PEARC19 Program

Saturday 6:30pm-7:30pm
Sunday 10:00am-5:00pm
Sunday 12:00pm-8:00pm
Sunday 1:00pm-5:00pm
Sunday 5:00pm-7:00pm
Sunday 9:00pm-11:59pm
Monday 7:00am-6:30pm
Monday 8:00am-5:00pm
Monday 8:30am-12:00pm
Monday 8:30am-5:00pm
Monday 10:00am-10:30am
Monday 10:00am-5:00pm
Monday 12:00pm-1:30pm
Monday 1:30pm-5:00pm
Monday 3:00pm-3:30pm
Monday 5:15pm-6:15pm
Monday 5:30pm-6:30pm
Monday 6:00pm-8:00pm
Monday 9:00pm-11:59pm
Tuesday 7:00am-8:00pm
Tuesday 8:00am-5:00pm
Tuesday 8:15am-10:30am
Tuesday 10:00am-5:00pm
Tuesday 10:30am-11:00am
Tuesday 11:00am-12:30pm
Tuesday 11:00am-5:00pm
Tuesday 12:30pm-1:30pm
Tuesday 1:30pm-3:00pm
Tuesday 1:30pm-5:00pm
Tuesday 2:00pm-4:00pm
Tuesday 3:00pm-3:30pm
Tuesday 3:30pm-5:00pm
Tuesday 4:00pm-6:00pm
Tuesday 4:30pm-6:30pm
Tuesday 5:15pm-6:15pm
Tuesday 6:30pm-8:30pm
Tuesday 9:00pm-11:59pm
Wednesday 7:00am-5:00pm
Wednesday 8:00am-3:00pm
Wednesday 8:00am-5:00pm
Wednesday 8:15am-10:30am
Wednesday 10:00am-5:00pm
Wednesday 10:30am-11:00am
Wednesday 10:45am-12:30pm
Wednesday 11:00am-12:30pm
Wednesday 12:30pm-1:30pm
Wednesday 1:30pm-3:00pm
Wednesday 1:30pm-3:30pm
Wednesday 1:30pm-5:00pm
Wednesday 3:00pm-3:30pm
Wednesday 3:30pm-5:00pm
Wednesday 4:00pm-5:00pm
Wednesday 5:00pm-10:00pm
Wednesday 6:30pm-10:00pm
Thursday 7:00am-12:00pm
Thursday 8:15am-10:00am
Thursday 10:00am-10:30am
Thursday 10:00am-12:00pm
Thursday 10:30am-12:00pm
Thursday 12:00pm-1:30pm
Thursday 1:30pm-3:00pm
Thursday 3:00pm-4:00pm

search

Saturday 6:30pm-7:30pm

Saturday 6:30pm-7:30pm
Hyatt Living Room
Special Event · Private, Closed Meeting, Special Event

Advanced Computing For Social Change Welcome and Networking Mixer (Participants Only)

Advanced Computing for Social Change Institute
Kelly Gaither and Rosalia Gomez (TACC), Ruby Mendenhall (University of Illinois at Urbana-Champaign), Linda Akli (SURA), and Sue Fratkin (Fratkin Associates)
▸ Biographies
▾ Abstract

Advanced Computing for Social Change is a competition for undergraduate students from diverse disciplines and backgrounds who want to work collaboratively to:

Learn to apply data analysis and computational thinking to a social challenge Experience the latest tools and techniques for exploring data through visualization Expand skills in team-based problem solving Learn how to communicate ideas more effectively to the general public

To be eligible for this exciting program, you must:

Participants will be selected via an application/selection process. To be eligible, students must be full time undergraduate students at an accredited college/university; Not plan to graduate before May 2020; Have a minimum overall GPA of at least 2.5/4.0 (or equivalent); Be able to attend the full challenge, July 27-August 1; Complete the online application form before the deadline.

The costs of airfare, lodging, meals, and conference registration will be provided.

Return to Top

Sunday 10:00am-5:00pm

Sunday 10:00am-5:00pm
Stetson F
Special Event · Private, Closed Meeting, Primary Room, Special Event

Advanced Computing for Social Change Institute

Kelly Gaither and Rosalia Gomez (TACC), Ruby Mendenhall (University of Illinois at Urbana-Champaign), Linda Akli (SURA), and Sue Fratkin (Fratkin Associates)
▸ Biographies
▾ Abstract

Advanced Computing for Social Change is a competition for undergraduate students from diverse disciplines and backgrounds who want to work collaboratively to:

Learn to apply data analysis and computational thinking to a social challenge Experience the latest tools and techniques for exploring data through visualization Expand skills in team-based problem solving Learn how to communicate ideas more effectively to the general public

To be eligible for this exciting program, you must:

Participants will be selected via an application/selection process. To be eligible, students must be full time undergraduate students at an accredited college/university; Not plan to graduate before May 2020; Have a minimum overall GPA of at least 2.5/4.0 (or equivalent); Be able to attend the full challenge, July 27-August 1; Complete the online application form before the deadline.

