-Stepping Versus Implicit Schemes with Krylov Solvers

Ronald Caplan (Predictive Science Inc), Zoran Mikic (Predictive Science Inc), Jon Linker (Predictive Science Inc), Roberto Lionello (Predictive Science Inc)

We explore the performance/scaling of using explicit super time-stepping (STS) algorithms versus implicit schemes with Krylov solvers for integrating parabolic operators in thermodynamic MHD models of the solar corona. Specifically, we compare the second-order Runge-Kutta Legendre (RKL2) method with implicit backward Euler computed using the preconditioned conjugate gradient (PCG) solver with both a point-Jacobi and a non-overlapping domain decomposition ILU0 preconditioner. The algorithms are used to integrate anisotropic Spitzer thermal conduction and artificial kinematic viscosity at time-steps much larger than the explicit Euler limit.

A key component of the comparison is the use of a real-world simulation on large HPC systems, with special attention placed on the parallel scaling of the algorithms. It is shown that, for the specific problem and model used, the RKL2 method generally surpasses the implicit method with PCG solvers in performance and scaling, but suffers from accuracy limitations in some localized regions of the grid.

Poster #47

GPU-STREAM: Now in 2D!

Tom Deakin (University of Bristol), James Price (University of Bristol), Matt Martineau (University of Bristol), Simon McIntosh-Smith (University of Bristol)

We present a major update to the GPU-STREAM benchmark implementation, first shown at SC15. The original benchmark allowed comparison of achievable memory bandwidth performance through the STREAM kernels on OpenCL devices. GPU-STREAM v2.0 extends the benchmark to another dimension: the kernels are implemented in a wide range of popular state-of-the-art parallel programming models. This allows an intuitive comparison of performance across a diverse set of programming models and devices, investigating whether choice of model matters to performance and performance portability. In particular we investigate 7 parallel programming languages (OpenMP 4.x, OpenACC, Kokkos, RAJA, SYCL, CUDA and OpenCL) across 12 devices (6 GPUs from NVIDIA and AMD, Intel Xeon Phi (Knights Landing), 4 generations of Intel Xeon CPUs, and IBM Power 8).

Poster #48

Lightweight, Reusable Models for Dynamically Tuning Data-Dependent Code

David A. Beckingsale (Lawrence Livermore National Laboratory), Olga Pearce (Lawrence Livermore National Laboratory), Todd Gamblin (Lawrence Livermore National Laboratory)

Increasing architectural diversity has made performance portability extremely important for parallel simulation codes. Emerging C++ frameworks like RAJA and Kokkos allow developers to write kernels that can be tuned for different architectures at compile time. However, production codes use adaptive techniques like AMR, and a code's behavior depends not only on the host architecture, but also on the input problem and dynamic factors like run-time adaptation. Existing auto-tuning approaches can handle slowly evolving applications effectively, but make decisions too slowly to be executed each time control passes over a loop. We have developed Apollo, an auto-tuning extension for RAJA that uses pre-trained decision models to dynamically tune data-dependent code at runtime, with auto-generated decision models being evaluated before each kernel. We apply Apollo to applications ranging from simple hydrodynamics benchmarks to a production multi-physics code, and show that it can achieve speedups from 1.1x to 4.8x, depending on the application.

Poster #49

Simulating Batch and Application Level Scheduling Using GridSim and SimGrid

Ahmed Eleliemy (University of Basel, Switzerland), Ali Mohammed (University of Basel, Switzerland), Florina M. Ciorba (University of Basel, Switzerland)

Modern high performance computing (HPC) systems are increasing in the complexity of their design and in the levels of parallelism they offer. Studying and enhancing scheduling in HPC became very interesting for two main aspects. First, scheduling decisions are taken by different types of schedulers such as batch, application, process, and thread schedulers. Second, simulation has become an important tool to examine the design of HPC systems. Therefore, in this work, we study the simulation of different scheduling levels. We used two well-known simulation toolkits, SimGrid and GridSim, in order to support two different scheduling levels, batch and application level scheduling. Each toolkit is extended to support both levels. Moreover, three different scheduling algorithms for each level are implemented and their performance is examined through a real workload dataset. Finally, a comparison for the extension challenges of the two simulators is conducted.

Poster #50

DeepROAD: A Multifaceted Deep Learning Suite for Real-Time Optimized Autonomous Driving

Edwin L. Weill (Clemson University), Jesse Tetreault (Clemson University), Varun Praveen (Clemson University), Melissa Smith (Clemson University)

Autonomous driving is undergoing intensive study in various fields of research. The number of research topics directly related to the success of an effective autonomous vehicle is proliferating. Online assessment of surroundings, one of the most important actions for an autonomous vehicle, has had a large number of research project attempting to tackle different aspects of driving; for instance, there is research geared toward control of the vehicle while other research is focused on understanding the environment. DeepROAD intends to leverage deep learning in a manner conducive to perception and understanding of surroundings in real-time. This research employs deep learning for detection and segmentation of surroundings for decision making as well as compression, allowing for smaller networks and quicker inference times.

Poster #51

Open XDMoD Job Viewer: A Tool to Monitor Job Performance

Joseph P. White (University at Buffalo), Ryan Rathsam (University at Buffalo), Cynthia D. Cornelius (University at Buffalo), Robert L. DeLeon (University at Buffalo), Thomas R. Furlani (University at Buffalo), Steven M. Gallo (University at Buffalo), Matthew D. Jones (University at Buffalo), Abani K. Patra (University at Buffalo), Jeanette M. Sperhac (University at Buffalo), Thomas Yearke (University at Buffalo), Jeffrey T. Palmer (University at Buffalo), Nikolay Simakov (University at Buffalo), Martins Innus (University at Buffalo), Benjamin D. Plessinger (University at Buffalo)

Open XDMoD is a web-based tool intended to provide HPC stakeholders with a variety of usage and performance data. One new feature of Open XDMoD is the Job Viewer that allows users and user support personnel to view detailed job-level performance information. The user can inspect detailed time-dependent performance data (obtained from performance co-pilot or an analogous utility) to determine how efficiently their job ran and, if performance was poor, to gain insight into possible causes of the problem and how to fix it. This poster provides basic information regarding the job-level performance data that is available through the Open XDMoD Job Viewer and how to use it to assess job performance.

Poster #62

LIKWID 4: Lightweight Performance Tools

Jan Eitzinger (University of Erlangen-Nuremberg), Thomas Roehl (University of Erlangen-Nuremberg), Georg Hager (University of Erlangen-Nuremberg), Gerhard Wellein (University of Erlangen-Nuremberg)

LIKWID is a collection of command-line tools for performance-aware programmers of multicore and manycore CPUs. It follows the UNIX design philosophy of “one task, one tool.” Among its capabilities are system topology reporting, enforcement of thread-core affinity for threading, MPI, and hybrid programming models, setting clock speeds, hardware performance event counting, energy measurements, and low-level benchmarking. In this poster, we describe the feature set of the current LIKWID version and elaborate on the developments added in recent years: a new software architecture with script language APIs, a LIKWID core library for tool development, systematic validation of hardware performance event counts, and more. We aim for LIKWID to provide a one-stop solution for running and analyzing high-performance software on current multi- and many-core systems and for developing more advanced tools on top of the library interface.

Poster #63

Complexity and Accuracy Tradeoff Analysis of Parallel Application Simulation Using SST/Macro

Zhou Tong (Florida State University)

Modeling and simulating of HPC applications on contemporary supercomputing platforms is challenging due to the growing size of systems and applications. To understand the performance characteristics, parallel applications are often simulated at different levels including flow-level, packet-level, and flit-level. It is commonly believed that more detailed simulation will result in more useful information about the applications. However, it remains unclear how much more accurate simulation results finer-grained simulations can deliver at the cost of higher simulation complexity. In this study, we investigate the trade-off between modeling and simulation complexity and accuracy using the macroscale component of the Structured Simulation Toolkit (SST) and a modeling tool. We measure the performance of a set of parallel applications on three current generation supercomputers, analyze the modeling and simulation results produced by the tools, and draw conclusions from the analysis of the statistics about the trade-off between complexity and accuracy of modeling and simulation.

Poster #64

On the Path to the Holy Grail: Predicting Onset of System Failure with Log Files

Robert E. Settlage (Virginia Polytechnic Institute and State University), Michael B. Marshall (Virginia Polytechnic Institute and State University), Karthik R. Senthilvel (Virginia Polytechnic Institute and State University), Vijay K. Agarwala (Virginia Polytechnic Institute and State University), Joshua D. Akers (Virginia Polytechnic Institute and State University), Rajiv D. Bendale (Engility Corporation), Kimberly Robertson (Engility Corporation)

High performance computing environments are challenging to maintain and optimize. Indeed, in large systems, mean time to the next failure can be on the order of milliseconds. In production environments, detecting eminent failure or suboptimum performance is critical. Here, we report first steps in predicting failures – anomaly detection within system log events. Further, we demonstrate the utility of anomaly detection as a diagnostic tool for identifying misconfigurations or performance issues using stale file handles as our target message. We then propose a system to monitor system logs and alert on anomalies in real-time to improve HPC operational efficiency.

Poster #65

Understanding Ineffectiveness of the Application-Level Fault Injection

Luanzheng Guo (University of California, Merced), Jing Liang (University of California, Merced), Dong Li (University of California, Merced)

Extreme-scale applications are at a significant risk of being hit by soft errors on supercomputers, as the scale of these systems and the component density continues to increase. In order to better understand soft error vulnerabilities in those applications, the application-level fault injection is widely employed to evaluate applications. This poster reveals that the application-level fault injection has some inherent uncertainties due to the random nature of fault injection. First, the fault injection result has a strong correlation with the number of fault injection tests. What is a good number of fault injection tests is uncertain. Second, given a specific application, the fault injection result can vary as the input problem of the application varies. How to interpret the fault injection result is uncertain. Those uncertainties can make fault injection ineffective for accurately modeling application vulnerability.

Poster #66

Reducing Communication Costs in the Parallel SpMV

Amanda J. Bienz (University of Illinois), Luke Olson (University of Illinois)

Sparse matrix-vector multiplication (SpMV) is a dominant operation in many linear solvers. The large communication requirements associated with parallel SpMVs often yield inefficent methods at large scales. Data communicated across the network is much more costly than messages sent between processes which lie on the same node. Therefore, the overall cost of communication can be greatly reduced by decreasing both the number and size of inter-node messages, while increasing the amount of intra-node communication. Parallel SpMVs can be improved by gathering data among a node before sending messages across the network.

Poster #67

Parallel Performance-Energy Predictive Modeling of Browsers: Case Study of Servo

Rohit Zambre (University of California, Irvine), Lars Bergstrom (Mozilla), Laleh Beni (University of California, Irvine), Aparna Chandramowlishwaran (University of California, Irvine)

Mozilla Research is developing Servo, a parallel web browser engine, to exploit the benefits of parallelism and concurrency in the web rendering pipeline. Parallelization results in improved performance for pinterest.com but not for google.com. Occasionally, the overhead of creating, deleting, and

coordinating parallel work outweighs its benefits. In this poster, we showcase results of our models, generated through supervised learning, that capture the relationship between key web page primitives and a browser's parallel performance. Such a model allows us to predict the degree of parallelism in a web page. Additionally, we consider energy usage trade-offs for different levels of speedups in our automated labeling algorithm. This is critical for improving the browser's performance and minimizing its energy usage. Experiments on a quad-core Intel Ivy Bridge (i7-3615QM) laptop, with 535 pages on Servo's layout stage, show performance and energy improvements of up to 94.52% and 46.32% respectively.

Poster #68

DynoGraph: Benchmarking Dynamic Graph Analytics

Eric Hein (Georgia Institute of Technology), Tom Conte (Georgia Institute of Technology)

Large scale graph processing is the key to understanding complex relationships, ranging from the interaction of people on social media to the flow of data through a corporate computer network. Graph analytics present unique challenges for HPC system designers since they lack data locality and are difficult to partition into equally-sized units of work. In addressing these challenges, many researchers have opted to work with static graphs instead of the more difficult case of dynamic graphs that change rapidly during the analysis. Dynamic graphs require an entirely different memory layout, leading to degraded performance as algorithms traverse a fragmented, unsorted graph data structure. DynoGraph provides a standard for benchmarking dynamic graph analytics engines, bringing needed focus to this important class of applications.

Poster #69

Black-Box Kernel-Level Performance Modeling For GPUs

James D. Stevens (University of Illinois), Andreas Klöckner (University of Illinois)

We present a mechanism to symbolically gather performance-relevant operation counts from numerically-oriented subprograms ("kernels") expressed in the Loopy programming system, and apply these counts in a simple, linear model of kernel run time. We use a series of "performance-instructive" kernels to fit the parameters of a unified model to the performance characteristics of GPU hardware from multiple hardware generations and vendors. We evaluate predictive accuracy on a broad array of computational kernels relevant to scientific computing. In terms of geometric mean, our simple, vendor- and GPU-type-independent model achieves relative accuracy comparable to that of previously published work using hardware specific models.

Poster #70

Demonstrating the Impact of OpenMP Overheads in Multi-Physics Using a Mini App

Dylan McKinney (Lawrence Livermore National Laboratory), Ian Karlin (Lawrence Livermore National Laboratory), Riyaz Haque (Lawrence Livermore National Laboratory)

Pure MPI programming approaches only strong scale so far due to surface to volume issues. Switching to an MPI+OpenMP programming model can exploit growing concurrency while minimizing surface to volume challenges. However, for small problem sizes, the amount of work within an OpenMP region may not amortize thread overheads. KULL, a multi-physics production code, has this problem. In this poster, we demonstrate that large OpenMP overheads limit the thread scaling of KULL. We show that the Lightweight OpenMP (LOMP) compiler reduces overheads enough to allow KULL to scale further. We duplicate this issue in LULESH, a hydrodynamics proxy application. LULESH, at small problem sizes, sees up to 5x speed-up using LOMP. By demonstrating the impact overheads have on small problems, and duplicating this issue in a proxy application, we indicate to vendors that OpenMP overhead limits performance and provide them with benchmark test problems that show our OpenMP challenges.

Programming Systems

Poster #39

Devito: Fast Finite Difference Computation

Marcos de Aguiar (SENAI CIMATEC), Gerard Gorman (Imperial College London), Navjot Kukreja (SENAI CIMATEC), Michael Lange (Imperial College London), Mathias Louboutin (University of British Columbia), Felippe Zacarias (SENAI CIMATEC)

Seismic imaging, used in energy exploration, is arguably the most compute and data intensive application in the private sector. The commonly used methods involve solving the wave equations numerically using finite difference formulations. Writing optimized code for these applications involves multiple man-years of effort that need to be repeated every time a new development needs to be factored in – for every target platform.

DeVito is a new tool for performing optimized Finite Difference (FD) computation from high-level symbolic problem definitions. The application developer needs to provide a differential equation in symbolic form. DeVito performs automated code generation and Just-In-Time (JIT) compilation based on this symbolic equation to create and execute highly optimized Finite Difference kernels on multiple computer platforms. DeVito has been designed to be used as part of complex workflows involving data flows across multiple applications over different nodes of a cluster.

Poster #40

STView: An Eclipse Plug-in Tool for Visualizing Program Structures in Fortran Source Codes

Tomomi Ohichi (Kobe University), Masaaki Terai (RIKEN), Mitsuo Yokokawa (Kobe University), Kazuo Minami (RIKEN)

We developed an Eclipse plug-in tool named STView for visualizing the program structures of Fortran source codes to help improve the performance of the program on a supercomputer. To create a tree that represents program structures, STView uses an abstract syntax tree (AST) generated by Photran and filters the tree because the AST has many nontrivial nodes for tokens. While ordinary visualization tools such as profiler and refactoring tools represent a call tree that only includes the relationship among procedures, STView can visualize loops and branches in addition to a call tree. Moreover, STView can show the time-consuming parts or hotspots of a program by profiling data. We evaluated the capability of STView using 13 scientific applications collected from websites and confirmed that it robustly visualizes a tree for all applications.

Poster #41

Making a Legacy Code Auto-Tunable Without Messing It Up

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

Since computing platforms are diverging, it is becoming more important to make a legacy code “auto-tunable” so as to run it efficiently on various platforms. However, such a code is likely to be considerably complex because it is supposed to change its code structure according to several parameters. In this work, we discuss how to use auto-tuning technologies without overcomplicating the code itself. This work focuses on user-defined code transformations that can be defined separately from an application code. Then, this work demonstrates that the Xevolver framework allows users to easily define code transformation rules for transforming a particular code to its auto-tunable version. As a result, application developers can usually maintain the original version, and the code is transformed just before auto-tuning. If code transformation rules are properly defined by experts, the application developers can benefit from auto-tuning technologies without considering the complex auto-tunable code generated by the transformation.

Poster #42

Hobbes Node Virtualization Layer: System Software Infrastructure for Application Composition and Performance Isolation

Noah Evans (Sandia National Laboratories), Brian Kocoloski (University of Pittsburgh), John R. Lange (University of Pittsburgh), Kevin Pedretti (Sandia National Laboratories), Shyamali Mukherjee (Sandia National Laboratories), Ron Brightwell (Sandia National Laboratories), Patrick G. Bridges (University of New Mexico)

As computation outstrips I/O performance, moving as much of the end-to-end HPC workflow onto individual nodes is a potential solution to solving data movement bottlenecks. However this integration is difficult, most HPC simulations were designed to run in isolation therefore on-node composition potentially breaks many assumptions applications have about scheduling, resource usage, and performance. To address this problem we have developed an MPI runtime and PMI that utilize the Node Virtualization Layer provided by the Hobbes project to enable the composition of applications on-node using MPI. This approach makes it possible to use common MPI intracommunicators as a composition mechanism across heterogeneous operating systems on the same node.

Poster #71

DSL and Autotuning Tools for Code Optimization on HPC Inspired by Navigation Use Case

Jan Martinovic (IT4Innovations, VSB - Technical University of Ostrava), Katerina Slaninova (IT4Innovations, VSB - Technical University of Ostrava), Martin Golasowski (IT4Innovations, VSB - Technical University of Ostrava), Radim Cmar (Sygic), Joao M. P. Cardoso (University of Porto), Joao Bispo (University of Porto), Gianluca Palermo (Polytechnic University of Milan), Davide Gadioli (Polytechnic University of Milan), Cristina Silvano (Polytechnic University of Milan)

Improving performance and scalability of a source code for the HPC platform is a tedious and time consuming task. The code has to be executed many times under different conditions in order to observe its behavior. This optimization task is best performed automatically, but due to the heterogeneous nature of the source codes for the HPC platform, full automation is often hard to implement. We present a novel approach to this problem which uses a domain specific language (DSL) and a custom compiler toolset LARA proposed in the ANTAREX project. The DSL allows to specify strategies for code transformations, including required instrumentation for various tools and libraries. Code used in our example is part of a extensive code base of a server-side navigation software. In our poster, we present LARA strategies for scalability and performance measurements of a source code and for integration of the autotuning framework.

Poster #72

Multi-GPU Graph Analytics

Yuechao Pan (University of California, Davis), Yangzihao Wang (University of California, Davis), Yuduo Wu (University of California, Davis), Carl Yang (University of California, Davis), John D. Owens (University of California, Davis)

We present a single-node, multi-GPU programmable graph processing library that allows programmers to easily extend single-GPU graph algorithms to achieve scalable performance on large graphs with billions of edges. Our design only requires users to specify a few algorithm-dependent concerns, hiding most multi-GPU related implementation details. We analyze the theoretical and practical limits to scalability in the context of varying graph primitives and datasets. We describe several optimizations, including kernel fusion, direction optimized traversal, idempotence, and a just-enough memory allocation scheme, for better performance and smaller memory consumption. Compared to previous work, we achieve best-of-class performance across operations and datasets, including excellent strong and weak scalability on most primitives as we increase the number of GPUs in the system.

Poster #73

Tapas: An Implicitly Parallel Programming Framework For Hierarchical N-Body Algorithms

Keisuke Fukuda (Tokyo Institute of Technology), Motohiko Matsuda (RIKEN), Naoya Maruyama (RIKEN), Rio Yokota (Tokyo Institute of Technology), Kenjiro Taura (University of Tokyo), Satoshi Matsuoka (Tokyo Institute of Technology)

Tapas is our new C++ programming framework for hierarchical algorithms such as n-body, on large scale heterogeneous supercomputers. Encapsulating the algorithms' complexities in a library or a framework has been challenging due to irregular data access over distributed memory. Tapas solves this by converting the user's implicit-style parallel code into an inspector-executor style program solely by the use of C++ template metaprogramming. A prototype implementation of the Fast Multipole Method on Tapas demonstrates a comparable performance and scaling as ExaFMM, the fastest hand-tuned implementation of FMM, as well as efficient usage of hundreds of GPUs. Specifically, the serial execution achieves 15% faster performance than ExaFMM, whereas the distributed-memory strong-scaling evaluation using up to 1500 CPU cores demonstrates 64% to 81% of ExaFMM. The multi-GPU version achieves a 2.5x speedup over the CPU version when executed on 100 nodes of TSUBAME2.5 with 300 GPU.

Poster #74

Meta-Balancer: Automating Load Balancing Decisions

Harshitha Menon (University of Illinois), Kavitha Chandrasekar (University of Illinois), Laxmikant Kale (University of Illinois)

HPC applications are increasingly becoming complex and dynamic. Many applications require dynamic load balancing to achieve high performance. Different applications have different characteristics and hence need to use different load balancing strategies. There are many load balancing algorithms available. However, invocation of an unsuited load balancing strategy can lead to inefficient execution. Most commonly, the application programmer decides which load balancer to use based on some educated guess. We propose Meta-Balancer, a framework to automatically decide the best suited load balancing strategy. Meta-Balancer monitors application characteristics and based on that, it chooses an ideal load balancing algorithm to use. In order to predict the best load balancing strategy, Meta-Balancer uses a supervised random forest machine learning technique with the application characteristics as the features. Using this, we are able to achieve high prediction accuracy of 82% on the test set to demonstrate performance benefits of up to 3X.

State of the Practice

Poster #59

Data-Driven Workflows on Crays with Hybrid Scheduling: A Case Study of Celera on Magnus

Charlene Yang (Pawsey Supercomputing Center), Seyhan Yazar (University of Western Australia), George Gooden (University of Western Australia), Alex Hewitt (University of Western Australia)

This poster presents an example of data-driven workflows running on a SLURM+ALPS scheduled Cray machine at Pawsey Supercomputing Centre. This example adapts the de novo assembly software WGS-Celera, which originally has no SLURM support, to be able to dynamically request resources and distribute work through the hybrid scheduling system on Magnus, a Cray XC40 at Pawsey. The changes made to the code are uploaded to GitHub, and some insights about the adaptation are offered in this poster to benefit other users of Celera or other similarly-structured workflows around the world.