The costs of airfare, lodging, meals, and conference registration will be provided.

Return to Top

Sunday 12:00pm-8:00pm

Sunday 12:00pm-8:00pm
Regency Ballroom Registration Desk - Regency Foyer
Registration · Registration

Conference Registration and Information Desk

▾ Abstract

Conference Registration and Information Desk Hours: Sunday – 12:00pm to 8:00pm Monday – 7:00am to 6:30pm Tuesday – 7:00am to 8:00pm Wednesday – 7:00am to 5:00pm Thursday – 7:00am to 12:00pm

Return to Top

Sunday 1:00pm-5:00pm

Sunday 1:00pm-5:00pm
Water Tower
Co-located Meeting · Private, Closed Meeting, Co-located Meeting

1st Annual Clowder All Paws - Community Meeting

Kenton McHenry (University of Illinois at Urbana-Champaign)
▾ Abstract

Data sharing and publication are essential for modern scientific research, both in terms of leveraging past work towards novel discoveries as well as reproducibility. The proposed workshop will serve to bring together the community of users around Clowder, a novel cross domain open source data sharing, publication, and analysis framework based on the notion of Active Curation, an approach that significantly leverages machine learning and other analysis tools to automate much of the otherwise tedious curation process. A new $5 million NSF Cyberinfrastructure for Sustained Scientific Innovation (CSSI) award will work to bring together the grass roots cross disciplinary, international, community that has formed around Clowder towards ensuring its sustainability as a true open source framework, with distributed stakeholders and development activities addressing current scientific needs. PEARC offers an ideal venue to bring these very diverse cross disciplinary stakeholders together. The workshop will aim to kick-off an annual meeting, which may be at other venues in subsequent years, bringing together the current Clowder community to share developments, roadmaps, and support new users/communities.

Return to Top

Sunday 5:00pm-7:00pm

Sunday 5:00pm-7:00pm
Crystal A
Student Program · Student Program

Student Program Opening and Welcome

▾ Abstract

Bringing together students and chairs of the various student programs to meet and greet. There will be a variety of speakers from students, committee, and invited guests. This gathering will allow the students to ask questions about the conference and gain insight into how they may best maximize their experience at this conference. Appetizers and refreshments will be served.

Sunday 5:00pm-7:00pm
Addams
Co-located Meeting · Private, Closed Meeting, Co-located Meeting

Campus Champions Leadership Team Meeting

Marisa Brazil (Purdue University) and Dana Brunson (Internet2)
▾ Abstract

This is a closed strategic planning session for the Campus Champions Leadership Team and Staff.

Return to Top

Sunday 9:00pm-11:59pm

Sunday 9:00pm-11:59pm
Field
Special Event · Hackathon

SGCI Hack@PEARC19

SGCI Hackathon (sciencegateways.org)
▾ Abstract

Join us for three days of science gateway hacking with some of the best programmers in the field!

In conjunction with PEARC19 in Chicago, SGCI is sponsoring a Hackathon to challenge up-and-coming programmers for tomorrow's science gateways. Undergraduate and graduate students will compete on gateway projects established by SGCI partners and community members.

Join us on the social media app SLACK to get the latest updates! Follow us on Twitter at @sciencegateways The competition will take place from Sunday, July 28, to Wednesday, July 31, in the Hyatt Regency Chicago where the PEARC19 conference is being held. Final awards will be announced on Thursday, August 1.