Poster #60

Modernizing a Long-Lived Production Physics Code

Charles Ferenbaugh (Los Alamos National Laboratory), Sriram Swaminarayan (Los Alamos National Laboratory), Chuck Aldrich (Los Alamos National Laboratory), Matthew Calef (Los Alamos National Laboratory), Joann Campbell (Los Alamos National Laboratory), Marcus Daniels (Los Alamos National Laboratory), Michael Hall (Los Alamos National Laboratory), Scot Halverson (Los Alamos National Laboratory), Thomas Masser (Los Alamos National Laboratory), Michael McKay (Los Alamos National Laboratory), Zachary Medin (Los Alamos National Laboratory), Ralph Menikoff (Los Alamos National Laboratory), David Nicholaeff (Los Alamos National Laboratory), Robert Robey (Los Alamos National Laboratory), Gabriel Rockefeller (Los Alamos National Laboratory), Jeremy Sauer (Los Alamos National Laboratory), John Wohlbier (CI Software Associates)

LANL's Eulerian Applications Project is working on modernizing its code base to perform at large scale on Trinity and other future platforms. To make this possible, we are untangling the complicated dependencies that have built up over the years, and dividing the functionality into smaller, self-contained packages. These packages will be well-documented, have well-defined APIs, and be unit-testable. They will allow us to refactor code in individual packages with minimal impact to other parts of the code. We will then be able to work on the optimizations for Trinity and other advanced architectures. This poster will describe the packagization strategy, our progress to date, initial optimization results, and future plans.

Poster #61

A Tool for Semi-Automatic Application-Level Checkpointing

Trung Nguyen Ba (University of Texas at Austin), Ritu Arora (University of Texas at Austin)

Computational jobs running on supercomputing resources at open-science data centers are often limited to a maximum number of compute-nodes and wall-clock time. However, many jobs need longer than the maximum allowed wall-clock time to complete. To overcome this limitation, applications can be checkpointed such that their execution state is saved before they time-out from the job-queue. Using their saved state, the applications can resume their computation from the point where they stopped in the previous run. When the checkpointing-and-restart mechanism is built within the application, it is called Application-Level Checkpointing (ALC). We are developing a tool for semi-automatic ALC of existing applications without requiring any manual reengineering of the applications. The memory footprint of the checkpoints written using our tool is small. Applications written in C/C++/MPI/OpenMP will be supported in the upcoming release of our tool, and in future, the tool will support Fortran and Python applications too.

System Software

Poster #52

Cerberus: A 3-Phase Burst Buffer Aware Batch Scheduler for HPC Systems

Jiaqi Yan (Illinois Institute of Technology), Xu Yang (Illinois Institute of Technology), Dong Jin (Illinois Institute of Technology), Zhiling Lan (Illinois Institute of Technology)

Burst buffer drastically improves the performance of computational applications by providing high perceived IO bandwidth. However, traditional batch schedulers have barely embraced the full potential of burst buffer technology in practice. In this poster, we model the execution of scientific applications that generate tens of TB of data on a supercomputer equipped with burst buffer. We characterize the lifetime of generic applications into three phases: stage-in, running, and stage-out. We develop a novel burst-buffer-aware batch scheduler Cerberus to manage resource allocation in different phases. In both stage-in and stage-out phases, Cerberus allocates the burst buffer resources to achieve the maximum data transfer throughput between burst buffer and the external storage system. In the running phase, Cerberus maximizes the predefined objectives of job scheduling.

Poster #53

Turning Privacy Constraints into Syslog Analysis Advantage

Siavash Ghiasvand (Technical University Dresden), Florina M. Ciorba (University of Basel, Switzerland), Wolfgang E. Nagel (Technical University Dresden)

The mean time between failures (MTBF) of HPC systems is rapidly reducing, and that current failure recovery mechanisms e.g., checkpoint-restart, will no longer be able to recover the systems from failures. Early failure detection is a new class of failure recovery methods that can be beneficial for HPC systems with short MTBF.

System logs (syslogs) are invaluable source of information which give us a deep insight about system behavior and make the early failure detection possible. Beside normal information, syslogs contain sensitive data which might endanger users' privacy. Even though analyzing various syslogs is necessary for creating a general failure detection/prediction method, privacy concerns discourage system administrators

to publish their syslogs. Herein, we ensure user privacy via de-identifying syslogs, and then turning the applied constraint for addressing users' privacy into an advantage for system behavior analysis. Results indicate significant reduction in required storage space and 3 times shorter processing time.

Poster #54

Accessing GPUs from Containers in HPC

Lucas Benedicic (Swiss National Supercomputing Center), Miguel Gila (Swiss National Supercomputing Center)

Thanks to the significant popularity gained lately by Docker, the HPC community has recently started exploring container technology and the potential benefits its use would bring to the users of supercomputing systems like the Cray XC30 Piz Daint of the Swiss National Supercomputing Centre (CSCS). Working in close collaboration with NERSC and an engineering team at Nvidia, CSCS is extending the Shifter framework in order to enable seamless GPU access from containers.

We explore the feasibility of using Shifter as a container framework for deploying applications using GPU devices with practically no overhead, thus achieving native execution performance. The tested examples include two GPU programming environments, namely CUDA and OpenCL.

Poster #55

The High Performance Open Community Runtime: Explorations on Asynchronous Many Task Runtime Systems

Joshua Landwehr (Pacific Northwest National Laboratory), Joshua Suetterlein (University of Delaware), Andres Marquez (Pacific Northwest National Laboratory), Joseph Manzano (Pacific Northwest National Laboratory), Kevin Barker (Pacific Northwest National Laboratory), Guang Gao (University of Delaware)

This poster presents the Performance Open Community Runtime, an asynchronous many task runtime which aims to provide a scalable, efficient, and productive platform to exercise novel runtime ideas that will exploit the massive parallel resources in today's HPC systems. This mature platform goes beyond the proof-of-concept phase by demonstrating that it can scale up to thousands of cores. Moreover, it is highly configurable, aiding in the controlled exploration and validation of novel ideas and concepts. We have proven its scalability for both strong and weak scaling experimental setups with selected kernels running on leadership clusters. We have conducted preliminary studies with different memory models and used our own characterization / introspective framework for an in-depth attribution of performance.

Poster #56

Software-level Fault Tolerant Framework for Task-based Applications

Joy Yeh (University of Bristol), Grzegorz Pawelczak (University of Bristol), James Sewart (University of Bristol), James Price (University of Bristol), Ferad Zyulkyarov (Barcelona Supercomputing Center), Leonardo Bautista-Gomez (Barcelona Supercomputing Center), Osman Unsal (Barcelona Supercomputing Center), Simon McIntosh-Smith (University of Bristol), Amaurys Avila Ibarra (University of Bristol)

Fault tolerance has been identified as one of the major challenges for exascale computing. In addition to fail-stop errors, silent data corruptions (SDCs) can perturb applications and produce incorrect results. Software-based fault tolerance mechanisms have the advantage of being capable of leveraging some of the properties of the applications to improve their reliability. In this poster, we present a fault tolerance framework that implements multiple resiliency schemes to cope with both fail-stop errors and data corruption. Our techniques are tested with two real scientific applications: BUDE, a molecular docking engine, and TeaLeaf, a heat conduction code. Using this framework we have successfully detected and recovered from real data corruptions. We have also performed error injection experiments, which clearly demonstrated the efficacy of our framework.

Poster #57

Power-Aware Heterogeneous Computing Through CPU-GPU Hybridization

Kyle Siehl (Washington State University Vancouver), Xinghui Zhao (Washington State University Vancouver)

Graphic Processing Units (GPUs) have recently been widely used in general purpose computing, aiming for improving the performance of applications. However, this performance gain often comes with higher power consumption. In this paper, we present Archon, a framework for power-aware CPU-GPU hybridization. Specifically, Archon takes user's programs as input, automatically distribute the workload between CPU and GPU, and dynamically tunes the distribution ratio at runtime for an energy-efficient execution. Experiments have been carried out using a matrix multiplication application, and the results show that Archon can provide considerable energy savings, comparing to the CPU-only and GPU-only executions. These energy savings are achieved without extra efforts from the programmers.

Poster #58

Pin-Pointing Node Failures in HPC Systems

Anwesha Das (North Carolina State University), Frank Mueller (North Carolina State University), Paul Hargrove (Lawrence Berkeley National Laboratory), Eric Roman (Lawrence Berkeley National Laboratory)

Automated fault prediction and diagnosis in HPC systems needs to be efficient for better system resilience. With increasing scalability required for exascale, accurate fault prediction aiding in quick remedy is hard. With changing supercomputer architectures, distilling fault data from the noisy raw logs requires substantial efforts. Predicting node failures in such voluminous system logs is challenging. To this end, we investigate an interesting way to pin-point node failures in such supercomputing systems. Our study on Cray system data with automated machine learning tools suggests that specific patterns of event messages on node unavailability can be indicator to node failures. This data extraction coupled with system and job data correlation helps in devising a methodology to predict node failures and their location over a specific time frame. This work aims to enable broader applicability for a generic fault prediction framework.

Scientific Visualization & Data Analytics Showcase

SC16's Visualization and Data Analytics Showcase Program provides a forum for the year's most instrumental movies in HPC. This year, there will be both a live display throughout the conference so that attendees can experience and enjoy the latest in science and engineering HPC results expressed through state-of-the-art visualization technologies, and a session at SC16 dedicated to the best of the submissions.

Selected entries will be displayed live in a museum/art gallery format in the South foyer on the Lower Concourse of the Convention Center. Six finalists will compete for the Best Visualization Award, with each finalist presenting his or her movie in a 15-minute presentation on Wednesday, November 16.

Movies are judged based on how their movie illuminates science, by the quality of the movie, and for innovations in the process used for creating the movie.

Tuesday, November 15

Scientific Visualization & Data Analytics/ Poster Reception

5:15pm-7pm
Room: South Foyer

The Scientific Visualization & Data Analytics gallery will be part of SC16's poster reception and is an opportunity for attendees to interact with the authors. The reception is open to all attendees. Complimentary refreshments and appetizers are available.

Scientific Visualization & Data Analytics Exhibition

Tuesday, November 15 – Thursday, November 17

8:30am-5:00pm
Room: South Foyer

The Scientific Visualization & Data Analytics movies will be displayed live in a museum/art gallery format.

High Resolution Visualization and Simulation of High-Pressure Primary Atomization Using a Complex Diesel Injector

Authors: *Simon M. Su (US Army Research Laboratory), Luis G. Bravo (US Army Research Laboratory), Richard C. Angelini (US Army Research Laboratory), Michael M. Stephens (US Army Engineer Research and Development Center), Richard I. Walters (Lockheed Martin), Miguel A. Valenciano (Lockheed Martin), Vu H. Tran (Lockheed Martin), David G. Longmire (Lockheed Martin)*

In combustion modeling research, one of the main challenges is the inaccessibility of the optically dense region and the steep resolution demands for observation. Scientific data visualization enables visualization of the data at the necessary scale and resolution. In our visualization laboratory, we used ParaSAGE and a large high resolution tiled display system to enable high resolution data visualization. We also extended ParaView to run on zSpace virtual reality system to enable 3D immersive and interactive exploration of the scientific data. Animations created using the time varying data also provided additional insight. Using all three visualization techniques, we are able to have a better understand of the combustion model.

Visualization and Analysis of Threats from Asteroid Ocean Impacts

Authors: *John Patchett (Los Alamos National Laboratory), Francesca Samsel (University of Texas at Austin), Karen Tsai (University of Texas at Austin), David Rogers (Los Alamos National Laboratory), Greg Abram (University of Texas at Austin), Terece Turton (University of Texas at Austin), Boonthanome Nouanesengsy (Los Alamos National Laboratory), Galen Gisler (Los Alamos National Laboratory)*

An asteroid colliding with earth can have grave consequences. An impact in the ocean has complex effects as the kinetic energy of the asteroid is transferred to the water, so scientists at Los Alamos National Laboratories are using XRAGE simulations on HPC systems to understand the behavior of these ocean impacts. By running ensembles of large scale 3D simulations, scientists study the range of simulation variables such as angle of impact, asteroid mass, and whether or not the asteroid explodes in the atmosphere or impacts the ocean. These variables help scientists understand the magnitude of impacts such as dispersement of water into the atmosphere - where it can impact the global climate, and tsunami creation that can place population centers at risk.

Data Mining Tornadogenesis Precursors

Authors: *Greg Foss (University of Texas at Austin), Amy McGovern (University of Oklahoma), Corey Potvin (National Oceanic and Atmospheric Administration), Greg Abram (University of Texas at Austin), Anne Bowen (University of Texas at Austin), Neena Hulkoti (University of Texas at Austin), Arnav Kaul (University of Texas at Austin), Nick Suey (University of Texas at Austin)*

We investigate using 3D visualization techniques and an interactive user interface with data mining for identifying tornadogenesis precursors in supercell thunderstorm simulations. We've found results will assist defining storm features (objects) extracted and input to the data mining: ensuring automatically extracted objects match visually identified ones. This video shows samples of hook echoes, BWERs, updrafts, cold pools, helicity with regions of strong vorticity, and vertical pressure perturbation gradients.

Video: https://dl.dropboxusercontent.com/u/94130018/xsede16_foss-tacc_StormSimsDataMining-Tornadogenesis.mov

Interpretation of Simultaneous Mechanical-Electrical-Thermal Failure in a Lithium-Ion Battery Module

Authors: *Chao Zhang (National Renewable Energy Laboratory), Shriram Santhanagopalan (National Renewable Energy Laboratory), Mark J. Stock (National Renewable Energy Laboratory), Nicholas Brunhart-Lupo (National Renewable Energy Laboratory), Kenny Gruchalla (National Renewable Energy Laboratory)*

Lithium-ion batteries are currently the state-of- the-art power sources for electric vehicles, and their safety behavior when subjected to abuse, such as a mechanical impact, is of critical concern. A coupled mechanical-electrical-thermal model for simulating the behavior of a lithium-ion battery under a mechanical crush has been developed. We present a series of production-quality visualizations to illustrate the complex mechanical and electrical interactions in this model.

Visualization of the Eastern Renewable Generation Integration Study

Authors: *Kenny Gruchalla (National Renewable Energy Laboratory), Joshua Novacheck (National Renewable Energy Laboratory), Aaron Bloom (National Renewable Energy Laboratory)*

The Eastern Renewable Generation Integration Study (ERGIS), explores the operational impacts of the wide spread adoption of wind and solar photovoltaics (PV) resources in the U.S. Eastern Interconnection and Québec Interconnection (collectively, EI). In order to understand some of the economic and reliability challenges of managing hundreds of gigawatts of wind and PV generation, we developed state of the art tools, data, and models for simulating power system operations using hourly unit commitment and 5-minute economic dispatch over an entire year. Using NREL's high-performance computing capabilities and new methodologies to model operations, we found that the EI could balance the variability and uncertainty of high penetrations of wind and PV at a 5-minute level under a variety of conditions. A large-scale display and a combination of multiple coordinated views and small multiples were used to visually analyze the four large highly multivariate scenarios with high spatial and temporal resolutions.

Nanocarbon Synthesis by High-Temperature Oxidation of Nanoparticles

Authors: *Joseph A. Insley (Argonne National Laboratory), Ying Li (Argonne National Laboratory), Rajiv K. Kalia (University of Southern California), Aiichiro Nakano (University of Southern California), Pankaj Rajak (University of Southern California), Ken-ichi Nomura (University of Southern California), Priya Vashishta (University of Southern California)*

High-temperature oxidation of silicon-carbide nanoparticles (nSiC) underlies a wide range of technologies, such as high-power electronic switches, thermal protection of space vehicles and self- healing ceramic nanocomposites. Multimillion-atom reactive molecular dynamics simulations computed on Argonne National Laboratory's Mira IBM Blue Gene/Q supercomputer and validated by ab initio quantum simulations predict unexpected condensation of nanocarbon with fractal geometry during high-temperature oxidation of nSiC. In this work, we present the visualization of the growth of graphene-like carbon clusters produced on the surface of an oxidizing nSiC, highlighting their connective structure.

SC16 Salt Lake City, Utah | hpc matters.

Student Programs

Saturday, November 12

Building Professional Networks Workshop
5:30pm-7:30pm
Room: Exhibit Hall 5

Raquell Holmes (improvscience)

Do you want to create more functional relationships with your colleagues? Would you like to find more comfort and freedom in the way you talk about what you do? Led by Dr. Raquell Holmes, this workshop uses improvisational theater in which participants create collaborative conversations and interviews. Through these playful exercises students develop the professional communication skills necessary for building networks among themselves and others. Students will come away with greater flexibility in talking about their interests and engaging confidently with their colleagues. The 2-hour session is open to conference attendees who are registered as students or are part of the student programs offered with the conference.

Sunday, November 13

Diversity and Inclusion: Views from the Field
8:30am-10am
Room: 260

Raquell Holmes (improvscience), Ritu Arora (University of Texas at Austin), George Thiruvathukal (Loyola University Chicago)

The year is 2016, and for the United States the nation is focused, polarized on issues of gender and race. Is this the same as discussions on diversity and inclusion? How does this national focus play out in our day to day? What does diversity and inclusion have to do in our day to day work as professionals within HPC related fields? In these talks on Diversity and Inclusion, we hear from professionals within HPC. They are colleagues- supervisors, researchers -who pay attention to, have learned from and poignantly speak of their experience of gender and race in their work in HPC fields.

Panel: Coming Out on Behalf of Diversity
10:30am-12pm
Room: 260

Raquell Holmes (improvscience), Jamika Burge (Pennsylvania State University), Jeanine Cook (Sandia National Laboratories), Rebecca Hartman-Baker (Lawrence Berkeley National Laboratory)

What does it mean to come out on behalf of diversity, when there is already a focus on issues of gender? Together, our panelists and audience will explore passions, questions, and challenges that come with taking a stand for diversity. What's difficult about coming out as an advocate for full inclusion across race, class, ability, and gender lines? What can and are we doing to build the field we want to be in? Bring your desire for an inclusive computing enterprise and come out on behalf of diversity.

Education/Career Keynote and Pitch-It Workshop
1:30pm-3pm
Room: 260

Education/Career Keynote
Wen-Mei Hwu (University of Illinois)

Pitch-It Workshop: Learn to Sell Yourself and Your Research
Interactive session for students to put together an elevator pitch on themselves and / or their research. Students will develop pitches under guidance for the first 30 minutes followed by 15 minutes of paired practicing.

Education & Career Panels
3:30pm-5pm
Room: 260

Mentoring: Learning from the Pros
Three brief talks on mentoring.

Life on the Other Side: Early Career Panel
Panel session discussion careers and opportunities in HPC including academia, government, and industry.

Monday, November 14

Finding Your Way: The Possibilities, Pitfalls, and Practice of HPC Research

8:30am-10am
Room: 260

In a series of short talks, mid-career and senior researchers in the HPC community will describe some of the challenges they faced and what they did to chart their career paths. Topics will include research issues for those who are not yet engaged in HPC research; ways to get your desired research project funded; differences in career paths for researchers between industry, government labs, and academia; and "best practices" for initiating, conducting, and publishing on HPC research.

HOWTO: Getting Started In The HPC Research Community

10:30am-12pm
Room: 260

This panel discussion aims to help students understand what they can do right now to better prepare themselves to contribute to the HPC research community. Questions from the audience are encouraged. The panel will discuss topics including: choices facing undergraduate students (graduate school or not? Ph.D. or not?); classes vs research; choices facing graduate students (academia or industry? R1 or R2 academia? Government laboratory vs commercial research lab?); the importance of starting networking early; and other things our panelists wish their younger selves had known.

Navigating SC16, Getting Involved with SC17... and Beyond

1:30pm-3pm
Room: 260

This session will provide an overview of the different programs and activities that will take place during SC16. Members of the SC16 planning committee will describe areas including: the Technical Program, HPC Matters, Visualization Showcase, Emerging Technologies, among others. In addition, presenters will cover topics that will help attendees start preparing for SC17. Are you a student interested in submitting to the Tech Program? Do you want to learn more about opportunities for students at SC? Come to this session and learn more SC. Members of the SC17 planning committee will be available to answer all your SC related questions.

Experiencing HPC for Undergraduates Opening Reception

2pm-3pm
Room: 151-B

This session is a "meet and greet" for students in the program and the HPC for Undergraduates Chair and Vice Chair.

Experiencing HPC for Undergraduates Orientation

3pm-5pm
Room: 151-B

Alan Sussman (University of Maryland), Erik Saule (University of North Carolina, Charlotte)

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.

Mentor-Protégé

3:30pm-5pm
Room: Exhibit Hall 5

The Mentor-Protege program connects newcomers at SC with experienced conference attendees. This event serves as an opportunity for mentor-protege pairs to meet each other and others. Matches are made well before the conference. Attendance is by ticket only.

Student Cluster Competition Kickoff

7:30pm-9pm
Room: Exhibit Hall 1

SC16 is excited to hold another nail-biting Student Cluster Competition (SCC), now in its tenth year! Held as part of SC16's Students@SC, the Student Cluster Competition is designed to introduce the next generation of students to the high-performance computing (HPC) community and showcase student expertise in a friendly yet spirited competition.

The Student Cluster Competition provides a multi-disciplinary experience integrated within the HPC community's biggest gathering, the Supercomputing Conference. The competition is a microcosm of a modern HPC center that teaches and inspires students to pursue careers in the field. Student teams design and build small clusters, with hardware and software vendor partners, learn the designated scientific applications and apply optimization techniques for their chosen architectures. Come see the student teams compete at booth #4723.

This year the SCC is excited to announce a new activity. Students will be replicating results from the award winning paper from SC15: "A parallel connectivity algorithm for de Bruijn graphs in metagenomic applications" and it's accompanying application, ParConnect. This is done as a part of the Three Rs initiative that aims to increase the level of repeatability, reproducibility and replicability of results at the SC conference as well as other ACM conferences.

The Student Cluster Competition would like to thank Raytheon, SAIC, Micron, Schlumberger and Allinea for their generous support.

Tuesday, November 15

Student Cluster Competition
10am-6pm
Room: Exhibit Hall 1

Experiencing HPC for Undergraduates: Introduction to HPC Research
10:30am-12pm
Room: 250-D

John Shalf (Lawrence Berkeley National Laboratory), Mary Hall (University of Utah), Martin Schulz (Lawrence Livermore National Laboratory), Edmond Chow (Georgia Institute of Technology)

A panel of leading practitioners in HPC will introduce the various aspects of HPC, including architecture, applications, programming models and tools.

Wednesday, November 16

Student/Postdoc Job Fair
10am-3pm
Room: 251-ABDE

The SC16 Students@SC Program will again host the Student/Postdoc Job Fair. The Student/Postdoc Job Fair will be open to all students and postdocs attending SC16, giving them an opportunity to meet with potential employers. Previous SC job fair employers include research labs, academic institutions, recruiting agencies and private industry who met with students to discuss graduate fellowships, internships, summer jobs, co-op opportunities, graduate school assistant positions and permanent employment. We expect this year's event to be equally successful. Resumes submitted by students will be provided to organizations participating in the job fair. Students should also bring copies of their resumes to the event.

Student Cluster Competition
10am-5pm
Room: Exhibit Hall 1

Experiencing HPC for Undergraduates: Graduate Student Perspective
10:30am-12pm
Room: 250-D

Ke Wen (Columbia University), Qingrui Lu (Virginia Polytechnic Institute and State University), Shaden Smith (University of Minnesota), Kanika Sood (University of Oregon), Markus Hohnerbach (RWTH Aachen University), Pierre Matri (Technical University of Madrid), Chao Li (North Carolina State University)

This session will be held as a panel discussion. Current graduate students, some of whom are candidates for the Best Student Paper Award in the Technical Papers program at SC16, 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 papers.

Thursday, November 17

Experiencing HPC for Undergraduates: Careers in HPC
10:30am-12pm
Room: 250-D

Jennifer Schopf (Indiana University), Kate Keahey (Argonne National Laboratory), Jerome Vienne (University of Texas at Austin), Tim Mattson (Intel Corporation)

This panel will feature a number of distinguished members of the HPC research community discussing their varied career paths. The panel includes representatives from industry, government labs and universities. The session will include ample time for questions from the audience.

SC16 Salt Lake City, Utah | hpc matters.

Tutorials

Sunday, November 13

Building the Modern Research Data Portal Using the Globus Platform

8:30am-5pm
Room: 250-C

Presenters: *Steve Tuecke (University of Chicago), Vas Vasiliadis (University of Chicago), Rachana Ananthakrishnan (University of Chicago), Stuart Martin (University of Chicago)*

New Globus REST APIs, combined with high-speed networks and Science DMZs, create a research data platform on which developers can create entirely new classes of scientific applications, portals, and gateways. Globus is an established service that is widely used for managing research data on XSEDE, DOE, and campus computing resources, and it continues to evolve with the addition of data publication capabilities, and enhancement of the core data transfer and sharing functions. Over the past year, we have added new identity and access management functionality that will simplify access to Globus using campus logins, and facilitate the integration of Globus, XSEDE, and other research cyberinfrastructure services into web and mobile applications. In this tutorial we will describe and demonstrate how web applications can leverage Globus and Science DMZs to provide a broad range of researchers with access to advanced data management capabilities using existing organizational credentials. A combination of presentation and hands-on exercises will result in attendees building a simple, yet fully functional, web application that can be used in their own work.

Efficient Parallel Debugging for MPI, Threads, and Beyond

8:30am-5pm
Room: 250-D

Presenters: *Matthias S. Mueller (RWTH Aachen University), Joachim Protze (RWTH Aachen University), Bronis R. de Supinski (Lawrence Livermore National Laboratory), David Lecomber (Allinea), Mark O'Connor (Allinea), Ganesh Gopalakrishnan (University of Utah), Tobias Hilbrich (Technical University Dresden)*

Parallel software enables modern simulations on high performance computing systems. Defects—or commonly bugs—in these simulations can have dramatic consequences on matters of utmost importance. This is especially true if defects remain unnoticed and silently corrupt results. The parallelism that enables these simulations gives rise to a wide range of parallelism related defects. Thus developers must both cope with mastering complex HPC architectures and the pitfalls of parallel programming.

We invite attendees to a tutorial that addresses the efficient removal of software defects. The tutorial provides systematic debugging techniques that are tailored to defects that revolve around parallelism. We present leading edge tools that aid developers in pinpointing, understanding, and removing defects in MPI, OpenMP, MPI-OpenMP, and further parallel programming paradigms. The tutorial tools include the widely used parallel debugger Allinea DDT, the data race detector Intel Inspector XE, the highly-scalable stack analysis tool STAT, and the MPI correctness tool MUST. Hands-on examples—make sure to bring a laptop computer—will guide attendees in the use of these tools. We conclude the tutorial with a discussion and provide pointers for debugging with paradigms such as CUDA or on architectures such as Xeon Phi.

High Performance Python for Scientific Computing

8:30am-5pm
Room: 355-E

Presenters: *Stan Seibert (Continuum Analytics), Michael Klemm (Intel Corporation), WIlliam Scullin (Argonne National Laboratory), Kevin O'Leary (Intel Corporation)*

Today's HPC mainstream languages and programming models consist mainly of Fortran, C/C++, OpenMP, and MPI. Simple-to-use scripting languages such as Python are being used beyond pre and post processing of data. The learning curve for Python is relatively low, and it is a natural first step for those scientist/domain specialists to dip their proverbial toes into computational science. The burgeoning use of high performance compute platforms as applied to big data analytics and machine learning raises the performance requirements as Python is well used in these areas. Scaling up performance of Python applications is challenging without having to rewrite the application in a native language.

In this tutorial participants will learn about the latest developments and tools for high performance Python for scikit-learn, NumPy/SciPy/Pandas, mpi4py, numba, etc. The audience will learn how to apply low overhead profiling tools to analyze mixed C/C++ and Python applications to detect performance bottlenecks in the code and to pinpoint hotspots as the target for performance tuning. Join us and learn the best known methods, tools, and libraries to get the best performance from your Python application.

HPC Acquisition and Commissioning

8:30am-12pm
Room: 355-F

Presenters: *Andrew Jones (Numerical Algorithms Group), Terry Hewitt (WTH Associates Ltd), Owen G. M. Thomas (Red Oak Consulting)*

This tutorial, part of the SC16 State of the Practice, will guide attendees through the process of purchasing and deploying an HPC system. It will cover the whole process from engaging with stakeholders in securing funding, requirements capture, market survey, specification of the tender/request for proposal documents, engaging with suppliers, evaluating proposals, and managing the installation.

Attendees will learn how to specify what they want, yet enable the suppliers to provide innovative solutions beyond their specification both in technology and in the price; how to demonstrate to stakeholders that the solution selected is best value for money; and the common risks, pitfalls and mitigation strategies essential to achieve an on-time and on-quality installation process.

The tutorial has 3 parts: the procurement process including RFP; bid evaluation, benchmarks, clarification processes; contracting, project management, commissioning, acceptance testing.

The presenters have been involved in numerous major HPC procurements since 1990, as service managers, bidders to funding agencies, as customers and as impartial expert advisors. The presenters are from the UK, but the tutorial will be applicable to HPC procurements anywhere. The tutorial is based on experience across a diverse set of real world cases in various countries, in private and public sectors.

Kokkos: Enabling Manycore Performance Portability for C++ Applications

8:30am-5pm
Room: 250-E

Presenters: *H. Carter Edwards (Sandia National Laboratories), Jeff Amelang (Google), Christian Trott (Sandia National Laboratories)*

The Kokkos library enables applications and domain specific libraries to implement intra-node thread scalable algorithms that are performance portable across diverse manycore architectures. Kokkos uses C++ template meta-programming, as opposed to compiler extensions or source-to-source translators, to map user code onto architecture-targeted mechanisms such as OpenMP, Pthreads, and CUDA. Kokkos' execution mapping inserts users' parallel code bodies into well-defined parallel patterns and then uses an architecture-appropriate scheduling to execute the computation. Kokkos' data mapping implements polymorphic layout multidimensional arrays that are allocated in architecture-abstracted memory spaces with a layout (e.g., row-major, column-major, tiled) appropriate for that architecture. By integrating execution and data mapping into a single programming model Kokkos eliminates the contemporary array-of-structures versus structure-of-arrays dilemma from user code. Kokkos' programming model consists of the following extensible abstractions: execution spaces where computations execute, execution policies for scheduling computations, parallel patterns, memory spaces where data is allocated, array layouts mapping multi-indices onto memory, and data access intent traits to portably map data accesses to architecture-specific mechanisms such as GPU texture cache. Tutorial participants will learn Kokkos' programming model through lectures, hands-on exercises, and presented examples.

Node-Level Performance Engineering

8:30am-5pm
Room: 255-C

Presenters: *Georg Hager (University of Erlangen-Nuremberg), Gerhard Wellein (University of Erlangen-Nuremberg)*

The advent of multi- and many-core chips has led to a further opening of the gap between peak and application performance for many scientific codes. This trend is accelerating as we move from petascale to exascale. Paradoxically, bad node-level performance helps to "efficiently" scale to massive parallelism, but at the price of increased overall time to solution. If the user cares about time to solution on any scale, optimal performance on the node level is often the key factor. We convey the architectural features of current processor chips, multiprocessor nodes, and accelerators, as far as they are relevant for the practitioner. Peculiarities like SIMD vectorization, shared vs. separate caches, bandwidth bottlenecks, and ccNUMA characteristics are introduced, and the influence of system topology and affinity on the performance of typical parallel programming constructs is demonstrated. Performance engineering and performance patterns are suggested as powerful tools that help the user understand the bottlenecks at hand and to assess the impact of possible code optimizations. A cornerstone of these concepts is the roofline model, which is described in detail, including useful case studies, limits of its applicability, and possible refinements.

Parallel Computing 101

8:30am-5pm
Room: 255-D

Presenters: *Quentin F. Stout (University of Michigan), Christiane Jablonowski (University of Michigan)*

This tutorial provides a comprehensive overview of parallel computing, emphasizing those aspects most relevant to the user. It is suitable for new users, managers, students, and anyone seeking an overview of parallel computing. It discusses software and hardware/software interaction with an emphasis on standards, portability, and systems that are widely available.

The tutorial surveys basic parallel computing concepts, using examples selected from multiple engineering, scientific, and data analysis problems. These examples illustrate using MPI on distributed memory systems; OpenMP on shared memory systems; MPI+OpenMP on hybrid systems; CUDA, OpenACC, OpenCL on GPUs; and Hadoop and Spark on big data. It discusses numerous parallelization and load balancing approaches, and software engineering and performance improvement aspects, including the use of state-of-the-art tools.

The tutorial helps attendees make intelligent decisions by covering the primary options that are available, explaining how they are used and what they are most suitable for. Extensive pointers to web-based resources are provided to facilitate follow-up studies.

Parallel I/O In Practice

8:30am-5pm
Room: 355-B

Presenters: *Robert Latham (Argonne National Laboratory), Robert Ross (Argonne National Laboratory), Brent Welch (Google)*

I/O on HPC systems is a black art. This tutorial sheds light on the state-of-the-art in parallel I/O and provides the knowledge necessary for attendees to best leverage I/O resources available to them. We cover the entire I/O software stack including storage and parallel file systems at the lowest layer, the role of burst buffers (nvram), intermediate layers (such as MPI-IO), and high-level I/O libraries (such as HDF5). We emphasize ways to use these interfaces that result in high performance and tools for generating insight into these stacks. Benchmarks on real systems are used throughout to show real-world results.

In the first third of the tutorial we cover the fundamentals of parallel I/O. We discuss storage technologies, both present and near-future. Our parallel file systems material covers general concepts and gives examples from Lustre, GPFS, PanFS, HDFS, Ceph, and Cloud Storage. Our second third takes a more application-oriented focus. We examine the upper library layers of the I/O stack, covering MPI-IO, parallel netCDF, and HDF5. We discuss interface features, show code examples, and describe how application calls translate into PFS operations. Finally we discuss tools for capturing and understanding I/O behavior.

Parallel Programming in Modern Fortran

8:30am-5pm
Room: 250-B

Presenters: *Karla Morris (Sandia National Laboratories), Damian Rouson (Sourcery Institute), Salvatore Filippone (Cranfield University), Fernanda S. Foertter (Oak Ridge National Laboratory)*

Fortran remains widely used in HPC, but most users describe their programming skills as self-taught, and most continue to use older versions of the language. The advent of an open-source compiler and runtime library supporting many of the parallel programming features of Fortran 2015 opens up the opportunity for widespread evaluation of these new features

by the HPC community. This tutorial teaches single-program, multiple-data (SPMD) programming with Fortran 2008/2015 coarrays. We also introduce Fortran's loop concurrency and pure procedure features and demonstrate their use in asynchronous expression evaluation for partial differential equation (PDE) solvers. We incorporate other language features, including object-oriented (OO) programming, when they support our chief aim of teaching parallel programming. In particular, we demonstrate OO design patterns that enable hybrid CPU/GPU calculations on sparse matrices in the Parallel Sparse Basic Linear Algebra Subroutines (PSBLAS) library. We will also cover the parallel collective subroutines that have been proposed for Fortran 2015 and are supported by the GNU and Cray compilers. Attendees will use the GCC 5.3 Fortran compiler and the OpenCoarrays library[1] to compile parallel executables inside virtual machines. Those interested in GPU computing will have access to Oak Ridge's Cray supercomputer Titan.

[1] http://www.opencoarrays.org

Power-Aware High Performance Computing: Challenges and Opportunities for Application and System Developers

8:30am-12pm
Room: 255-E

Presenters: *Dieter Kranzlmueller (Ludwig Maximilian University of Munich), David Lowenthal (University of Arizona), Barry Rountree (Lawrence Livermore National Laboratory), Martin Schulz (Lawrence Livermore National Laboratory)*

Power and energy consumption are critical design factors for next generation large-scale HPC systems. The costs for energy are shifting the budgets from investment to operating costs, and more and more often the size of systems will be determined by its power needs.

As a consequence, the US Department of Energy (DOE) has set the ambitious limit of 20MW for their first exascale system, and many other funding agencies around the world have expressed similar goals. Yet, with today's HPC architectures and systems, this is still far out of reach: the goal will only be achievable through a complex set of mechanisms at all levels of hardware and software, including buildings and infrastructure; all of these aspects will additionally and directly impact the application developer. On future HPC systems, running a code efficiently (as opposed to purely with high performance) will be a major requirement for every user.

In this tutorial, we will discuss the challenges caused by power and energy constraints, review available approaches in hardware and software, highlight impacts on HPC center and infrastructure design and operations, and ultimately show how this change in paradigm from "cycle awareness" to "power awareness" will impact application development.

Productive Programming in Chapel: A Computation-Driven Introduction

8:30am-5pm
Room: 250-A

Presenters: *Bradford Chamberlain (Cray Inc.), Michael Ferguson (Cray Inc.), Greg Titus (Cray Inc.), Ben Albrecht (Cray Inc.), Sung-Eun Choi (Cray Inc.), Engin Kayraklioglu (George Washington University)*

Chapel (http://chapel.cray.com) is an emerging open-source language whose goal is to vastly improve the programmability of parallel systems while also enhancing generality and portability compared to conventional techniques. Considered by many users to be the most promising of next-generation parallel languages, Chapel is seeing growing levels of interest not only among HPC programmers, but also in the data analytics, academic, and mainstream communities. Chapel's design and implementation are portable and open-source, supporting a wide spectrum of platforms from desktops to commodity clusters, the cloud, and supercomputers from Cray and other vendors.

This tutorial will provide an in-depth introduction to Chapel's features using a computation-driven approach: rather than simply lecturing on individual language features, we will motivate each Chapel concept by illustrating its use in computations taken from motivating benchmarks and proxy applications. A pair of hands-on segments will let participants write, compile, and execute parallel Chapel programs, either by logging into Cray accounts that we will provide, or directly on their laptops (gcc + make should be pre-installed). We'll end the tutorial by providing an overview of the project's status and activities, and by soliciting feedback from participants with the goal of improving Chapel to their benefit.

Programming Intel's 2nd Generation Xeon Phi (Knights Landing)

8:30am-5pm
Room: 255-B

Presenters: *Si Liu (University of Texas at Austin), Victor Eijkhout (University of Texas at Austin), Antonio Gómez-Iglesias (University of Texas at Austin), Jerome Vienne (University of Texas at Austin), Carlos Rosales (University of Texas at Austin), Dmitry Prohorov (Intel Corporation)*

Intel's next generation Xeon Phi, Knights Landing (KNL), brings many changes from the first generation, Knights Corner (KNC). The new processor supports self-hosted nodes, connects cores via mesh topology rather than a ring, and uses a new memory technology, MCDRAM. It is based on Intel's x86 technology with wide vector units and multiple hardware threads per core. KNL introduces new challenges for the programmer because of the multiple configuration options for MCDRAM memory.

This tutorial is designed for experienced programmers familiar with MPI and OpenMP. We'll review the KNL architecture, and discuss the differences between KNC and KNL. We'll discuss the impact of the different MCDRAM memory configurations (flat, cache, hybrid) and the different modes of cluster configuration (all-to-all, quadrant, sub-NUMA quadrant). Recommendations regarding MPI task layout when using KNL with the Intel OmniPath fabric will be provided.

As in past tutorials, we will focus on the use of reports and directives to improve vectorization and the implementation of proper memory access. We will also showcase new Intel VTune Amplifier XE capabilities that allow for in-depth memory access analysis and hybrid code profiling. Hands-on exercises will be executed on the KNL-upgraded Stampede system at the Texas Advanced Computing Center (TACC).

Scalable HPC Visualization and Data Analysis Using VisIt

8:30am-12pm
Room: 255-F

Presenters: *Cyrus Harrison (Lawrence Livermore National Laboratory), David Pugmire (Oak Ridge National Laboratory), Hank Childs (University of Oregon), Robert Sisneros (University of Illinois), Jens Henrik Göbbert (Forschungszentrum Juelich), Matthieu Dorier (Argonne National Laboratory)*

Visualization and data analysis are essential components of the scientific discovery process. Scientists and analysts running HPC simulations rely on visualization and analysis tools for data exploration, quantitative analysis, visual debugging, and communication of results. This half-day tutorial will provide SC16 attendees with a practical introduction to mesh-based HPC visualization and analysis using VisIt, an open source parallel scientific visualization and data analysis platform. We will provide an introduction to HPC visualization practices and couple this with hands-on experience using VisIt.

This tutorial includes: 1) An introduction to visualization techniques for mesh-based simulations 2) A guided tour of VisIt 3) Hands-on demonstrations visualizing a CFD Blood Flow simulation 4) Special Topic: Beyond post-hoc visualization and analysis using in situ processing

This tutorial builds on the past success of VisIt tutorials, updated for SC16 with new content focused on use cases where in situ processing has merits over traditional post-hoc file based analysis. Attendees will gain practical knowledge and recipes to help them effectively use VisIt to analyze data from their own simulations.

Solving and Sharing the Puzzle: Modeling and Simulation of Computer Architectures with SST and OCCAM

8:30am-12pm
Room: 250-F

Presenters: *Bruce Childers (University of Pittsburgh), Luis Oliveira (University of Pittsburgh), Arun Rodrigues (Sandia National Laboratories)*

Modeling and simulation are critical in the development and understanding of HPC hardware and design. Often, simulation is the foremost method of affordable exploration to understand the intricacies of design space trade-offs, novel architectures, emerging components, and interconnect infrastructures to generate vital statistics. There is an urgency for interoperability, consistency, and communication between simulation tools and developers. To facilitate this use of node-level simulations, a standard communication system between models is needed.

The Structural Simulation Toolkit (SST), developed by Sandia National Laboratories, is a scalable, open-source, parallel discrete-event framework designed to act as a unifying tool for hardware simulation and component analysis. The University of Pittsburgh's Open Curation for Computer Architecture Modeling (OCCAM) program allows for reproducible, customizable, and simplified integration of simulation tools for easing community development, deployment, and reproducibility.

This tutorial introduces the fundamental components and key features of SST. Examples are presented that will build on

introductory concepts and illustrate the creation of new simulation components. The OCCAM platform will be introduced, enabling participants to explore SST with an example set of simulations. Participants will also be invited to explore a series of instructor-led customization and testing techniques, with a focus on model expansion and analysis.

Big Data Meets HPC: Exploiting HPC Technologies for Accelerating Apache Hadoop, Spark, and Memcached

1:30pm-5pm
Room: 355-C

Presenters: *Dhabaleswar K. (DK) Panda (Ohio State University), Xiaoyi Lu (Ohio State University)*

Apache Hadoop and Spark are gaining prominence in handling Big Data and analytics. Similarly, Memcached in Web-2.0 environment is becoming important for large-scale query processing. Recent studies have shown that default Hadoop, Spark, and Memcached can not efficiently leverage the features of modern high-performance computing clusters, like Remote Direct Memory Access (RDMA) enabled high-performance interconnects, high-throughput and large-capacity parallel storage systems (e.g. Lustre). These middleware are traditionally written with sockets, and do not deliver best performance on modern networks. In this tutorial, we will provide an in-depth overview of the architecture of Hadoop components (HDFS, MapReduce, RPC, HBase, etc.), Spark, and Memcached. We will examine the challenges in re-designing networking and I/O components of these middleware with modern interconnects, protocols (such as InfiniBand, iWARP, RoCE, and RSocket) with RDMA and storage architectures. Using the publicly available software packages in the High-Performance Big Data project (HiBD, http://hibd.cse.ohio-state.edu), we will provide case studies of the new designs for several Hadoop/Spark/Memcached components and their associated benefits. Through these case studies, we will also examine the interplay between high-performance interconnects, storage (HDD, NVM, and SSD), and multi-core platforms to achieve the best solutions for these components and Big Data applications on modern HPC clusters.

Essential HPC Finance Practice: Total Cost of Ownership (TCO), Internal Funding, and Cost-Recovery Models

1:30pm-5pm
Room: 355-F

Presenters: *Andrew Jones (Numerical Algorithms Group), Owen G. M. Thomas (Red Oak Consulting)*

The tutorial provides an impartial, practical, non-sales focused guide to the financial aspects of HPC facilities and service. It presents a rare and invaluable opportunity for HPC managers, practitioners, and stakeholders to learn more about calculating and using TCO models; along with the pros and cons of different internal cost recovery and funding models.

Well-managed TCO, return on investment and cost recovery models can be hugely beneficial to HPC managers and operators by demonstrating the value of HPC to the organization, driving the continuation and growth of HPC investment. They can also help uncover practical improvements to deliver better services to users.

Attendees will benefit from exploration of the main issues, pros and cons of differing approaches, practical tips, hard-earned experience and potential pitfalls. After the tutorial, attendees will be in a stronger position to calculate and use TCO within their organizations, and to design and use internal cost-recovery models.

The tutorial, part of the SC16 State of the Practice theme, is based on experience across a diverse set of real world cases in various countries, in both private and public sectors, with projects of all sizes and shapes.

Insightful Automatic Performance Modeling

1:30pm-5pm
Room: 250-F

Presenters: *Alexandru Calotoiu (Technical University Darmstadt), Felix Wolf (Technical University Darmstadt), Torsten Hoefler (ETH Zurich), Martin Schulz (Lawrence Livermore National Laboratory)*

Many applications suffer from latent performance limitations that may cause them to consume too many resources under certain conditions. Examples include the input-dependent growth of the execution time. Solving this problem requires the ability to properly model the performance of a program to understand its optimization potential in different scenarios. In this tutorial, we will present a method to automatically

generate such models for individual parts of a program from a small set of measurements. We will further introduce a tool that implements our method and teach how to use it in practice. The learning objective of this tutorial is to familiarize the attendees with the ideas behind our modeling approach and to enable them to repeat experiments at home.

Managing HPC Software Complexity with Spack

1:30pm-5pm
Room: 255-F

Presenters: *Todd Gamblin (Lawrence Livermore National Laboratory), Massimiliano Culpo (EPFL), Gregory Becker (Lawrence Livermore National Laboratory), Matthew P. LeGendre (Lawrence Livermore National Laboratory), Gregory L. Lee (Lawrence Livermore National Laboratory), Elizabeth Fischer (Columbia University), Benedikt Hegner (CERN)*

HPC software is becoming increasingly complex. The largest applications require over 70 dependency libraries, compilers, build options, Python, and MPI implementations to achieve good performance. However, the space of possible build configurations is combinatorial, and existing package management tools do not handle these complexities well. Because of this, most HPC software is built by hand. Developers waste countless hours porting and rebuilding software instead of producing new scientific results.

This tutorial focuses on Spack, our open-source tool for HPC package management. Spack uses concise package recipes written in Python to automate builds with arbitrary combinations of compilers, MPI versions, and dependency libraries. With Spack, users can rapidly install software without knowing how to build it; developers can efficiently manage automatic builds of tens or hundreds of dependency libraries; and HPC center staff can deploy many versions of software for thousands of users. We provide a thorough introduction to Spack's capabilities: from simple installation, to creating custom packages, to advanced multi-user deployment and Python installations. We also provide detailed use cases based on our experiences at LLNL, EPFL, NASA, CERN, and other HPC sites. Most sessions involve hands-on demonstrations, so attendees should bring a laptop computer.

Monday, November 14

Advanced MPI Programming

8:30am-5pm
Room: 355-E

Presenters: *Pavan Balaji (Argonne National Laboratory), William Gropp (University of Illinois), 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. At the same time, the MPI standard itself is evolving to address the needs and challenges of future extreme-scale platforms as well as applications. This tutorial will cover several advanced features of MPI, including new MPI-3 features, that can help users program modern systems effectively. Using code examples based on scenarios found in real applications, we will cover several topics including efficient ways of doing 2D and 3D stencil computation, derived datatypes, one- sided communication, hybrid (MPI + shared memory) programming, topologies and topology mapping, and neighborhood and nonblocking collectives. Attendees will leave the tutorial with an understanding of how to use these advanced features of MPI and guidelines on how they might perform on different platforms and architectures.

Advanced OpenMP: Performance and 4.5 Features

8:30am-5pm
Room: 255-C

Presenters: *Christian Terboven (RWTH Aachen University), Ruud van der Pas (Oracle Corporation), Eric Stotzer (Texas Instruments), Bronis R. de Supinski (Lawrence Livermore National Laboratory), Michael Klemm (Intel Corporation)*

With the increasing prevalence of multicore processors, shared-memory programming models are essential. OpenMP is a popular, portable, widely supported, and easy-to-use shared-memory model. Developers usually find OpenMP easy to learn. However, they are often disappointed with the performance and scalability of the resulting code. This disappointment stems not from shortcomings of OpenMP, but rather with the lack of depth with which it is employed. Our "Advanced OpenMP Programming" tutorial addresses this critical need by exploring the implications of possible OpenMP parallelization strategies, both in terms of correctness and performance.

While we quickly review the basics of OpenMP programming, we assume attendees understand basic parallelization concepts. We focus on performance aspects, such as data

and thread locality on NUMA architectures, false sharing, and exploitation of vector units. We discuss language features in-depth, with emphasis on advanced features like tasking or cancellation. We close with the presentation of the directives for attached compute accelerators. Throughout all topics we present the recent additions of OpenMP 4.5 and extensions that have been subsequently adopted by the OpenMP Language Committee.

Application Porting and Optimization on GPU-Accelerated POWER Architectures

8:30am-5pm
Room: 255-F

Presenters: *Oscar Hernandez (Oak Ridge National Laboratory), Archana Ravindar (IBM), Jiri Kraus (NVIDIA Corporation), Andreas Herten (Forschungszentrum Juelich), Costas Bekas (IBM), Bronson Messer (Oak Ridge National Laboratory)*

The POWER processor has re-emerged as a technology for supercomputer architectures. One major reason is the tight integration of processor and GPU accelerator through the new NVLink technology. Two major sites in the US, ORNL and LLNL, have already decided to have their pre-exascale systems being based on this new architecture. This tutorial will give an opportunity to obtain in-depth knowledge and experience with GPU-accelerated POWER nodes. It focuses on porting applications to a single node and covers the topics architecture, compilers, performance analysis, and multi-GPU programming. The tutorial will include an overview of the new NVLink based node architectures, lectures on first-hand experience in porting to this architecture, and will conclude with exercises using tools to focus on performance.

Container Computing for HPC and Scientific Workflows

8:30am-12pm
Room: 355-B

Presenters: *Luiz Carlos Irber Jr (University of California, Davis), Shane Canon (Lawrence Berkeley National Laboratory), Lisa Gerhardt (Lawrence Berkeley National Laboratory)*

Container computing is revolutionizing the way applications are developed and delivered. It offers opportunities that never existed before for significantly improving efficiency of scientific workflows and easily moving these workflows from the laptop to the supercomputer. Tools like Docker and Shifter enable a new paradigm for scientific and technical computing. However, to fully unlock its potential, users and administrators need to understand how to utilize these new approaches. This tutorial will introduce attendees to the basics of creating container images, explain best practices, and cover more advanced topics such as creating images to be run on HPC platforms using Shifter. The tutorial will also explain how research scientists can utilize container-based computing to accelerate their research and how these tools can boost the impact of their research by enabling better reproducibility and sharing of their scientific process without compromising security.

Debugging and Performance Analysis on Native and Offload HPC Architectures

8:30am-5pm
Room: 250-A

Presenters: *Sergi Siso (Hartree Centre), Damian Alvarez (Juelich Supercomputing Center), Sandra Wienke (RWTH Aachen University), Nikolay Piskun (Rogue Wave Software), Chris Gottbrath (NVIDIA Corporation), Woo-Sun Yang (National Energy Research Scientific Computing Center)*

Scientists and engineers from different domains are turning to simulation in high-end supercomputers due to their ever increasing capabilities, thus helping them to understand physical phenomena. However, that means that in addition to understanding the complex phenomena studied, scientist and engineers have to understand the intricacies of modern HPC systems. This tutorial is aimed at those professionals that need to understand which tools can help them to efficiently use these new systems.

To fulfill the tutorial targets, we focus on Knights Landing (KNL) as well as GPU-based hybrid architectures. We cover debugging techniques on these architectures using the TotalView parallel debugger. We also discuss profiling tools for GPUs and coprocessors. Here the training focuses on commercial (VTune, NVProf) and open source (Extrae/Paraver) tools. The programming environment used in the tutorial is a combination of MPI, OpenACC, and OpenMP 4.0.

The tutorial is highly practical. Demos on different platforms, like Power8+GPU, will be shown during the lecture to include this new platform and also to help the audience retain the ideas explained in the slides. The tutorial will include hands-on exercises with a set of tasks for each covered platform. Participants need SSH and NX & VNC capable laptops.

Fault-Tolerance for HPC: Theory and Practice

8:30am-5pm
Room: 250-F

Presenters: *George Bosilca (University of Tennessee), Aurelien Bouteiller (University of Tennessee), Thomas Herault (University of Tennessee), Yves Robert (ENS Lyon)*

Resilience is a critical issue for large-scale platforms, and this tutorial provides a comprehensive survey of fault-tolerant techniques for HPC, with a fair balance between practice and theory.

This tutorial is organized along four main topics: (i) An overview of failure types (software/hardware, transient/fail-stop), and typical probability distributions (Exponential, Weibull, Log-Normal); (ii) General-purpose techniques, which include several checkpoint and rollback recovery protocols, replication, prediction, and silent error detection; (iii) Application-specific techniques, such as ABFT for grid-based algorithms or fixed-point convergence for iterative applications, user-level checkpointing in memory; and (iv) Practical deployment of fault tolerant techniques with User Level Fault Mitigation (a proposed MPI standard extension).

Relevant examples based on ubiquitous computational solver routines will be protected with a mix of checkpoint-restart and advanced recovery techniques in a hands-on session.

The tutorial is open to all SC16 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. However, basic knowledge of MPI will be helpful for the hands-on session.

Hands-On Practical Hybrid Parallel Application Performance Engineering

8:30am-5pm
Room: 250-B

Presenters: *Christian Feld (Juelich Supercomputing Center), Markus Geimer (Juelich Supercomputing Center), Sameer Shende (University of Oregon), Bert Wesarg (Technical University Dresden), Brian J. N. Wylie (Juelich Supercomputing Center)*

This tutorial presents state-of-the-art performance tools for leading-edge HPC systems founded on the community-developed Score-P instrumentation and measurement infrastructure, demonstrating how they can be used for performance engineering of effective scientific applications based on standard MPI, OpenMP, hybrid combination of both, and increasingly common usage of accelerators. Parallel performance tools from the Virtual Institute – High Productivity Supercomputing (VI-HPS) are introduced and featured in 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 HPC Linux [http://www.hpclinux.org] OVA image containing all of the necessary tools (running within a virtual machine), participants will conduct exercises on the Stampede system at TACC where remote access to Intel Xeon and Intel Xeon Phi accelerator-based nodes will be provided for the hands-on sessions. This will help to prepare participants to locate and diagnose performance bottlenecks in their own parallel programs. Go to http://www.vi-hps.org/training/other/SC16.html for directions to prepare for the hands-on exercises and a detailed agenda.

Harnessing the Power of FPGAs with Altera's SDK for OpenCL

8:30am-12pm
Room: 355-F

Presenters: *Andrew Ling (Intel Corporation), Byron Sinclair (Intel Corporation)*

In this tutorial, you will learn why Field Programmable Gate Arrays have become so popular in HPC and have been outpacing the overall semiconductor industry in terms of adoption and growth. Furthermore, we will cover architectural features, such as hardened floating point DSPs and on-chip memory, of FPGAs that make them well suited to many applications traditionally run on multicore CPUs and GPUs at much higher performance per watt. In addition, we will introduce programming FPGAs using Altera's SDK for OpenCL and how specific OpenCL coding techniques can lead to efficient circuits implemented on the FPGA. Finally, we will go over several case studies where FPGAs have shown very competitive performance when programmed using OpenCL, including convolutional neural nets, FFTs, and astronomy de-dispersion algorithms.

How to Analyze the Performance of Parallel Codes 101

8:30am-5pm
Room: 255-E

Presenters: *Martin Schulz (Lawrence Livermore National Laboratory), Jim Galarowicz (Krell Institute), Don Maghrak (Krell Institute), Jennifer Green (Los Alamos National Laboratory), David Montoya (Los Alamos National Laboratory), Doug Pase (Sandia National Laboratories)*

Performance analysis is an essential step in the development of HPC codes. It will even gain in importance with the rising complexity of machines and applications that we are seeing today. Many tools exist to help with this analysis, but the user is too often left alone with interpreting the results.

We will provide a practical roadmap for the performance analysis of HPC codes and will provide users step-by-step advice on how to detect and optimize common performance problems, covering both on-node performance and communication optimization as well as issues on threaded and accelerator-based architectures. Throughout this tutorial, we will show live demos using Open|SpeedShop, a comprehensive and easy-to-use tool set. Additionally, at the end of each section we will provide hands-on exercises for attendees to try out the new techniques learned. All techniques will, however, apply broadly to any tool, and we will point out alternative tools where useful.

InfiniBand and High-Speed Ethernet for Dummies

8:30am-12pm
Room: 355-C

Presenters: *Dhabaleswar K. (DK) Panda (Ohio State University), Hari Subramoni (Ohio State University)*

InfiniBand (IB) and High-Speed Ethernet (HSE) technologies are generating a lot of excitement toward building next generation High-End Computing (HEC) systems including clusters, datacenters, file systems, storage, cloud computing, and Big Data (Hadoop, Spark, HBase and Memcached) environments. RDMA over Converged Enhanced Ethernet (RoCE) technology is also emerging. This tutorial will provide an overview of these emerging technologies, their offered architectural features, their current market standing, and their suitability for designing HEC systems. It will start with a brief overview of IB and HSE. In-depth overview of the architectural features of IB and HSE (including iWARP and RoCE), their similarities and differences, and the associated protocols will be presented. An overview of the emerging Omni-Path architecture will be provided. Next, an overview of the OpenFabrics stack which encapsulates IB, HSE and RoCE (v1/v2) in a unified manner will be presented. An overview of libfabrics stack will also be provided. Hardware/software solutions and the market trends behind IB, HSE and RoCE will be highlighted. Finally, sample performance numbers of these technologies and protocols for different environments will be presented.

Large Scale Visualization with ParaView

8:30am-5pm
Room: 250-D

Presenters: *Kenneth Moreland (Sandia National Laboratories), W. Alan Scott (Sandia National Laboratories), David E. DeMarle (Kitware Inc), Joe Insley (Argonne National Laboratory), Jonathan Woodring (Los Alamos National Laboratory), John Patchett (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.

Linear Algebra Libraries for High-Performance Computing: Scientific Computing with Multicore and Accelerators

8:30am-5pm
Room: 250-E

Presenters: *Jack Dongarra (University of Tennessee), Jakub Kurzak (University of Tennessee), Michael Heroux (Sandia National Laboratories), James Demmel (University of California, Berkeley)*

Today, a desktop with a multicore processor and a GPU accelerator can already provide a TeraFlop/s of performance, while the performance of the high-end systems, based on multicores and accelerators, is already measured in tens of PetaFlop/s. This tremendous computational power can only be fully utilized with the appropriate software infrastructure, both at the low end (desktop, server) and at the high end (supercomputer installation). Most often a major part of the

computational effort in scientific and engineering computing goes in solving linear algebra subproblems. After providing a historical overview of legacy software packages, the tutorial surveys the current state-of-the-art numerical libraries for solving problems in linear algebra, both dense and sparse. (D) PLASMA, MAGMA, PETSc, Trilinos, SuperLU, Hypre, and other software packages are discussed. The tutorial also highlights recent advances in algorithms that minimize communication, i.e. data motion, which is much more expensive than arithmetic.

Programming Your GPU with OpenMP: A Hands-On Introduction

8:30am-12pm
Room: 255-B

Presenters: *Simon McIntosh-Smith (University of Bristol), James Reinders (Independent), Timothy Mattson (Intel Corporation)*

OpenMP 1.0 was released in 1997. Back then the concern was symmetric multiprocessors. Over time hardware evolved with more complex memory hierarchies (NUMA computers) and distributed memory systems which forced us to understand how OpenMP mixes with MPI.

Current trends bring co-processors such as GPUs into the fold. A modern platform is often best viewed as a heterogeneous system with CPU cores, GPU cores, and over time other specialized accelerators. OpenMP has responded by adding a set of directives that map code and data onto a device. We refer to this family of directives as the target directives.

In this tutorial, we will explore these directives as they apply to programming GPUs. We assume attendees know the fundamentals of OpenMP so we can use our time to deeply understand the target directives and their use in complex applications. This will be a hands-on tutorial. Students will use their own laptops (with Windows, Linux, or OS X). We will have access to remote servers with GPUs, but the best option is for attendees to load an OpenMP compiler onto their laptops before the tutorial. Information about OpenMP compilers is available at www.openmp.org.

Secure Coding Practices and Automated Assessment Tools

8:30am-12pm
Room: 250-C

Presenters: *Barton Miller (University of Wisconsin), Elisa Heymann (Autonomous University of Barcelona)*

Security is crucial to the software that we develop and use. With the incredible growth of both web and cloud services, security is becoming even more critical. Securing your network is not enough! Every service that you deploy is a window into your data center from the outside world and is a window that could be exploited by an attacker.

This tutorial is relevant to anyone wanting to learn about minimizing security flaws in the software they develop or manage. We share our experiences gained from performing vulnerability assessments of critical middleware. You will learn skills critical for software developers and analysts concerned with security. Software assurance tools – tools that scan the source or binary code of a program to find weaknesses – are the first line of defense in assessing the security of a software project. These tools can catch flaws in a program that affect both the correctness and safety of the code. This tutorial is also relevant to anyone wanting to learn how to use these automated assessment tools to minimize security flaws in the software they develop or manage.

Data Center Design and Planning for HPC Folks

1:30pm-5pm
Room: 355-F

Presenters: *Sharan Kalwani (DataSwing), Michael Thomas (ESD Global), Bernie Woytek (Gensler Inc)*

HPC usage continues to rise and is no longer the domain of the special centers. The mainstreaming of HPC will place demands on the environment needed to house it. During an SC15 meeting of HPC admins, the second most lamented item (survey available) was that HPC folks are often without warning, tasked with planning data center needs. This tutorial addresses that, since no other SCxy tutorial is available to impart the basics of planning for a data center. We may know how to run codes efficiently but have no clue as to how to plan for or manage the support infrastructure. Data center design is critical, as otherwise expensive HPC resources can be hobbled by lack of proper data center facilities.

In this tutorial, the authors propose to train users on how to begin the planning process, take stock of the current environmental, space, cooling, electricity, security, operations, and

maintenance as well as foresee potential problems, point out methods and resources. The tutorial goals are to provide the attendees with skills to: do specific HPC data center design and tailored planning; equip them with knowledge about various choices/capabilities; maximize resources, and allow for growth without tears.

InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage

1:30pm-5pm
Room: 355-C

Presenters: *Dhabaleswar K. (DK) Panda (Ohio State University), Hari Subramoni (Ohio State University)*

As InfiniBand (IB) and High-Speed Ethernet (HSE) technologies mature, they are being used to design and deploy various High-End Computing (HEC) systems: HPC clusters with accelerators (GPGPUs and MIC) supporting MPI, Storage and Parallel File Systems, Cloud Computing systems with SR-IOV 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, portability, reliability and network congestion. 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. Advanced hardware and software features of IB, HSE and RoCE and their capabilities to address these challenges will be emphasized. Next, we will focus on Open Fabrics RDMA and Libfabrics programming, and network management infrastructure and tools to effectively use these systems. A common set of challenges being faced while designing these systems will be presented. Finally, 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.

Programming Irregular Applications with OpenMP: A Hands-On Introduction

1:30pm-5pm
Room: 250-C

Presenters: *Tim Mattson (Intel Corporation), Alice Koniges (Lawrence Berkeley National Laboratory), Clay Breshears (Intel Corporation)*

OpenMP 1.0 (released in 1997) focused on loop-level parallelism. This works well for “regular applications” where the parallelism maps nicely onto blocks of loop iterations. For problems based on traversing linked lists, recursive algorithms, or other “irregular applications”, a loop based approach is difficult to use.

To handle these irregular problems, tasks were added to OpenMP 3.0 (released in 2008). A task is an instance of executable code and its data environment. They weren’t just grafted onto the older, loop-oriented OpenMP core. The entire specification had to be carefully restructured to make sure threads, tasks, worksharing constructs, and the rest of OpenMP all worked together in a single system.

In this tutorial, we assume students already know the fundamentals of OpenMP. We will explore tasks through a sequence of exercises designed to help people learn how to use tasks in application programs. 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.

Testing of HPC Scientific Software

1:30pm-5pm
Room: 255-B

Presenters: *Anshu Dubey (Argonne National Laboratory), Alicia Klinvex (Sandia National Laboratories)*

Testing at various granularities has recently acquired an urgency in high-performance scientific computing communities because of the need for refactoring caused by changing platform architectures. Projects that need to refactor are often lacking in the necessary expertise and resources to acquire such expertise. Testing is also critical for producing credible results and for code maintenance. The IDEAS (www.ideas-productivity.org) scientific software productivity project aims toward increasing software productivity and sustainabil-

ity with participants from many projects that define the state of practice in software engineering in the HPC community. We offer a tutorial that distills the combined knowledge of IDEAS team members in the area of scientific software testing. The tutorial will be useful to all projects that recognize the importance of testing in general and will provide tremendous help to projects in need of refactoring their software in particular.

Vectorization Strategies for Intel's 2nd Generation Intel Xeon® Phi™ Architecture Named Knights Landing

1:30pm-5pm
Room: 355-B

Presenters: *Ashish Jha (Intel Corporation), Vitali Morozov (Argonne National Laboratory), Jack Deslippe (Lawrence Berkeley National Laboratory)*

Intel® has recently introduced the second generation Xeon® Phi™ architecture code named Knights Landing (KNL). KNL incorporates new features, such as self-booting capabilities and a new memory architecture. KNL also leverages the new AVX-512 ISA and has two floating point units that can execute two AVX-512 instructions in the same cycle.

Exploiting vector capabilities is a key to effective programming on KNL. The Intel programming environment can be best harnessed by picking the right level of abstraction and obtaining vectorization with minimal effort. In some cases, the compiler needs only to get a pragma hint. While, for experienced users, the environment provides a set of intrinsics which is abstraction of low-level assembly instructions, the full control over the code can be gained by assembly programming.

This tutorial will give attendees an understanding of vectorization. We will show why codes often fail to auto-vectorize, which features of the compiler they can leverage to help with auto-vectorization, and how to apply AVX-512 vector intrinsics in writing their own low-level vectorized code which cannot be easily expressed in high-level language.

Workshops

Sunday, November 13

3rd International Workshop on HPC User Support Tools (HUST-16)

9am-12:30pm
Room: 155-E

Ralph C. Bording (Pawsey Supercomputing Center), Todd Gamblin (Lawrence Livermore National Laboratory), Vera Hansper (CSC - IT Center for Science)

Supercomputing centers exist to drive scientific discovery by supporting researchers in computational science fields. To make users more productive, in the complex HPC environment, HPC centers employ user support teams. These teams serve many roles, from setting up accounts, to consulting on math libraries and code optimization, to managing HPC software stacks. Often, support teams struggle to adequately support scientists. HPC environments are extremely complex, and combined with the complexity of multi-user installations, exotic hardware, and maintaining research software, supporting HPC users can be extremely demanding.

With the third HUST workshop, we will continue to provide a necessary forum for system administrators, user support team members, tool developers, policy makers and end users. We will provide fora to discuss support issues, and we will provide a publication venue for current support developments. Best practices, user support tools, and any ideas to streamline user support at supercomputing centers are in scope.

Computational Approaches for Cancer

9am-5:30pm
Room: 251-B

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

As the size, number, variety, and complexity of cancer datasets have grown in recent years, new computational challenges and opportunities have emerged within the cancer research and clinical application areas. The workshop focuses on bringing together interested individuals ranging from clinicians, mathematicians, data scientists, computational scientists, hardware experts, engineers, developers, leaders, and others with an interest in advancing the use of computation at all levels to better understand, diagnose, treat, and prevent cancer. With an interdisciplinary focus, the workshop provides opportunities for participants to learn about how computation is employed across multiple areas including imaging, genomics, analytics, modeling, pathology, and drug discovery with a focus on impacting cancer. The forward focus of the workshop looks at challenges and opportunities for large scale HPC, including potential for exascale applications involving cancer.

ExaMPI16

9am-5:30pm
Room: 155-A

Ryan E. Grant (Sandia National Laboratories), Patrick G. Bridges (University of New Mexico), Anthony Skjellum (Auburn University), Purushotham V. Bangalore (University of Alabama), Ronald B. Brightwell (Sandia National Laboratories), Pavan Balaji (Argonne National Laboratory), Masamichi Takagi (RIKEN)

The aim of workshop is to bring together researchers and developers to present and discuss innovative algorithms and concepts in the Message Passing programming model and to create a forum for open and potentially controversial discussions on the future of MPI in the Exascale era. Possible workshop topics include innovative algorithms for collective operations, extensions to MPI, including datacentric models, scheduling/routing to avoid network congestion, "fault-tolerant" communication, interoperability of MPI and PGAS models, integration of task-parallel models in MPI, and use of MPI in large scale simulations.

HPGDMP'16: First International Workshop on High Performance Graph Data Management and Processing

9am-12:30pm
Room: 251-C

Toyotaro Suzumura (IBM), Dario Garcia-Gasulla (Barcelona Supercomputing Center), Miyuru Dayarathna (WSO2 Inc), Julian Shun (University of California, Berkeley)

Applications that need to manage and process large scale graph data have become prominent in recent times. Social network analysis, semantic web, bioinformatics, and cheminformatics are some examples of application domains that deal with large graphs containing billions of vertices and edges. Graph processing has attracted significant attention from the high performance computing community due to the complexities involved with the processing and storage of large graphs. Data management systems such as in-memory, distributed graph databases have been introduced for storing and managing large graphs. Graph processing frameworks and libraries have been developed to simplify high performance large scale graph processing. Furthermore, large scale distributed-memory compute clusters, single shared-memory high performance computers, and heterogeneous hardware containing GPUs and FPGAs have been used for carrying out large scale graph data processing tasks. These efforts have been bolstered by graph related benchmarking initiatives such as the Graph 500 and Green Graph 500 benchmarks. Despite these significant research efforts, there still exist significant issues and technical gaps which need to be solved in the area of high performance graph data management and processing. The High Performance Graph Data Management and Processing 2016 (HPGDMP16) workshop aims to provide a unified platform for discussing the latest state-of-the-art efforts conducted to address research issues related to high performance large graph management and processing.

IA^3 2016 - 6th Workshop on Irregular Applications: Architectures and Algorithms

9am-5:30pm
Room: 251-D

Antonino Tumeo (Pacific Northwest National Laboratory), John Feo (Pacific Northwest National Laboratory), Oreste Villa (NVIDIA Corporation)

Due to the heterogeneous data sets they process, data intensive applications employ a diverse set of methods and data structures. Consequently, they abound with irregular memory accesses, control flows, and communication patterns. Current supercomputing systems are organized around components optimized for data locality and bulk synchronous computations. Managing any form of irregularity on them demands substantial effort and often leads to poor performance.

Holistic solutions to address these challenges can emerge only 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. Strong collaborative efforts among researchers with different expertise, including domain experts and end users, could lead to significant breakthroughs. This workshop brings together scientists with these different backgrounds to discuss methods and technologies for efficiently supporting irregular applications on current and future architectures.

Innovating the Network for Data Intensive Science (INDIS 2016)

9am-5:30pm
Room: 155-F

Cees de Laat (University of Amsterdam), Malathi Veeraraghavan (University of Virginia), Brian Tierney (Energy Sciences Network), Paola Grosso (University of Amsterdam), Sylvia Kuijpers (SURFnet)

Wide area networks are now an integral and essential part of the data-driven supercomputing ecosystem connecting information sources, processing, simulation, visualization, and user communities together. Networks for data-intensive science have more extreme requirements than general-purpose networks. These requirements closely impact the design of processor interconnects in supercomputers/cluster computers, but they also impact campus networks, regional networks and national backbone networks. The ability to move large datasets in and out of supercomputing centers are an integral and essential part of the data-driven supercomputing ecosystem. Further, networks are required to connect research instruments such as photon sources and large visualization displays. This workshop brings together the network researchers and innovators to present challenges and novel ideas that stretch SCinet and network research even further. This workshop presents papers that propose new and novel techniques regarding capacity and functionality of networks, its control, and its architecture, possibly to be demonstrated at current and future supercomputing conferences.

Every year SCinet develops and implements the network for the SC conference. This network is state-of-the-art, connects many demonstrators of big science data processing infrastructures at the highest line speeds and newest technologies available, and demonstrates novel functionality. The show

floor network connects to many laboratories and universities worldwide using high-bandwidth connections. The NRE part of the workshop will showcase a representative subset of these demonstrations and host a panel on future innovations and involvement of SCinet.

ISAV 2016: In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization

9am-5:30pm
Room: 251-A

Wes Bethel (Lawrence Berkeley National Laboratory), Patrick O'Leary (Kitware Inc), Gunther Weber (Lawrence Berkeley National Laboratory), Nicola Ferrier (Argonne National Laboratory), Matthew Wolf (Georgia Institute of Technology), Earl Duque (Intelligent Light)

The considerable interest in the HPC community regarding in situ analysis and visualization is due to several factors. First is an I/O cost savings, where data is analyzed/visualized while being generated, without first storing to a file system. Second is the potential for increased accuracy, where fine temporal sampling of transient analysis might expose some complex behavior missed in coarse temporal sampling. Third is the ability to use all available resources, CPU's and accelerators, in the computation of analysis products.

The workshop brings together researchers, developers, and practitioners from industry, academia, and government laboratories developing, applying, and deploying in situ methods in extreme-scale, high performance computing. The goal is to present research findings, lessons learned, and insights related to developing and applying in situ methods and infrastructure across a range of science and engineering applications in HPC environments; to discuss topics like opportunities presented by new architectures, existing infrastructure needs, requirements and gaps, and experiences to foster and enable in situ analysis and visualization; to serve as a "center of gravity" for researchers, practitioners, and users/consumers of in situ methods and infrastructure in the HPC space.

Seventh Annual Workshop for the Energy Efficient HPC Working Group (EE HPC WG)

9am-5:30pm
Room: 155-C

Michael Patterson (Intel Corporation), Torsten Wilde (Leibniz Supercomputing Centre), Natalie Bates (Lawrence Berkeley National Laboratory), Stephen Poole (US Department of Defense), Daniel Hackenberg (Technical University Dresden), Thomas Durbin (University of Illinois), Steve Martin (Cray Inc.), Dale Sartor (Lawrence Berkeley National Laboratory), Anna Maria Bailey (Lawrence Livermore National Laboratory), Ramkumar Nagappan (Intel Corporation), Nicolas Dube (Hewlett Packard Enterprise), Andres Marquez (Pacific Northwest National Laboratory), David Grant (Oak Ridge National Laboratory), Josip Loncaric (Los Alamos National Laboratory), James H. Laros (Sandia National Laboratories), David J. Martinez (Sandia National Laboratories), Susan Coghlan (Argonne National Laboratory), James H. Rogers (Oak Ridge National Laboratory), Toshio Endo (Tokyo Institute of Technology), Fumiyoshi Shoji (RIKEN), Francois Robin (CEA-DAM), Caleb Phillips (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 600 members from ~20 different countries. This workshop is unique in that it provides a forum for sharing best practices of major supercomputing centers from around the world. Discussion and audience participation is encouraged. There are presentations, panels and discussions instead of papers and posters. Presenters are mostly from major governmental and academic supercomputing centers. The panels encourage discussion around more controversial topics and include panelists from supercomputing centers, academic institutions as well as the vendor community. This year's topics will include: future proofing your data-center design, energy efficiency economics, open-interfaces and frameworks for power and energy management, optical interconnects and photonics, controlling power and energy through-out the software stack. This year's keynote speaker is Thomas Schulthess, Director of the Swiss National Supercomputing Centre.

The 7th International Workshop on Performance Modeling, Benchmarking and Simulation of HPC Systems (PMBS16)

9am-5:30pm
Room: 155-B

Simon D. Hammond (Sandia National Laboratories), Stephen A. Jarvis (University of Warwick), Steven A. Wright (University of Warwick)

The PMBS16 workshop is concerned with the comparison of high-performance computing systems through performance modeling, benchmarking, or through the use of tools such as simulators. We are particularly interested in research which reports the ability to measure and make tradeoffs in software/hardware co-design to improve sustained application performance. We are also keen to capture the assessment of future systems, for example, through work that ensures continued application scalability through peta- and exa-scale systems.

The aim of this workshop is to bring together researchers, from industry and academia, concerned with the qualitative and quantitative evaluation and modeling of high-performance computing systems. Authors are invited to submit novel research in all areas of performance modeling, benchmarking and simulation, and we welcome research that brings together current theory and practice. We recognize that the coverage of the term 'performance' has broadened to include power consumption and reliability, and that performance modeling is practiced through analytical methods and approaches based on software tools and simulators.

Performance modeling, benchmarking and simulation will underpin software and hardware design choices as we advance towards the exa-scale era. This workshop continues to attract high quality input from industry, government and academia.

WOLFHPC: Sixth International Workshop on Domain-Specific Languages and High-Level Frameworks for HPC

9am-12:30pm
Room: 251-F

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. 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. Specifically, the workshop considers these areas: application frameworks, domain-specific languages and libraries, high-level data structures, streaming languages, high-level aspects of HPCS languages, and directive-based optimization approaches.

Women in HPC: Diversifying the HPC Community

9am-5:30pm
Room: 355-D

Toni Collis (EPCC at the University of Edinburgh), Julia Andrys (Murdoch University), Sunita Chandrasekaran (University of Delaware), Trish Damkroger (Lawrence Livermore National Laboratory), Rebecca Hartman-Baker (Lawrence Berkeley National Laboratory), Daniel Holmes (EPCC at the University of Edinburgh), Adrian Jackson (EPCC at the University of Edinburgh), Kimberly McMahon (McMahon Consulting), Lorna Rivera (University of Illinois), Lorna Smith (EPCC at the University of Edinburgh)

The fifth international Women in HPC workshop will be held at SC16, Salt Lake City, USA. The workshop will address a variety of issues relevant to both employers and to employees, specifically to identify particular challenges faced by women, outline opportunities and strategies for broadening participation, and share information on the steps being taken to encourage women into the field and retain a diverse workforce. A keynote address and panel discussions will focus on strategies adopted by employers to diversify the HPC workforce and the efficacy of the approaches. Audience participation is encouraged by the use of breakout sessions.

The second part of the day will focus on providing career development for women working in High Performance Computing. Early career women will showcase their work, and we will provide breakout discussions on 'Skills to Thrive: sharing experiences and advice to succeed in the workplace'. Attendees will be provided with the opportunity to build their network and meet peers and potential mentors.

The day will conclude with a career coaching session led LLNL's Acting Associate Director of Computation, Trish Damkroger, who is also a certified executive coach, to help the attendees be the very best version of themselves as a leader, executive, and whole person, providing early career and mid career women with the tools to enhance their future careers.

Following its successful conclusion, the workshop organizers will publish a whitepaper summarizing the outcomes, outlining best practices, and quantifying the benefits to the HPC community in improving diversity.

Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems

9am-5:30pm
Room: 251-E

Vassil Alexandrov (Barcelona Supercomputing Center), Jack Dongarra (University of Tennessee), Al Geist (Oak Ridge National Laboratory), Christian Engelmann (Oak Ridge National Laboratory)

Novel scalable scientific algorithms are needed in order 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. With the advent of Big Data in the past few years, the need of such scalable mathematical methods and algorithms able to handle data and compute intensive applications at scale becomes even more important.

Scientific algorithms for multi-petaflop and exaflop systems also need to be fault tolerant and fault resilient, since the probability of faults increases with scale. Resilience at the system software and at the algorithmic level is needed as a crosscutting effort. Finally, with the advent of heterogeneous compute nodes that employ standard processors as well as GPGPUs, scientific algorithms need to match these architectures to extract the most performance. This includes different system-specific levels of parallelism as well as co-scheduling of computation. Key science applications require novel mathematics and mathematical models and system software that address the scalability and resilience challenges of current- and future-generation extreme-scale HPC systems.

The Fourth International Workshop on Software Engineering for HPC in Computational Science and Engineering (SE-HPCCSE 2016)

2pm-5:30pm
Room: 251-C

Jeffrey Carver (University of Alabama), Neil Chue Hong (University of Edinburgh), Selim Ciraci (Microsoft Corporation)

Researchers are increasingly using HPC, including GPGPUs and computing clusters, for computational science and engineering (CSE) applications. Unfortunately, when developing HPC software, developers must solve reliability, availability, and maintainability problems in extreme scales, understand domain specific constraints, deal with uncertainties inherent in scientific exploration, and develop algorithms that use computing resources efficiently. Software engineering (SE) researchers have developed tools and practices to support development tasks, including: validation and verification, design, requirements management, and maintenance. HPC CSE software requires appropriately tailored SE tools/methods. The SE-HPCCSE workshop addresses this need by bringing together members of the SE and HPC CSE communities to share perspectives, present findings from research and practice, and generate an agenda to improve tools and practices for developing HPC CSE software. This workshop builds on the success of the 2013-2015 editions.

Workshop on Extreme-Scale Programming Tools

2pm-5:30pm
Room: 155-E

Allen D. Malony (University of Oregon), Felix Wolf (Technical University Darmstadt), Martin Schulz (Lawrence Livermore National Laboratory), William Jalby (University of Versailles)

The path to exascale computing will challenge HPC application developers in their quest to achieve the maximum potential that the machines have to offer. Factors such as limited power budgets, clock frequency variability, heterogeneous load imbalance, hierarchical memories, and shrinking I/O bandwidths will make it increasingly difficult to create high-performance applications. Tools for debugging, performance

measurement and analysis, and tuning will be needed to overcome the architectural, system, and programming complexities envisioned in exascale environments. At the same time, research and development progress for HPC tools faces equally difficult challenges from exascale factors. Increased emphasis on autotuning, dynamic monitoring and adaptation, heterogeneous analysis, and so on will require new methodologies, techniques, and engagement with application teams. This workshop will serve as a forum for HPC application developers, system designers, and tools researchers to discuss the requirements for exascale-enabled tools and the roadblocks that need to be addressed. The workshop is the fifth in a series of SC conference workshops organized by the Virtual Institute - High Productivity Supercomputing (VI-HPS), an international initiative of HPC researchers and developers focused on parallel programming and performance tools for large-scale systems. The workshop includes a keynote address, peer-reviewed technical papers, and a lively panel session.

Monday, November 14

Early Career Program I

8:30am-10am
Room: 251-D

Introduction and Peer Networking

Jeff Hollingsworth (University of Maryland)

This session will start with a brief introduction about the workshop and why the SC conference is hosting it. Your will then have a chance to get to know other participants and learn about each other's career path to date.

Time Management

Liz Jessup (University of Colorado, Boulder), Michael A. Heroux (Sandia National Laboratories)

This session will discuss how to manage your time as you transition from a student to a professional. Topics such as work-life balance, prioritization, teaching vs. research, and service will be discussed.

7th SC Workshop on Big Data Analytics: Challenges and Opportunities

9am-5:30pm
Room: 155-A

Ranga Raju Vatsavai (North Carolina State University), Scott Klasky (Oak Ridge National Laboratory), Manish Parashar (Rutgers University)

Recent decades have witnessed data explosion, and petabyte sized data archives are not uncommon anymore. It is estimated that organizations with high end computing (HEC) infrastructures and data centers are doubling the amount of data that they are archiving every year. On the other hand, computing infrastructures are becoming more heterogeneous. The first three workshops held with SC10, SC11, and SC12 were a great success. Continuing on this success, in addition to the cloud focus, we expanded the scope of this workshop with an emphasis on middleware infrastructure that facilitates efficient data analytics on big data starting from SC13. This 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.

Energy Efficient Supercomputing (E2SC)

9am-5:30pm
Room: 255-D

Darren J. Kerbyson (Pacific Northwest National Laboratory), Kirk Cameron (Virginia Polytechnic Institute and State University), Adolfy Hoisie (Pacific Northwest National Laboratory), David Lowenthal (University of Arizona), Dimitris Nikolopoulos (Queen's 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 a viable exaflop HPC capability, revolutionary methods are required with a stronger integration among hardware features, system software, and applications. Equally important are the capabilities for fine-grained spatial and temporal measurement and control to facilitate these layers for energy efficient computing across all layers. Current approaches for energy efficient computing rely heavily on power efficient hardware in isolation. However, it is pivotal for hardware to expose mechanisms for energy efficiency to optimize power and energy consumption for various workloads. At the same time, high fidelity measurement techniques, typically ignored in data-center level measure-

ment, are of high importance for scalable and energy efficient interplay in different layers of application, system software, and hardware.

Joint International Workshop on Parallel Data Storage and Data Intensive Scalable Computing Systems (PDSW-DISCS)

9am-5:30pm
Room: 155-C

Garth Gibson (Carnegie Mellon University), Yong Chen (Texas Tech University)

The Joint International Workshop on Parallel Data Storage and Data Intensive Scalable Computing Systems (PDSW-DISCS) is a merger of two successful SC workshop communities, the Parallel Data Storage Workshop (PDSW) and Data Intensive Scalable Computing Systems (DISCS). Our goal is to better promote and stimulate researchers' interactions and to better address some of most critical challenges for scientific data storage, management, devices, and processing infrastructure for both traditional compute intensive simulations as well as data-intensive high performance computing solutions. Novel submitted papers are peer-reviewed in the late summer. About 8-12 are selected for presentation at the full day SC workshop, published in either the ACM or IEEE Digital Library, and shepherded for proposal to a journal special issue. The workshop web site can be reached through either www.pdsw.org or discl.cs.ttu.edu/discs.

LLVM-HPC2016: Third Workshop on the LLVM Compiler Infrastructure in HPC

9am-5:30pm
Room: 251-F

Hal Finkel (Argonne National Laboratory)

LLVM, winner of the 2012 ACM Software System Award, has become an integral part of the software-development ecosystem for optimizing compilers, dynamic-language execution engines, source-code analysis and transformation tools, debuggers and linkers, and a whole host of programming language and toolchain-related components. Now heavily used in both academia and industry, where it allows for rapid development of production-quality tools, LLVM is increasingly used in work targeted at high-performance computing. Research in and implementation of programming language analysis, compilation, execution, and profiling has clearly benefited from the availability of a high-quality, freely-available infrastructure on which to build. This workshop will focus on recent developments, from both academia and industry, that build on LLVM to advance the state-of-the-art in HPC.

Machine Learning in HPC Environments

9am-12:30pm
Room: 355-D

Robert Patton (Oak Ridge National Laboratory), Barry Chen (Lawrence Livermore National Laboratory), Lawrence Carin (Duke University)

The intent of this workshop is to bring together researchers, practitioners, and scientific communities to discuss methods that utilize extreme scale systems for machine learning. This workshop will focus on the greatest challenges in utilizing HPC for machine learning and methods for exploiting data parallelism, model parallelism, ensembles, and parameter search. We invite researchers and practitioners to participate in this workshop to discuss the challenges in using HPC for machine learning and to share the wide range of applications that would benefit from HPC powered machine learning.

In recent years, the models and data available for machine learning (ML) applications have grown dramatically. High performance computing (HPC) offers the opportunity to accelerate performance and deepen understanding of large data sets through machine learning. Current literature and public implementations focus on either cloud-based or small-scale GPU environments. These implementations do not scale well in HPC environments due to inefficient data movement and network communication within the compute cluster, originating from the significant disparity in the level of parallelism. Additionally, applying machine learning to extreme scale scientific data is largely unexplored. To leverage HPC for ML applications, serious advances will be required in both algorithms and their scalable, parallel implementations.

MTAGS16: 9th Workshop on Many-Task Computing on Clouds, Grids, and Supercomputers

9am-12:30pm
Room: 251-B

Ke Wang (Intel Corporation), Justin Wozniak (Argonne National Laboratory), Ioan Raicu (Illinois Institute of Technology)

The 9th workshop on Many-Task Computing on Clouds, Grids, and Supercomputers (MTAGS) will provide the scientific community a dedicated forum for presenting new research, development, and deployment efforts of large-scale many-task computing (MTC) applications on large scale clusters, clouds, grids, and supercomputers. MTC, the theme of the workshop encompasses loosely coupled applications, which are generally composed of many tasks to achieve some larger application goal. This workshop will cover challenges that can hamper efficiency and utilization in running applications on

large-scale systems, such as local resource manager scalability and granularity, efficient utilization of raw hardware, parallel file-system contention and scalability, data management, I/O management, reliability at scale, and application scalability. We welcome paper submissions in theoretical, simulations, and systems topics with special consideration to papers addressing the intersection of petascale/exascale challenges with large-scale cloud computing. Papers will be peer-reviewed, and accepted papers will be published in the workshop proceedings as part of the ACM SIGHPC. The workshop will be co-located with the IEEE/ACM Supercomputing 2016 Conference in Salt Lake City on November 14th, 2016. More information could be found at: https://sites.google.com/site/mtags2016/.

PyHPC 2016: 6th Workshop on Python for High-Performance and Scientific Computing

9am-5:30pm
Room: 155-E

Andreas Schreiber (German Aerospace Center), William Scullin (Argonne National Laboratory), Bill Spotz (Sandia National Laboratories), Andy R. Terrel (Fashion Metric Inc)

The high-level programming language Python is well established with a large community in academia and industry. It is a general-purpose language adopted by many scientific applications. Examples are computational fluid dynamics, biomolecular simulation, machine learning, data analysis, and scientific visualization. Scientists, engineers, and educators use Python for scientific computing, high-performance computing, and distributed computing. Traditionally, system administrators use Python for system management and automating administration tasks. Python is extremely easy to learn due to its very clean syntax and great readability. Therefore, developers love Python as it facilitates writing sustainable and maintainable software systems. For the same reasons, Python is well suited for education at all levels.

The overarching theme of the workshop is productivity vs. performance in HPC and scientific programming. While Python is extremely strong in supporting human productivity, it still lacks in computational performance compared to 'traditional' HPC languages such as Fortran or C. For the workshop, we encourage authors to submit novel research in improving performance of Python applications as well as research on productivity of development with Python.

The workshop will bring together researchers and practitioners using Python in all aspects of high performance and scientific computing. The goal is to present Python applications from mathematics, science, and engineering, to discuss general topics regarding the use of Python, and to share experiences using Python in scientific computing education.

More at: http://www.dlr.de/sc/pyhpc2016

Second International Workshop on Heterogeneous Computing with Reconfigurable Logic

9am-12:30pm
Room: 251-C

Michaela Blott (Xilinx), Torsten Hoefler (ETH Zurich), Michael Lysaght (Irish Centre for HighEnd Computing), Jason Bakos (University of South Carolina)

The advent of high-level synthesis (HLS) creates exciting new opportunities for using FPGAs in HPC. HLS allows programs written in OpenCL, C, etc. to be mapped directly and effectively to FPGAs, without the need for low-level RTL design. At the same time, FPGA-based acceleration presents the opportunity for dramatic improvements in performance and energy-efficiency for a variety of HPC applications. This workshop will bring together application experts, FPGA experts, and researchers in heterogeneous computing to present cutting-edge research and explore opportunities and needs for future research in this area.

More at: http://h2rc.cse.sc.edu/

Taking Supercomputing to the Clinic: Medical Image Analysis and Visualization

9am-5:30pm
Room: 251-A

Anthony Costa (Icahn School of Medicine at Mount Sinai), Patricia Kovatch (Icahn School of Medicine at Mount Sinai), Chris Johnson (University of Utah), Alan Tannenbaum (Stony Brook University), Zahi Fayad (Icahn School of Medicine at Mount Sinai), Lalitha Shankar (National Institutes of Health), Tom Fogal (NVIDIA Corporation)

Interest in quantitative image analysis and visualization in healthcare, from population-level academic research studies to patient-specific analysis has grown precipitously in the recent past. While applications to these tools for large-scale research have found significant support from the research community developing these tools, opportunities at smaller scale, especially at the clinical level where most healthcare effort and research is executed, have been extremely limited. Tools appropriate at the population scale are generally inaccessible to these communities, while their potential value has exploded. Interest and awareness of these tools has never

been higher, as physicians and researchers have recognized how HPC and visualization can complement and even drive their own work. Lines of communication between these intersecting fields, however, have been extremely limited, stifling engagement and deployment of the most advanced imaging analysis and technology for healthcare research and clinical practice. Our workshop will bring together these disparate communities to engage on problems and solutions encountered in the practice of image analysis and visualization in healthcare. We will use this joint collaborative development to specifically address the needs of both communities. Thought leaders from across the aisle divide will come together to discuss problems where HPC, imaging analysis, and visualization can dramatically impact patient care. In the afternoon, we will invite proposals for lightning talks on a key technology, tool, or application involving medical image analysis and visualization together with an interactive demonstration of that technology. Submissions will be linked to the SC16 Visualization Showcase.

The 1st International Workshop on Post-Moore Era Supercomputing (PMES)

9am-5:30pm
Room: 155-B

Satoshi Matsuoka (Tokyo Institute of Technology), Jeffrey S. Vetter (Oak Ridge National Laboratory), Koji Inoue (Kyushu University), Takeshi Iwashita (Hokkaido University), Franz Franchetti (Carnegie Mellon University), John Shalf (Lawrence Berkeley National Laboratory), Kengo Nakajima (University of Tokyo), Richard Vuduc (Georgia Institute of Technology), Keren Bergman (Columbia University), Tom Conte (Georgia Institute of Technology), Gerhard Wellein (University of Erlangen-Nuremberg), Erik P. DeBenedictis (Sandia National Laboratories)

This interdisciplinary workshop is organized to explore the scientific issues, challenges, and opportunities for supercomputing beyond the scaling limits of Moore's Law, with the ultimate goal of keeping supercomputing at the forefront of computing technologies beyond the physical and conceptual limits of current systems. Moore's Law—the doubling the number of transistors in a chip every two years—has so far contributed to every aspect of supercomputing system architectures, including GPU and many-core accelerators, large on-chip caches, integrated special purpose hardware, and increasing memory capacities. However, it is now well accepted that current approaches will reach their limits in next decade due to the confluence of several limitations including both fundamental physics and economics. Although device and manufacturing technologies continue to make progress, most experts predict that CMOS transistor shrinking may stop at around 2025 to 2030 due to these limits. Nevertheless, continuing the progress of supercomputing beyond the scaling limits of Moore's Law is likely to require a comprehensive re-thinking of technologies, ranging from innovative materials and devices, circuits, system architectures, programming systems, system software, and applications. In this regard, the goal of the workshop is to explore the technological directions of supercomputing to prepare for this "Post Moore's Law" era by fostering interdisciplinary dialog across the spectrum of stakeholders: applications, algorithms, software, and hardware. Experts from academia, government, and industry in the fields of computational science, mathematics, engineering, and computer science will participate in the workshop as invited speakers, position papers, and panelists.

The Seventh International Workshop on Data-Intensive Computing in the Cloud

9am-12:30pm
Room: 251-E

Yong Zhao (University of Electronic Science and Technology of China), Boyu Zhang (Purdue University), Wei Tang (Google)

Applications and experiments in all areas of science are becoming increasingly complex and more demanding in terms of their computational and data requirements. Some applications generate data volumes reaching hundreds of terabytes and even petabytes. As scientific applications become more data intensive, the management of data resources and dataflow between the storage and compute resources is becoming the main bottleneck. Analyzing, visualizing, and disseminating these large datasets 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.

Third SC Workshop on Best Practices for HPC Training

9am-12:30pm
Room: 155-F

Scott Lathrop (University of Illinois), Jay Alameda (University of Illinois), Carl Albing (US Naval Academy), Nia Alexandrov (Barcelona Supercomputing Center), Dana Brunson (Oklahoma State University), Fernanda Foertter (Oak Ridge National Laboratory), Richard Gerber (Lawrence Berkeley National Laboratory), Bilel Hadri (King Abdullah University of Science and Technology), Rebecca Hartman-Baker (Lawrence Berkeley National Laboratory), Susan Mehringer (Cornell University), Henry Neeman (University of Oklahoma), Maria-Ribera Sancho (Barcelona Supercomputing Center), Nitin Sukhija (Mississippi State University), Robert Whitten (University of Tennessee)

This is a request to conduct a third annual half-day workshop on HPC training during the SC16 Conference.

Community interest in this topic continues to grow as the organizers work to promote this important topic internationally. The SC15 workshop organizers actively recruited more than 170 people to the SC15 event, more than twice the SC14 workshop attendance. We believe attendance also increased as a result of the workshop being conducted on Monday during SC15, rather than being held on Friday during SC14.

As a result of this effort, the International HPC Training Consortium was established to facilitate planning for the SC workshops, and for coordinating year-long efforts to sustain the interests identified during the SC workshops. Membership in the Consortium continues to grow and includes over 70 people from 15 countries.

The SC16 workshop will be used to highlight the results of collaborative efforts during 2016 to develop and deploy HPC training, to identify new challenges and opportunities, and to foster new, enhanced and expanded collaborations to pursue during 2017.

Early Career Program II

10:30am-12pm
Room: 251-D

Planning your Career – Thinking Strategically

Dan Reed (University of Iowa), Cherri Pancake (Oregon State University)

You career should be more than a series of positions that just happened. Likewise, careers rarely take an exactly planned route. In this session, two senior members of the HPC community will discuss their career paths, and how they happened to evolve the way the did. They will also provide some tips on thinking about how to reach your long term goals for your career.

Career Networking

Richard Vuduc (Georgia Institute of Technology), Bernd Mohr (Forschungszentrum Juelich)

It has been said that who you know matters more than what you know. While that is probably not true, developing a diverse network of professional colleagues can be helpful in having both a productive and enjoyable career. In this session, the presenters will describe the process of developing and curating your professional network of friends, colleagues, and mentors.

Early Career Program III

1:30pm-3pm
Room: 251-D

Funding Your Research and Grant Writing

Padma Raghavan (Vanderbilt University), Almadena Y. Chtchelkanova (National Science Foundation)

The reality is it takes money to conduct your research. In this session, presenters will describe the various types of funding available to support your research. Examples of different federal agencies, corporations, and laboratory funding sources will be covered. In addition, they will discuss the elements of a successful grant proposal.

Publication Venues and the Review Process

Jeffrey K. Hollingsworth (University of Maryland)

Publish or Perish has been the motto of research for decades. However, there are many possible venues: web, workshops, conferences, journals, and books. In this session the speaker will describe how the eco-system of publishing venues co-exist and which ones might be appropriate for different types of work. He will also discuss the manuscript review process and

provide suggestions to help get your work the best chance of getting a favorable review.

HPC Systems Professional Workshop

2pm-5:30pm
Room: 155-F

Isaac Traxler (Louisiana State University), Jenett Tillotson (Indiana University), Randy Herban (Purdue University), Henry Neeman (University of Oklahoma), Stephen Lien Harrell (Purdue University), Prentice Bisbal (Rutgers University), William Scullin (Argonne National Laboratory), Robert Ping (Indiana University)

In order to meet the demands of HPC researchers, large-scale computational and storage machines require many staff members who design, install, and maintain these systems. These HPC systems professionals include system engineers, system administrators, network administrators, storage administrators, and operations staff who face problems that are unique to HPC systems. While many conferences exist for the HPC field and the system administration field, none exist that focus on the needs of HPC systems professionals. Support resources can be difficult to find to help with the issues encountered in this specialized field. Often times, systems staff turn to the community as a support resource and opportunities to strengthen and grow those relationships are highly beneficial. This workshop is designed to share solutions to common problems, provide a platform to discuss upcoming technologies, and present state of the practice techniques so that HPC centers will get a better return on their investment, increase performance and reliability of systems, and researchers will be more productive.

PAW: PGAS Applications Workshop

2pm-5:30pm
Room: 251-B

Karla Morris (Sandia National Laboratories), Katherine A. Yelick (Lawrence Berkeley National Laboratory), Yili Zheng (Lawrence Berkeley National Laboratory), Salvatore Filippone (Cranfield University), Bill Long (Cray Inc.), Bradford L. Chamberlain (Cray Inc.)

With the advent of extreme-scale computing and the accompanying complexity of emerging architectures and technologies, it is vital to provide software application researchers with effective approaches for developing parallel applications. This workshop focuses on one of these approaches, the Partitioned Global Access Space (PGAS) parallel programming model. The PGAS model simplifies programming complexity while enhancing performance by providing a shared address space abstraction that exposes data locality. This workshop will bring together application experts who will present concrete practical examples to illustrate the benefits of exploiting PGAS languages and libraries.

The 11th Workshop on Workflows in Support of Large-Scale Science

2pm-5:30pm
Room: 355-D

Sandra Gesing (University of Notre Dame), Rizos Sakellariou (University of Manchester)

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.

Third International Workshop on Accelerator Programming Using Directives (WACCPD)

2pm-5:30pm
Room: 251-C

Sunita Chandrasekaran (University of Delaware), Guido Juckeland (Helmholtz-Zentrum Dresden-Rossendorf)

One of the hard realities is that the hardware continues to evolve very rapidly with diverse memory subsystems or cores with different ISAs or accelerators of varied types. The HPC community is in constant need for sophisticated software tools and techniques to port legacy code to these emerging platforms. Maintaining a single code base yet achieving performance portable solution continues to pose a daunting task. Directive-based programming models such as OpenACC,

OpenMP tackle this issue by offering scientists a high-level approach to accelerate scientific applications and develop performance portable solutions. This enables accelerators to be first-class citizens for HPC! To address the rapid pace of hardware evolution, developers continue to explore and add richer features to the various (parallel) programming standards. Domain scientists continue to explore the programming and tools space while preparing themselves for future exascale systems.

This workshop aims to explore innovative language features - their implementations, compilation, and runtime scheduling techniques, performance optimization strategies, autotuning tools exploring the optimization space, and so on. We are looking forward to continuing to host this workshop at SC16. WACCPD has been a major forum for bringing together the users, developers, and tools community to share their knowledge and experiences of using directives and similar approaches to program emerging complex systems.

Workshop on Education for High Performance Computing (EduHPC)

2pm-5:30pm
Room: 251-E

Anshul Gupta (IBM), Sushil Prasad (Georgia State University), Charles Weems (University of Massachusetts), Alan Sussman (University of Maryland), Arnold Rosenberg (Northeastern University), Ioana Banicescu (Mississippi State University)

The EduHPC Workshop is devoted to the development and assessment of educational resources for undergraduate education in High Performance Computing (HPC) and Parallel and Distributed Computing (PDC). Both PDC and HPC now permeate the world of computing to a degree that makes it imperative for even entry-level computer professionals to incorporate these computing modalities into their computing kitbags, no matter what aspect of computing they work on. This workshop focuses on the state-of-the-art in HPC and PDC education, by means of both contributed and invited papers from academia, industry, and other educational and research institutions. Topics of interest include all topics pertaining to the teaching of PDC and HPC within Computer Science and Engineering, Computational Science, and Domain Science and Engineering curricula. The emphasis of the workshop is undergraduate education, but fundamental issues related to graduate education are also welcome. The target audience will broadly include SC16 attendees from academia and industry, including both researchers and educators, as well as the early adopters of NSF/TCPP curriculum on PDC (http://www.cs.gsu.edu/~tcpp/curriculum/index.php). The workshop is coordinated by the CDER Center for PDC Education and highlights the NSF/TCPP curriculum initiative. EduHPC-13 was the first education-related regular workshop held at SC and had excellent participation, but was focused on curriculum design. EduHPC-14 and EduHPC-15 shifted their emphasis to sharing resources for undergraduate education and were very successful.

Speed Mentoring

3:30pm-5pm
Room: 251-D

This session is only open to accepted members of the SC16 Early Career program. In this session, program participants will have a chance to meet 5-7 potential career mentors. Each round will allow potential mentors and protégés 5-10 minutes to get to know each other before the bell rings and it is on to meet more people!

Friday, November 18

3rd International Workshop on Visual Performance Analytics – VPA 2016

8:30am-12pm
Room: 355-E

Martin Schulz (Lawrence Livermore National Laboratory), Peer-Timo Bremer (Lawrence Livermore National Laboratory), Judit Gimenez (Barcelona Supercomputing Center), Joshua Levine (University of Arizona)

Over the last decades an incredible resources have been devoted to building ever more powerful supercomputers. However, exploiting the full capabilities of these machines is becoming exponentially more difficult with each new generation of hardware. To help understand and optimize the behavior of massively parallel simulations, the performance analysis community has created a wide range of tools and APIs to collect performance data, such as flop counts, network traffic, and cache behavior at the largest scale. However, this success has created a new challenge, as the resulting data is far too large and too complex to be analyzed in a straightforward manner. Therefore, new automatic analysis and visualization approaches must be developed to allow application developers to intuitively understand the multiple, interdependent effects that their algorithmic choices have on the final performance. This workshop intends to bring together researchers from the fields of performance analysis and visualization to discuss new approaches of combining both areas to analyze and optimize large-scale applications.

ESPM2 2016: Second International Workshop on Extreme Scale Programming Models and Middleware

8:30am-12pm
Room: 155-E

Khaled Hamidouche (Ohio State University), Karl Schulz (Intel Corporation), Hari Subramoni (Ohio State University), Dhabaleswar K. (DK) Panda (Ohio State University)

Next generation architectures and systems being deployed are characterized by high concurrency, low memory per-core, and multiple levels of hierarchy and heterogeneity. These characteristics bring out new challenges in energy efficiency, fault-tolerance and scalability. It is commonly believed that software has the biggest share of the responsibility to tackle these challenges. In other words, this responsibility is delegated to the next generation programming models and their associated middleware/runtimes. This workshop focuses on different aspects of programming models such as task-based parallelism (Charm++, OCR, X10, HPX, etc), PGAS (OpenSHMEM, UPC, CAF, Chapel, UPC++, etc.), directive-based languages (OpenMP, OpenACC) and hybrid MPI+X, etc. It also focuses on their associated middleware (unified runtimes, interoperability for hybrid programming, tight integration of MPI+X, and support for accelerators) for next generation systems and architectures. The ultimate objective of the ESPM2 workshop is to serve as a forum that brings together researchers from academia and industry working on the areas of programming models, runtime systems, compilation and languages, and application development.

Exascale I/O: Challenges, Innovations and Solutions (ExaIO)

8:30am-12pm
Room: 355-BC

Michele Weiland (EPCC at the University of Edinburgh), Mark Parsons (EPCC at the University of Edinburgh), Hans-Christian Hoppe (Intel Corporation), Sai Narasimhamurthy (Seagate Technology LLC), Nobert Eicker (Forschungszentrum Juelich)

Many of today's modeling and simulation applications struggle to achieve good parallel efficiency and performance on current petascale systems. The high discrepancy between a system's peak performance and the performance that can be achieved by applications is often blamed on the inter-node communications networks on these systems. In truth, internal communications networks represent very large improvements on the inter-node communication networks of even 5 years ago. The limited efficiency seen is chiefly a function of the lack of explicit parallelism exhibited by the underlying algorithms, or utilized by the application codes. More recently, I/O performance has started to play a critical role, so that even if the computational part of a workload can scale to extreme parallel systems at exascale, reading and writing the data associated with it will present a major challenge. It is rarely acknowledged is that, as core-counts have increased, the performance of I/O subsystems have struggled to keep up with computational performance and have, over the past few years, become a key bottleneck on today's largest systems.

This workshop will: * Describe the I/O related challenges that both HPC and high performance data analytics (HPDA) face going forward; * Examine how increasingly complex memory and storage hierarchies, and recent technological innovations such as Intel's 3D XPointTM non-volatile memory, could have an impact on the I/O challenge; and * Discuss potential hardware and software solutions for dealing with the extreme volumes of data involved in scientific (and industrial) computing today and in the future

First International Workshop on Communication Optimizations in HPC

8:30am-12pm
Room: 355-D

Akhil Langer (Intel Corporation), Maria Garzaran (Intel Corporation), Gengbin Zheng (Intel Corporation), Daniel Faraj (Intel Corporation), Malek Musleh (Intel Corporation), Michael Chuvelev (Intel Corporation)

As HPC applications scale to large supercomputing systems, their communication and synchronization need to be optimized in order to deliver high performance. To achieve this, capabilities of modern network interconnect and parallel runtime systems need to be advanced and the existing ones to be leveraged optimally. The workshop will bring together researchers and developers to present and discuss work on optimizing communication and synchronization in HPC applications. This includes, but is not limited to, methodological and algorithmic advances in topology-aware or topology-oblivious blocking and non-blocking collective algorithms, offloading of communication to network interface cards, topology aware process mappings for minimizing communication overhead on different network topologies such as dragonfly, high-dimensional torus networks, fat trees, optimizations for persistent communication patterns, studies and solutions for inter-job network interference, overlapping of communication with computation, optimizing communication overhead in presence of process imbalance, static or runtime tuning of collective operations, scalable communication

endpoints for manycore architectures, network congestion studies and mitigation methods, communication optimizations on peta/exa-scale systems, heterogeneous systems, and GPUs, machine learning to optimize communication, and communication aspects of GPGPU, graph applications, or fault tolerance. The workshop also aims at bringing researchers together to foster discussion, collaboration, and ideas for optimizing communication and synchronization that drive design of future peta/exascale systems and of HPC applications. In addition, we expect that researchers and others looking for research directions in this area will get up-to-date with the state-of-the-art so that they can drive their research in a manner that will impact the future of HPC.

NRE2016: Numerical Reproducibility at Exascale

8:30am-12pm
Room: 155-C

Michael Mascagni (Florida State University), Walid Keyrouz (National Institute of Standards and Technology)

A cornerstone of the scientific method is experimental reproducibility. As computation has grown into a powerful tool for scientific inquiry, the assumption of computational reproducibility has been at the heart of numerical analysis in support of scientific computing. With ordinary CPUs, supporting a single, serial, computation, the ability to document a numerical result has been a straight-forward process. However, as computer hardware continues to develop, it is becoming harder to ensure computational reproducibility, or to even completely document a given computation. This workshop will explore the current state of computational reproducibility in HPC, and will seek to organize solutions at different levels. The workshop will conclude with a panel discussion aimed at defining the current state of computational reproducibility for the exascale. We seek contributions in the areas of computational reproducibility in HPC.

This workshop will produce a refereed proceedings that will be available through the ACM Digital Library and IEEE Xplore (free of charge during and immediately after SC, and free after that to SIGHPC members). More at: http://www.cs.fsu.edu/~nre/nre2016.html

Runtime Systems for Extreme Scale Programming Models and Architectures (RESPA)

8:30am-12pm
Room: 155-A

Vivek Sarkar (Rice University), Siegfried Benkner (University of Vienna), Ron Brightwell (Sandia National Laboratories), Patrick S. McCormick (Los Alamos National Laboratory)

Extreme-scale and exascale systems impose new requirements on application developers and programming systems to target platforms with hundreds of homogeneous and heterogeneous cores, as well as energy, data movement, and resiliency constraints within and across nodes. Runtime systems can play a critical role in enabling future programming models, execution models, and hardware architectures to address these challenges, and in reducing the widening gap between peak performance and the performance achieved by real applications.

The goal of this workshop is to attract leading international researchers to share their latest results involving runtime approaches to address these extreme-scale and exascale software challenges. The scope of the workshop includes (but is not limited to) runtime system support for: high-level programming models and domain-specific languages; scalable intra-node and inter-node scheduling; memory management across coherence domains and vertical hierarchies of volatile/non-volatile storage; optimized locality and data movement; energy management and optimization; performance tuning; and resilience.

Acknowledgements

SC16 Conference Management

Conference Chair
John West, University of Texas at Austin

Vice Chair
Cherri M. Pancake, Oregon State University

Executive Director
Christine E. Cuicchi, ~

Deputy Chair
Bernd Mohr, Juelich Supercomputing Center

Special Assistant to the Chair
Patty Gropp, Independent
Matt Link, Indiana University

Executive

Conference Chair
John West, University of Texas at Austin

Vice Chair
Cherri M. Pancake, Oregon State University

Deputy Chair
Bernd Mohr, Juelich Supercomputing Center

Diverse HPC Workforce Chair
Trish Damkroger, Lawrence Livermore National Laboratory

Executive Director
Christine E. Cuicchi, ~

Finance Chair
Janet M. McCord, University of Texas at Austin

Infrastructure Chair
Jay Cliburn, SMS

Communications Chair
Leah Glick, Linklings LLC

Local Arrangements Chair
Becky Verastegui, Oak Ridge National Laboratory

SCinet Chair
Corby Schmitz, Argonne National Laboratory

Technical Program Chair
Lori Diachin, Lawrence Livermore National Laboratory

Special Assistant to the Chair
Patty Gropp, Independent
Matt Link, Indiana University

IEEE-CS
Anne Marie Kelly, IEEE Computer Society
June Leach-Barnaby, IEEE Computer Society

Exhibits Chair
Trey Breckenridge, Mississippi State University

Students @ SC Chair
Jeanine Cook, Sandia National Laboratories

ACM
Ashley Cozzi, Association for Computing Machinery

Diverse HPC Workforce

Diverse HPC Workforce Chair
Trish Damkroger, Lawrence Livermore National Laboratory

Diverse HPC Workforce Member
Mike Bernhardt, Oak Ridge National Laboratory
Toni Collis, EPCC at the University of Edinburgh
Kim Cupps, Lawrence Livermore National Laboratory
Kelly Gaither, University of Texas at Austin
Debra Goldfarb, Intel Corporation
Mary Hester, Lawrence Berkeley National Laboratory
Bruce Loftis, University of Colorado, Boulder
Kimberly A. McMahon, McMahon Consulting
Lorna Rivera, Georgia Institute of Technology
Lisa Smith, Intel Corporation
John West, University of Texas at Austin

HPC Matters

HPC Matters Chair
Wilfred Pinfold, Concurrent Systems LLC

HPC Matters Member
Trey Breckenridge, Mississippi State University
Christine E. Cuicchi, ~
Brent Gorda, Intel Corporation
Dimitrios Nikolopoulos, Queen's University Belfast
Ron Perrott, University of Oxford
Lauren Rotman, Lawrence Berkeley National Laboratory

Keynote

Keynote Chair
Dona Crawford, Lawrence Livermore National Laboratory

Keynote Contractor
Sponsoring Societies

ACM
Ashley Cozzi, Association for Computing Machinery

IEEE-CS
Anne Marie Kelly, IEEE Computer Society
June Leach-Barnaby, IEEE Computer Society

Communications

Communications Chair
Leah Glick, Linklings LLC

Communications Deputy Chair
Lauren Rotman, Lawrence Berkeley National Laboratory

Graphics & Website Contractor
Carlton Bruett, Carlton Bruett Design

Media Partnerships
Mike Bernhardt, Oak Ridge National Laboratory
April Hanks, Intel Corporation

Media Relations Contractor
Brian Ban, Ban Communications

Newsletter
Kathy Kincade, Lawrence Berkeley National Laboratory

Press & Media
April Hanks, Intel Corporation
Don Johnston, Lawrence Livermore National Laboratory

Publications
Christy Adkinson, Cray Inc.
Vivian M. Benton, Pittsburgh Supercomputing Center
Bruce Loftis, University of Colorado, Boulder

Wayfinding
Christy Adkinson, Cray Inc.

Web Content
Vivian M. Benton, Pittsburgh Supercomputing Center

Writer
Jamie Abel, Ohio Technology Consortium
Jon Bashor, Lawrence Berkeley National Laboratory
Douglas Beary, Lenovo
Don Johnston, Lawrence Livermore National Laboratory
Barbara Mack, Pingry Hill Enterprises Inc
Kimberly A. McMahon, McMahon Consulting

Plenary Production

Plenary Production Chair
Lauren Rotman, Lawrence Berkeley National Laboratory

Social Media

Social Media Member

Christine E. Harvey, MITRE Corporation
Gyongyi Horvath, GEANT
LeAnne Reed, McMahon Consulting
Faith Singer-Villalobos, University of Texas at Austin

Exhibits

Exhibits Chair
Trey Breckenridge, Mississippi State University

Deputy Exhibits Chair
Bronis R. de Supinski, Lawrence Livermore National Laboratory

Exhibitor Forum
Jackie Kern, University of Illinois

HPC Impact Showcase
David Martin, Argonne National Laboratory

HPC Impact Showcase Vice Chair
David Halstead, National Radio Astronomy Observatory

Exhibits Management Contractor
Paul Graller, Hall-Erickson Inc
Ann Maffeo, Hall-Erickson Inc
Mike Weil, Hall-Erickson Inc

Event Services Contractor
Stephen E. Hagstette Jr, Freeman Company
Regina Martin, Freeman Company
Darryl Monahan, Freeman Company

Financial Management

Finance Chair
Janet M. McCord, University of Texas at Austin

Finance Deputy Chair
Michela Taufer, University of Delaware

Finance Contractor
Danielle Bertrand, Talbot, Korvola & Warwick LLP
Anne Nottingham, Talbot, Korvola & Warwick LLP

Registration/Store/Merchandise

Registration/Store/Merchandise Chair
Michele Bianchini-Gunn, Lawrence Livermore National Laboratory

Deputy Registration Chair
Hai Ah Nam, Los Alamos National Laboratory

Registration/Store/Merchandise Contractor
Katie Adair, J. Spargo and Associates
Linda Sheeran, Linda Sheeran Promotions
Karen Shipe, J. Spargo and Associates

Infrastructure

Infrastructure Chair
Jay Cliburn, SMS

Infrastructure Vice Chair
Eric Sills, North Carolina State University

Infrastructure Deputy Chair
Tim Yeager, ~

Space
James W. Ferguson, University of Tennessee

Electrical
Gary New, National Center for Atmospheric Research

Signage
Dustin Leverman, Oak Ridge National Laboratory

Committee Office
Janie Bushn, University of Texas at Austin
Carolyn Peters, Argonne National Laboratory

Infrastructure Contractor
Mark Montague, Linklings LLC

AV/PC

AV/PC Chair
Douglas Fuller, Red Hat Inc

AV/PC Contractor
Gabe Blakeney, AV Concepts
Danielle Martinez, AV Concepts
James Raitt, AV Concepts
Ryan Raitt, Rezin8

Security

Security Chair
Tim Yeager, ~

Security Contractor
Peter Alexan, RA Consulting

Security Team
Jackie Kern, University of Illinois
Clea Marples, Lawrence Livermore National Laboratory
Ralph A. McEldowney, US Department of Defense HPC Modernization Program
Jamie Van Randwyk, Lawrence Livermore National Laboratory

Local Arrangements

Local Arrangements Chair
Becky Verastegui, Oak Ridge National Laboratory

Catering & Special Events
Barbara Horner-Miller, BHM Consulting LLC

Local Arrangements Contractor
Kevin Fox,
Karly Kopra,
Rebecca Mebane, MeetGreen
Shereen Vesalpour, MeetGreen

Housing

Housing Chair
Bronis R. de Supinski, Lawrence Livermore National Laboratory

Housing Deputy Chair
Ewa Deelman, University of Southern California

Housing Contractor
Ashley Garrett, Orchid Event Solutions

SCinet

SCinet Chair
Corby Schmitz, Argonne National Laboratory

SCinet Vice Chair
Linda Winkler, Argonne National Laboratory

SCinet Executive Director
Jason Zurawski, Energy Sciences Network

SCinet Deputy Chair
Jeffrey R. Schwab, Purdue University

SCinet Vendor Engagement
Jane Combs, University of Cincinnati
Jason Lee, Lawrence Berkeley National Laboratory

Architecture

Architecture Team Co-Chair
Marc Lyonnais, Ciena Corporation
Matt Smith, National Oceanic and Atmospheric Administration

Architecture Team Member
Kate P. Mace, Energy Sciences Network

Commodity Network

Commodity Network Team Co-Chair
Ryan Braithwaite, Los Alamos National Laboratory
Kevin Hayden, Argonne National Laboratory

Commodity Network Team Member
Angie Asmus, Colorado State University

Christine Baissac, Easy English 4 All
Zach Bare, Keystone Initiative for Network Based Education and Research
Aaron Botello, University of Colorado, Boulder
Parks Fields, Los Alamos National Laboratory
Mary Hester, Lawrence Berkeley National Laboratory
Dave Modl, Los Alamos National Laboratory

Commodity Network Vendor
Rey Crespo, Smart City Networks
Chris Pendleton, Smart City Networks

Communications

Communications Team Co-Chair
Mary Hester, Lawrence Berkeley National Laboratory
Gyongyi Horvath, GEANT

Communications Team Member
Rich Finlinson, Utah Education Network
Sylvia Kuijpers, SURFnet
Amber Rasche, North Dakota State University

Fiber

Fiber Team Co-Chair
Rebecca Hutchinson, SCinet
Annette E. Kitajima, Sandia National Laboratories

Fiber Team Member
Shane Archiquette, Hitachi Ltd
Kent M. Goldsmith, CABLExpress Corporation
Kristy Hamm, Pennsylvania State University
Christer Karlsson, South Dakota School of Mines and Technology
Julia Locke, Los Alamos National Laboratory
Camille Magno, JDS Uniphase Corporation
Miles McGee, US Air Force Research Laboratory

Help Desk

Help Desk Team Co-Chair
Virginia Bedford, US Army Engineer Research and Development Center
Timothy W. Dunaway, US Army Engineer Research and Development Center
Karmen Goddard, Independent

INDIS Workshop

INDIS Workshop Team Co-Chair
Cees de Laat, University of Amsterdam
Malathi Veeraraghavan, University of Virginia

INDIS Workshop Team Member
John Graham, University of California, San Diego
Paola Grosso, University of Amsterdam
Sylvia Kuijpers, SURFnet
Brian L. Tierney, Lawrence Berkeley National Laboratory

IT/Measurement

IT/Measurement Team Co-Chair
Benjamin Grover, Lawrence Livermore National Laboratory
Ezra Kissel, Indiana University
Ken Miller, Pennsylvania State University

IT/Measurement Team Member
Scott Chevalier, Indiana University
David Ediger, Georgia Institute of Technology
Greg Goddard, RedSeal Inc
Trevor Goodyear, Georgia Institute of Technology
Julian Hammer, University of Erlangen-Nuremberg
Jon Miller, Franklin and Marshall College
David Mitchell, Energy Sciences Network
Scott Richmond, Energy Sciences Network
Julia Staats, Corporation For Education Network Initiatives In California
Greg Veldman, Purdue University
Kevin Walsh, University of California, San Diego
Brendan White,
Miao Zhang, Indiana University
Brett Zimmerman, University of Oklahoma

IT/Measurement Vendor
Paul Anderson, Puppet Labs
Neil H. McKee, InMon Corporation

Interconnect

Interconnect Team Co-Chair
Brandon George, Data Direct Networks
Lance V. Hutchinson, Sandia National Laboratories

Logistics/Equipment

Logistics/Equipment Team Co-Chair
Ken Brice, US Army Research Laboratory
James H. Rogers, Oak Ridge National Laboratory

Logistics/Equipment Team Member
Frank Indiviglio, National Oceanic and Atmospheric Administration

Network Research Exhibition

Network Research Exhibition Team Co-Chair
Brian L. Tierney, Lawrence Berkeley National Laboratory
Caroline Carver Weilhamer, Indiana University
Matthew J. Zekauskas, Internet2

Network Research Exhibition Team Member
Nick Buraglio, Energy Sciences Network
Wallace Chase, Washington State University
Cees de Laat, University of Amsterdam
Alessandra Scicchitano, GEANT
Kevin Walsh, University of California, San Diego

Network Research Exhibition Vendor
Yatish Kumar, Corsa
Carolyn Raab, Corsa
Stacey Sheldon, Corsa
David Whittaker, Corsa
Network Security

Network Security

Network Security Team Co-Chair
Jeff W. Boote, Sandia National Laboratories
Alan Commike, Reservoir Labs

Network Security Team Member
Aaron Brown, Sandia National Laboratories
Michael Dopheide, Energy Sciences Network
Carrie Gates, Independent
Clark Gaylord, Virginia Polytechnic Institute and State University
Denise Grayson, Sandia National Laboratories
Thomas Kroeger, Sandia National Laboratories
Patrick Storm, University of Texas at Austin
Vince Urias, Sandia National Laboratories
Brian Wright, Sandia National Laboratories
Nick Young, US Army

Network Security Vendor
Neal Allen, Gigamon
Alan Commike, Reservoir Labs
Timothy, Covel,,
Traci Dawson, Splunk Corporation
Eric DeHaas, Juniper Networks
Josh Friedman, Palo Alto Networks Inc
Greg Goddard, RedSeal Inc
Allen Hebert, Gigamon
Robert Ricci, University of Utah
Livio, Ricciulli,, MetaFlows Inc
Susan Stover, Palo Alto Networks Inc
Craig Strait, Palo Alto Networks Inc
Noam Syrkin, Gigamon

Power

Power Team Co-Chair
Jay Harris, Clemson University
Jim Pepin, Clemson University

Routing

Routing Team Co-Chair
Conan Moore, University of Colorado, Boulder
JP Velders, University of Amsterdam

Routing Team Member
Evangelos Chaniotakis, Energy Sciences Network
Jamie Curtis, Research and Education Advanced Network New Zealand
Pieter de Boer, SURFnet
Jacob Evans, University of Utah
Debbie Fligor, University of Illinois
Jackson Gor, Lawrence Berkeley National Laboratory
Ryan Harden, University of Chicago
Zach Harlan, Energy Sciences Network
Thomas Hutton, University of California, San Diego
Dylan Jacob, Rice University
Indira Kassymkhanova, Lawrence Berkeley National Laboratory
Nathan Miller, Indiana University
Jessica Shaffer, Georgia Institute of Technology
Michael Smitasin, Lawrence Berkeley National Laboratory
Alan Verlo, Independent
Davey Wheeler, University of Illinois

Student Coordination

Student Coordination Team Co-Chair
Sarah M. Neuwirth, Heidelberg University
Laura Pettit, Indiana University

WAN Transport

WAN Transport Team Co-Chair
Akbar Kara, Lonestar Education and Research Network
Jim Stewart, Utah Education Network

WAN Transport Team Member
Scott Baily, Colorado State University
Sana Bellamine, Corporation For Education Network Initiatives In California
Barry Bryson, Utah Education Network
John Caffery, Louisiana State University
Mike Downie, Utah Education Network
Ryan Duitman, University of Arizona
Kevin Dutt, Utah Education Network
Bill Jensen, University of Wisconsin
Umar Kalim, Virginia Polytechnic Institute and State University
E. Paul Love, Internet Consulting of Vermont
Kali McLennan,
Brad Miller, Utah Education Network
Scott Oung, University of Texas at Dallas
Kevin Quire, Utah Education Network
Gonzalo P. Rodrigo, Umea University
Chris Tracy, Energy Sciences Network
Gi Vania, University of Texas at Dallas
Ryan Vaughn, Florida LambdaRail
Matthew J. Zekauskas, Internet2

WAN Transport Vendor
Steve Cox, ADVA Optical Networking
Chad Dennis, Infinera Corporation
Tom Felix, Zayo Group
Chris Hunt, Nokia Corporation
Matt Jary, Infinera Corporation
Scott Kohlert, Ciena Corporation
Bruce Larson, Marc Lyonnais, Ciena Corporation
Michael McAfee, CenturyLink
Matt Moser, ADVA Optical Networking
Carlos Pagan, Ciena Corporation
Bill Parr, Cisco Systems
Eric Petrosky, Infinera Corporation
Ori Shinberg, ECI Telecom Ltd
Michael Shuping, CenturyLink
Rod Wilson, Ciena Corporation
Peng Zhu, Alcatel-Lucent

Wireless

Wireless Team Co-Chair
Benny Sparks, Oak Ridge National Laboratory

Wireless Team Member
Jason Charcalla, University of Tennessee
Megan Sorensen, Idaho State University

Wireless Vendor
Matt Chrosniak, Cisco Systems
Zach Duncan, Leverage Information Systems

Students @ SC

Students @ SC Chair
Jeanine Cook, Sandia National Laboratories

Students @ SC Vice Chair
Dorian C. Arnold, University of New Mexico

Students @ SC Deputy Chair
Alan Sussman, University of Maryland

Student Job Fair
Verónica Vergara L., Oak Ridge National Laboratory

Student Outreach
Christine E. Harvey, MITRE Corporation

Mentor-Protégé Program
Christine Sweeney, Los Alamos National Laboratory

Finance & Arrangements

Finance & Arrangements Member
Jon Bashor, Lawrence Berkeley National Laboratory
Sandra Huskamp, ~
Yuho Jin, New Mexico State University
Bruce Loftis, University of Colorado, Boulder
Richard Murphy, Micron Technology Inc

Student Volunteers

Student Volunteers Chair
Yashema Mack, University of Tennessee

Student Volunteers Deputy Chair
Sally Ellingson, University of Kentucky

Student Volunteers Member
Trish Damkroger, Lawrence Livermore National Laboratory
Sally Ellingson, University of Kentucky
Jason Grant, University of Notre Dame
Christine E. Harvey, MITRE Corporation
Raquell Holmes, improvscience
Rubin H. Landau, Oregon State University
Elizabeth Leake, STEM-Trek
Mary Ann Leung, Sustainable Horizons Institute
Yashema Mack, University of Tennessee
Ralph A. McEldowney, US Department of Defense HPC Modernization Program
Sarah M. Neuwirth, Heidelberg University
Laura Pettit, Indiana University
Lorna Rivera, Georgia Institute of Technology
Tabitha K. Samuel, University of Tennessee
Kathryn Traxler, Louisiana State University
Verónica Vergara L., Oak Ridge National Laboratory
Robert Whitten, University of Tennessee

Student Development

Student Development Chair
Verónica Vergara L., Oak Ridge National Laboratory

Student Cluster Competition

Student Cluster Competition Chair
Stephen Lien Harrell, Purdue University

Student Cluster Competition Vice Chair
John Cazes, University of Texas at Austin

Student Cluster Competition Member
Bob Beck, Obtuse Systems Corporation
Christopher Bross, University of Erlangen-Nuremberg
John Cazes, University of Texas at Austin
Sam Coleman, Independent
Charles Garrett, Los Alamos National Laboratory
Kimberly A. McMahon, McMahon Consulting
Scott Michael, Indiana University
Hai Ah Nam, Los Alamos National Laboratory
Ramin Nosrat, IBM
Paul Peltz Jr, Los Alamos National Laboratory
Jenett Tillotson, Indiana University
Xiao Zhu, Purdue University

Student Cluster Competition Reproducibility

Student Cluster Competition Reproducibility Co-Chair
Stephen Lien Harrell, Purdue University
Scott Michael, Indiana University

Student Cluster Competition Reproducibility Committee Member
Charles Garrett, Los Alamos National Laboratory
Michael A. Heroux, Sandia National Laboratories
Hai Ah Nam, Los Alamos National Laboratory
Wilfred Pinfold, Concurrent Systems LLC
Michela Taufer, University of Delaware

Experiencing HPC for Undergraduates

Experiencing HPC for Undergraduates Chair
Alan Sussman, University of Maryland

Experiencing HPC for Undergraduates Vice Chair
Erik Saule, University of North Carolina, Charlotte

Technical Program

Technical Program Chair
Lori Diachin, Lawrence Livermore National Laboratory

Technical Program Deputy Chair
Padma Raghavan, Vanderbilt University

State of the Practice Liaison
William Kramer, University of Illinois
Dan C. Stanzione, University of Texas at Austin

Invited Speakers

Invited Speakers Chair
William D. Gropp, University of Illinois

Invited Speakers Vice Chair
David E. Keyes, King Abdullah University of Science and Technology

Invited Speakers Member
Elizabeth Jessup, University of Colorado, Boulder
Paul Messina, Argonne National Laboratory
Kengo Nakajima, University of Tokyo
Irene Qualters, National Science Foundation
Anne Trefethen, University of Oxford

Awards

Awards Co-Chair
Satoshi Matsuoka, Tokyo Institute of Technology
Rajeev Thakur, Argonne National Laboratory

Test of Time Award

Test of Time Award Chair
Michael A. Heroux, Sandia National Laboratories

Test of Time Award Vice Chair
Elizabeth Jessup, University of Colorado, Boulder

Test of Time Award Member
David Abramson, University of Queensland
William D. Gropp, University of Illinois
Mary Hall, University of Utah
Ewing Lusk, Argonne National Laboratory
Kengo Nakajima, University of Tokyo
Leonid Oliker, Lawrence Berkeley National Laboratory
Padma Raghavan, Vanderbilt University
Yves Robert, ENS Lyon
Valerie Taylor, Texas A&M University
Mateo Valero, Barcelona Supercomputing Center

Technical Papers

Technical Papers Chair
Lois Curfman McInnes, Argonne National Laboratory

Technical Papers Vice Chair
Richard Vuduc, Georgia Institute of Technology

Algorithms

Algorithms Chair
Anne Benoit, ENS Lyon

Algorithms Vice Chair
Julien Langou, University of Colorado, Denver

Algorithms Member
Emmanuel Agullo, French Institute for Research in Computer Science and Automation
Guillaume Aupy, Vanderbilt University
Grey Ballard, Sandia National Laboratories
Ioana Banicescu, Mississippi State University
Jed Brown, University of Colorado, Boulder
Aydin Buluc, Lawrence Berkeley National Laboratory
Erin Carson, New York University
Karen Devine, Sandia National Laboratories
Ana Gainaru, Mellanox Technologies
Murat Efe Guney, Intel Corporation
Thomas Herault, University of Tennessee
Judith C. Hill, Oak Ridge National Laboratory
Howie Huang, George Washington University
Kamer Kaya, Sabancı University
Hatem Ltaief, King Abdullah University of Science and Technology
Ulrike Meier Yang, Lawrence Livermore National Laboratory
Misun Min, Argonne National Laboratory
Sanjay Rajopadhye, Colorado State University
Amanda Randles, Duke University
Yves Robert, ENS Lyon
Krzysztof Rzadca, University of Warsaw
Ponnuswamy Sadayappan, Ohio State University
Erik Saule, University of North Carolina, Charlotte
Hari Sundar, University of Utah
Harsh Vardhan, Google
Frédéric Vivien, French Institute for Research in Computer Science and Automation
Rio Yokota, Tokyo Institute of Technology

Applications

Applications Chair
Carol Woodward, Lawrence Livermore National Laboratory

Applications Vice Chair
Michael Bader, Technical University Munich

Applications Member
H. Metin Aktulga, Michigan State University
Allison H. Baker, National Center for Atmospheric Research
Wolfgang Bangerth, Texas A&M University
George Biros, University of Texas at Austin
Xing Cai, Simula Research Laboratory
Donna Calhoun, Boise State University
Jean-Sylvain Camier, CEA-DAM
Anshu Dubey, Argonne National Laboratory
Pradeep Dubey, Intel Corporation
Joe Eaton, NVIDIA Corporation
Adrianna Gillman, Rice University
Gerard Gorman, Imperial College London

Boyce Griffith, University of North Carolina, Chapel Hill
Francois Gygi, University of California, Davis
David Ham, Imperial College London
Rich Hornung, Lawrence Livermore National Laboratory
Tobin Isaac, University of Texas at Austin
Kirk E. Jordan, IBM
Hector Klie, Sanchez Oil and Gas Corporation
Miriam Mehl, University of Stuttgart
Bronson Messer, Oak Ridge National Laboratory
Roger Pawlowski, Sandia National Laboratories
Noemi Petra, University of California, Merced
Daniel Reynolds, Southern Methodist University
Olaf Schenk, University of Lugano
Felix Schürmann, Swiss Federal Institute of Technology
Suzanne Shontz, University of Kansas
Masha Sosonkina, Old Dominion University
Alexandros Stamatakis, Heidelberg Institute for Theoretical Studies

Architectures & Networks

Architectures & Networks Chair
Mattan Erez, University of Texas at Austin

Architectures & Networks Vice Chair
Jung Ho Ahn, Seoul National University

Architectures & Networks Member
Ioana Baldini, IBM
Yungang Bao, Chinese Academy of Sciences
Christopher Batten, Cornell University
Jonathan Beard, ARM
Keren Bergman, Columbia University
Reetu Das, University of Michigan
Ronald Dreslinski Jr, University of Michigan
Paolo Faraboschi, Hewlett Packard Enterprise
Koji Inoue, Kyushu University
Samira Khan, University of Virginia
John Kim, Korea Advanced Institute of Science and Technology
Rakesh Kumar, University of Illinois
Jae W. Lee, Sungkyunkwan University
Moray McLaren, Google
Avi Mendelson, Technion - Israel Institute of Technology
George Michelogiannakis, Lawrence Berkeley National Laboratory
Vijay Nagarajan, University of Edinburgh
Scott Pakin, Los Alamos National Laboratory
Minsoo Rhu, NVIDIA Corporation
John Shalf, Lawrence Berkeley National Laboratory
Wei Wu, Intel Corporation
Jishen Zhao, University of California, Santa Clara

Clouds & Distributed Computing

Clouds & Distributed Computing Chair
Dick H.J. Epema, Delft University of Technology

Clouds & Distributed Computing Vice Chair
Lavanya Ramakrishnan, Lawrence Berkeley National Laboratory

Clouds & Distributed Computing Member
Ali R. Butt, Virginia Polytechnic Institute and State University
Kyle Chard, Argonne National Laboratory
Dean Hildebrand, IBM
Alexandru Iosup, Delft University of Technology
David Irwin, University of Massachusetts
Shantenu Jha, Rutgers University
Kate Keahey, Argonne National Laboratory
Kerstin Kleese van Dam, Brookhaven National Laboratory
Adrien Lèbre, French Institute for Research in Computer Science and Automation
Ruben S. Montero, Complutense University of Madrid
Christine Morin, French Institute for Research in Computer Science and Automation
Jelena Pjesivac-Grbovic, Google
Radu Prodan, University of Innsbruck
Ioan Raicu, Illinois Institute of Technology
Antony Rowstron, Microsoft Corporation
Evgenia Smirni, College of William and Mary
Rich Wolski, University of California, Santa Barbara
Yongwei Wu, Tsinghua University

Data Analytics, Visualization & Storage

Data Analytics, Visualization & Storage Chair
Yong Chen, Texas Tech University

Data Analytics, Visualization & Storage Vice Chair
Patricia Crossno, Sandia National Laboratories

Data Analytics, Visualization & Storage Member
Janine C. Bennett, Sandia National Laboratories
John Bent, EMC Corporation
Rita Borgo, King's College London
Suren Byna, Lawrence Berkeley National Laboratory
Toni Cortes, Barcelona Supercomputing Center
Kelly Gaither, University of Texas at Austin
Claire Guilbaud, CEA-DAM
Bingsheng He, Nanyang Technological University
Scott Klasky, Oak Ridge National Laboratory
Tevfik Kosar, University at Buffalo
Quincey Koziol, Lawrence Berkeley National Laboratory
Jens Krueger, University of Duisburg-Essen
Wei-keng Liao, Northwestern University
Xiaosong Ma, Qatar Computing Research Institute
Suzanne McIntosh, New York University
Kathryn Mohror, Lawrence Livermore National Laboratory
Xuanhua Shi, Huazhong University of Science and Technology
Daniela Ushizima, Lawrence Berkeley National Laboratory
Venkatram Vishwanath, Argonne National Laboratory
Jonathan L. Woodring, Lawrence Berkeley National Laboratory
Yinglong Xia, Huawei Research America
Jin Xiong, Institute of Computing Technology
Zheng Zhang, Rutgers University

Performance Measurement, Modeling, and Tools

Performance Chair
Gerhard Wellein, University of Erlangen-Nuremberg

Performance Vice Chair
Sameer Shende, University of Oregon

Performance Member
Richard Barrett, Sandia National Laboratories
Peter Beckman, Argonne National Laboratory
Jean-Baptiste Besnard, Paratools
Greg Bronevetsky, Google
Laura C. Carrington, University of California, San Diego
Bronis R. de Supinski, Lawrence Livermore National Laboratory
Rong Ge, Clemson University
Jeff R. Hammond, Intel Corporation
Heike Jagode, University of Tennessee
Tulin Kaman, University of Zurich
Andreas Knuepfer, Technical University Dresden
John Linford, Paratools
David Lowenthal, University of Arizona
Tareq Malas, Lawrence Berkeley National Laboratory
Allen Malony, University of Oregon
Frank Mueller, North Carolina State University
Catherine Olschanowsky, Colorado State University
Fabrizio Petrini, IBM
Vivek Sarkar, Rice University
Andrey Semin, Intel Corporation
Manu Shantharam, San Diego Supercomputer Center
Tyler Simon, US Department of Defense
Lauren L. Smith, National Security Agency
Nathan Tallent, Pacific Northwest National Laboratory
Jeffrey S. Vetter, Oak Ridge National Laboratory
Felix Wolf, Technical University Darmstadt
Michael Wolfe, NVIDIA Corporation

Programming Systems

Programming Systems Chair
Michelle Mills Strout, University of Arizona

Programming Systems Vice Chair
Ganesh Gopalakrishnan, University of Utah

Programming Systems Member
Scott B. Baden, University of California, San Diego
Michela Becchi, University of Missouri
Uday Bondhugula, Indian Institute of Science
James C. Brodman, Intel Corporation
Zoran Budimlic, Rice University
Aparna Chandramowlishwaran, University of California, Irvine
Albert Cohen, French Institute for Research in Computer Science and Automation
Stephen Fink, IBM
Hal Finkel, Argonne National Laboratory
Michael Garland, NVIDIA Corporation
Paul Hovland, Argonne National Laboratory
Costin Iancu, Lawrence Berkeley National Laboratory
Hironori Kasahara, Waseda University
Sriram Krishnamoorthy, Pacific Northwest National Laboratory
Ignacio Laguna, Lawrence Livermore National Laboratory
Jaejin Lee, Seoul National University
Maryam Mehri Dehnavi, Rutgers University
Boyana Norris, University of Oregon
Cindy Rubio-Gonzalez, University of California, Davis
Stephen F. Siegel, University of Delaware
Jeremy Siek, Indiana University
Didem Unat, Koc University
Qing Yi, University of Colorado, Colorado Springs
Tomofumi Yuki, French Institute for Research in Computer Science and Automation

State of the Practice

State of the Practice Chair
Costas Bekas, IBM

State of the Practice Vice Chair
Dirk Pleiter, Forschungszentrum Juelich

State of the Practice Member
Tina Butler, Lawrence Berkeley National Laboratory
Carlo Cavazzoni, CINECA
Susan Coghlan, Argonne National Laboratory
Winfield S. Futral, Lawrence Livermore National Laboratory
Rama Govindaraju, Google
Thomas Ludwig, German Climate Computing Center
Irene Qualters, National Science Foundation
James H. Rogers, Oak Ridge National Laboratory
Jennifer M. Schopf, Indiana University
Carsten Trinitis, Technical University Munich
Isabella Weger, European Centre for Medium-Range Weather Forecasts

System Software

System Software Chair
Christian Engelmann, Oak Ridge National Laboratory

System Software Vice Chair
Rolf Riesen, Intel Corporation

System Software Member
Sadaf R. Alam, Swiss National Supercomputing Center
Dorian C. Arnold, University of New Mexico
Rosa M. Badia, Barcelona Supercomputing Center
Brett Becker, University College Dublin
Zizhong Chen, University of California, Riverside
Andrew Chien, University of Chicago
Dilma Da Silva, Texas A&M University
Peter Dinda, Northwestern University
Ada Gavrilovska, Georgia Institute of Technology
Balazs Gerofi, RIKEN
Roberto Gioiosa, Pacific Northwest National Laboratory
Maya Gokhale, Lawrence Livermore National Laboratory
Hermann Haertig, Technical University Dresden
Trammell Hudson, Two Sigma
Kostas Katrinis, IBM
Ronald Minnich, Google
Dimitrios Nikolopoulos, Queen's University Belfast
Kevin Pedretti, Sandia National Laboratories
Eric Van Hensbergen, ARM

Tutorials

Tutorials Chair
Philip C. Roth, Oak Ridge National Laboratory

Tutorials Vice Chair
Naoya Maruyama, RIKEN

Tutorials Member
Katie Antypas, National Energy Research Scientific Computing Center
Ray S. Chen, University of Maryland
Hank Childs, University of Oregon
Lydia Duncan, Cray Inc.
Evan Felix, Pacific Northwest National Laboratory
Fernanda Foertter, Oak Ridge National Laboratory
Kevin Huck, University of Oregon
Takeshi Iwashita, Hokkaido University
Heike Jagode, University of Tennessee
Michael Kuhn, University of Hamburg
Stephen Leak, National Energy Research Scientific Computing Center
David Lecomber, Allinea
Graham Lopez, Oak Ridge National Laboratory
Robert F. Lucas, Information Sciences Institute
Maxime Martinasso, Swiss National Supercomputing Center
Marta Mattoso, Federal University of Rio de Janeiro
Kengo Nakajima, University of Tokyo
Suely Oliveira, University of Iowa
Andrey Ovsyannikov, Lawrence Berkeley National Laboratory
Srinivas Sridharan, Intel Corporation
Hiroyuki Takizawa, Tohoku University
Douglas Thain, University of Notre Dame
Didem Unat, Koc University
Jamie Van Randwyk, Lawrence Livermore National Laboratory
Ana Lucia Varbanescu, University of Amsterdam
Weichung Wang, National Taiwan University
Sandra Wienke, RWTH Aachen University
Paul Wilson, University of Wisconsin
Brian J. N. Wylie, Juelich Supercomputing Center
Rio Yokota, Tokyo Institute of Technology

Panels

Panels Chair
Michela Taufer, University of Delaware

Panels Vice Chair
Sandy Landsberg, US Department of Defense

Panels Member
David Abramson, University of Queensland
David H. Bailey, Lawrence Berkeley National Laboratory
Pavan Balaji, Argonne National Laboratory
George Biros, University of Texas at Austin
Aparna Chandramowlishwaran, University of California, Irvine
Jack Dongarra, University of Tennessee
Anne C. Elster, Norwegian University of Science and Technology
Bruce Hendrickson, Sandia National Laboratories
Torsten Hoefler, ETH Zurich
Seetharami Seelam, IBM
Katherine Yelick, Lawrence Berkeley National Laboratory

Workshops

Workshops Chair
Ewa Deelman, University of Southern California

Workshops Vice Chair
Maciej Malawski, AGH University of Science and Technology

Workshops Member
David Abramson, University of Queensland
Rosa M. Badia, Barcelona Supercomputing Center
Bartosz Balis, AGH University of Science and Technology
Claris Castillo, Renaissance Computing Institute
Almadena Chtchelkanova, National Science Foundation
Peter Coveney, University College London
Anne C. Elster, Norwegian University of Science and Technology
Trilce Estrada, University of New Mexico
Thomas Fahringer, University of Innsbruck
Soonwook Hwang, Korea Advanced Institute of Science and Technology
Ming Jiang, Lawrence Livermore National Laboratory
John Kim, Korea Advanced Institute of Science and Technology
Dalibor Klusáček, CESNET
Dieter Kranzlmueller, Ludwig Maximilian University of Munich
Valeria Krzhizhanovskaya, University of Amsterdam
Che-Rung Lee, National Tsing Hua University
Tim Mattson, Intel Corporation
Jarek Nabrzyski, University of Notre Dame
Ana-Maria Oprescu, University of Amsterdam
Christian Perez, French Institute for Research in Computer Science and Automation
Tom Peterka, Argonne National Laboratory
Jelena Pjesivac-Grbovic, Google
Lavanya Ramakrishnan, Lawrence Berkeley National Laboratory
Rizos Sakellariou, University of Manchester
Seetharami Seelam, IBM
Michela Taufer, University of Delaware
Kenjiro Taura, University of Tokyo

Posters

Posters Chair
Sadaf R. Alam, Swiss National Supercomputing Center

Posters Vice Chair
Chris J. Newburn, NVIDIA Corporation
Luke Olson, University of Illinois

Posters Member
Wonsun Ahn, University of Pittsburgh
Sadaf R. Alam, Swiss National Supercomputing Center
Damian Alvarez, Juelich Supercomputing Center
Pavan Balaji, Argonne National Laboratory
Muthu Baskaran, Reservoir Labs
James Beyer, NVIDIA Corporation
Paolo Bientinesi, Aachen University
R. Glenn Brook, University of Tennessee
Aydin Buluc, Lawrence Berkeley National Laboratory
Matteo Chesi, Swiss National Supercomputing Center
Florina M. Ciorba, University of Basel, Switzerland
Guillaume Colin de Verdiere, French Alternative Energies and Atomic Energy Commission

Doug Doerfler, Lawrence Berkeley National Laboratory
Anshu Dubey, Argonne National Laboratory
Joe Dummy, TryIt Inc.
Sameh Elnikety, Microsoft Corporation
Hal Finkel, Argonne National Laboratory
Fernanda Foertter, Oak Ridge National Laboratory
Victor Gergel, Nizhny Novgorod State University
Luc Giraud, French Institute for Research in Computer Science and Automation
Gerard Gorman, Imperial College London
Derek Groen, Brunel University
Bilel Hadri, King Abdullah University of Science and Technology
Andy Herdman, Atomic Weapons Establishment
Jan Hesthaven, EPFL
Judith C. Hill, Oak Ridge National Laboratory
Andreas Kloeckner, University of Illinois
Alice Koniges, Lawrence Berkeley National Laboratory
Jaejin Lee, Seoul National University
Olli-Pekka Lehto, CSC - IT Center for Science
Michael Lysaght, Irish Centre for HighEnd Computing
Maxime Martinasso, Swiss National Supercomputing Center
Simon McIntosh-Smith, University of Bristol
Kathryn Mohror, Lawrence Livermore National Laboratory
Diana Moise, Cray Inc.
Monique Morin, Los Alamos National Laboratory
Manoj Nambiar, Tata Consultancy Services
Stephen L. Olivier, Sandia National Laboratories
Luke Olson, University of Illinois
Ray Paden, Lenovo
Tapasya Patki, University of Arizona
Dirk Pleiter, Forschungszentrum Juelich
Fabrice Rastello, Inria
Carlos Rosales, University of Texas at Austin
Rahul Sampath, ExxonMobil Technical Computing Company
Ruchira Sasanka, Intel Corporation
Mitsuhisa Sato, RIKEN
Marie-Christine Sawley, Intel Corporation
Ivo F. Sbalzarini, Technical University Dresden
Matthieu Schaller, Durham University
Olaf Schenk, University of Lugano
Happy Sithole, Center for High Performance Computing in South Africa
Lauren L. Smith, National Security Agency
Thomas Steinke, Zuse Institute Berlin
Hari Sundar, University of Utah
Didem Unat, Koc University
Rob F. Van der Wijngaart, Intel Corporation

Birds of a Feather (BOFs)

Birds of a Feather Chair
Manish Parashar, Rutgers University

Birds of a Feather Vice Chair
Amy Apon, Clemson University

BOF Algorithms Chair
Daniel S. Katz, University of Illinois

BOF Applications Chair
Trilce Estrada, University of New Mexico

BOF Architecture and Networks Chair
Michela Becchi, University of Missouri

BOF Clouds & Distributed Computing Chair
Ivona Brandic, Vienna University of Technology

BOF Data Analytics, Visualization & Storage Chair
Hank Childs, University of Oregon

BOF Education Chair
David P. Bunde, Knox College

BOF Performance Measurement, Modeling and Tools Chair
Melissa C. Smith, Clemson University

BOF Programming Systems Chair
Omer Rana, Cardiff University

BOF State of the Practice Chair
Sharon Broude Geva, University of Michigan

BOF System Software Chair
Ron Brightwell, Sandia National Laboratories

Emerging Technologies

Emerging Technologies Chair
Wu Feng, Virginia Polytechnic Institute and State University

Emerging Technologies Vice Chair
Simon McIntosh-Smith, University of Bristol

Emerging Technologies Member
Pavan Balaji, Argonne National Laboratory
Ted Barragy, CGG
Taisuke Boku, University of Tsukuba
Franck Cappello, Argonne National Laboratory
Tom Deakin, University of Bristol
Wu Feng, Virginia Polytechnic Institute and State University
Kelly Gaither, University of Texas at Austin
Andy Herdman, Atomic Weapons Establishment
Torsten Hoefler, ETH Zurich
Yutong Lu, National University of Defense Technology
Filippo Mantovani, Barcelona Supercomputing Center
Simon McIntosh-Smith, University of Bristol
Indrani Paul, Advanced Micro Devices Inc
Amanda Randles, Duke University
Thomas Steinke, Zuse Institute Berlin
Neil Stringfellow, Pawsey Supercomputing Center
Michele Weiland, University of Edinburgh
Yunquan Zhang, Institute of Computing Technology

Scientific Visualization & Data Analytics Showcase

Scientific Visualization & Data Analytics Showcase Chair
Chris Johnson, University of Utah

Scientific Visualization & Data Analytics Showcase Vice Chair
Kristin Potter, University of Oregon

Scientific Visualization & Data Analytics Showcase Member
Utkarsh Ayachit, Kitware Inc
Rita Borgo, King's College London
David Camp, Lawrence Berkeley National Laboratory
Thierry Carrard, French Alternative Energies and Atomic Energy Commission
Amit Chourasia, University of California, San Diego
Kelly Gaither, University of Texas at Austin
Christoph Garth, University of Kaiserslautern
Ingrid Hotz,
Aaron Knoll, University of Utah
Heike Leitte,
Joshua A. Levine, University of Arizona
Liz Marai,
Paul A. Navratil, University of Texas at Austin
Tom Peterka, Argonne National Laboratory
David Pugmire, Oak Ridge National Laboratory
Paul Rosen, University of South Florida
Career Development

Doctoral Showcase

Doctoral Showcase Chair
Boyana Norris, University of Oregon

Doctoral Showcase Vice Chair
Manu Shantharam, San Diego Supercomputer Center

Doctoral Showcase Member
Guillaume Aupy, Vanderbilt University
Sunita Chandrasekaran, University of Delaware
Kamil Iskra, Argonne National Laboratory
Alice Koniges, Lawrence Berkeley National Laboratory
Kamesh Madduri, Pennsylvania State University
Boyana Norris, University of Oregon
Catherine Olschanowsky, Colorado State University
Shankar Sastry, University of Utah
Martin Schulz, Lawrence Livermore National Laboratory
Manu Shantharam, San Diego Supercomputer Center
Michelle Mills Strout, University of Arizona
Didem Unat, Koc University

Early Career Program

Early Career Program Chair
Jeffrey K. Hollingsworth, University of Maryland

Early Career Program Member
Dorian C. Arnold, University of New Mexico
Abhinav Bhatele, Lawrence Livermore National Laboratory
Sanjukta Bhowmick, University of Nebraska Omaha
Todd Gamblin, Lawrence Livermore National Laboratory
Karen Karavanic, Portland State University

Technical Archive

Proceedings
Karen Devine, Sandia National Laboratories

Conference Artifacts
Janet Brown, Carnegie Mellon University

SC17

Denver, CO | hpc connects.

Call for Participation

Denver, CO

Conference Dates:
Nov. 12 - 17, 2017

Exhibition Dates:
Nov. 13 - 16, 2017

http://sc17.supercomputing.org/

SC17: HPC Connects

Unlike anytime in history, HPC is allowing us to analyze mountains of data to prove long standing theories about our universe, our climate, our health, and our interests in many other research domains. Areas of the world now have access to new materials that help to slow down climate change... Coastal cities are evacuated before hurricanes claim lives...New precision medicine techniques bring new hope to our families and friends... Humanity is connected as never before through these discoveries. HPC brings together the computing systems and brilliant minds in this pursuit for a better world and our place in the cosmos.

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

Sponsors: acm sighpc IEEE computer society