Return to Top

Monday 7:00am-6:30pm

Monday 7:00am-6:30pm
Regency Ballroom Registration Desk - Regency Foyer
Registration · Registration

Conference Registration and Information Desk

▾ Abstract

Conference Registration and Information Desk Hours: Sunday – 12:00pm to 8:00pm Monday – 7:00am to 6:30pm Tuesday – 7:00am to 8:00pm Wednesday – 7:00am to 5:00pm Thursday – 7:00am to 12:00pm

Return to Top

Monday 8:00am-5:00pm

Monday 8:00am-5:00pm
Stetson F
Special Event · Private, Closed Meeting, Primary Room, Special Event

Advanced Computing for Social Change Institute

Kelly Gaither and Rosalia Gomez (TACC), Ruby Mendenhall (University of Illinois at Urbana-Champaign), Linda Akli (SURA), and Sue Fratkin (Fratkin Associates)
▸ Biographies
▾ Abstract

Advanced Computing for Social Change is a competition for undergraduate students from diverse disciplines and backgrounds who want to work collaboratively to:

Learn to apply data analysis and computational thinking to a social challenge Experience the latest tools and techniques for exploring data through visualization Expand skills in team-based problem solving Learn how to communicate ideas more effectively to the general public

To be eligible for this exciting program, you must:

Participants will be selected via an application/selection process. To be eligible, students must be full time undergraduate students at an accredited college/university; Not plan to graduate before May 2020; Have a minimum overall GPA of at least 2.5/4.0 (or equivalent); Be able to attend the full challenge, July 27-August 1; Complete the online application form before the deadline.

The costs of airfare, lodging, meals, and conference registration will be provided.

Return to Top

Monday 8:30am-12:00pm

Monday 8:30am-12:00pm
Atlanta
Tutorial · Tutorial Half-day

How to Accelerate Your Big Data and Associated Deep Learning Applications with Hadoop and Spark?

Dhabaleswar Panda and Xiaoyi Lu (The Ohio State University) and Mahidhar Tatineni (San Diego Supercomputer Center)
▾ Abstract

Apache Hadoop and Spark are gaining prominence in handling Big Data analytics. Recent studies have shown that default Hadoop and Spark can not leverage the high-performance networking and storage architectures efficiently, like Remote Direct Memory Access (RDMA) enabled high-performance interconnects and heterogeneous storage systems (e.g. HDD, SSD, NVMe-SSD, and Lustre). These middleware are traditionally written with sockets and do not deliver the best performance on modern high-performance networks. In this tutorial, we will provide an in-depth overview of the architecture of Hadoop components (HDFS, MapReduce, etc.) and Spark. We will examine the challenges in re-designing networking and I/O components of these middleware with modern interconnects, protocols (such as InfiniBand and RoCE) with RDMA and storage architectures. Using the publicly available software packages in the High-Performance Big Data (HiBD, http://hibd.cse.ohio-state.edu) project, we will provide case studies of the new designs for several Hadoop/Spark components and their associated benefits. Through these case studies, we will also examine the interplay between high-performance interconnects, high-speed storage systems, and multi-core platforms to achieve the best solutions for these components, Big Data processing, and Deep Learning applications on modern HPC clusters. This tutorial will provide hands-on sessions of Hadoop and Spark on SDSC Comet supercomputer.

http://web.cse.ohio-state.edu/~panda.2/pearc19_bigdata_tut.html

pdf

Monday 8:30am-12:00pm
Crystal A
Tutorial · Tutorial Half-day

Introduction to Python for Scientific Computing

Chris Myers (Cornell University, Center for Advanced Computing); Katharine Cahill (Ohio Supercomputing Center); and Aaron Weeden (Shodor)
▾ Abstract

This half-day tutorial is a quick immersion in the basics of the Python programming language and associated packages for scientific computing, including tools needed to participate in the Student Modeling Challenge, part of the PEARC19 student program. Topics covered will include key language features – such as variables, types, operators, control flow, input/output, functions, classes, built-in containers, and modules – as well as an overview of some of the important libraries and packages in the Python ecosystem for scientific computing (to support plotting, data access and manipulation, numerical algorithms, and the construction of integrated computational notebooks). Most techniques will be presented in live-demo mode, and each section will feature hands-on exercises so that participants can try out the commands or methods for themselves. To participate fully in the exercises, attendees should come with the latest version of the Anaconda Python 3 package downloaded and installed on their computer (https://www.anaconda.com/download).

pdf

Monday 8:30am-12:00pm
Crystal C
Tutorial · Tutorial Half-day

High Performance Distributed Deep Learning: A Beginner's Guide

Dhabaleswar K. Panda, Ammar Ahmad Awan, and Hari Subramoni (The Ohio State University)
▾ Abstract

The current wave of advances in Deep Learning (DL) has led to many exciting challenges and opportunities for Computer Science and Artificial Intelligence researchers alike. DL frameworks like TensorFlow, PyTorch, Caffe, and several others have emerged that offer ease of use and flexibility to describe, train, and deploy various types of Deep Neural Networks (DNNs). In this tutorial, we will provide an overview of interesting trends in DNN design and how cutting-edge hardware architectures are playing a key role in moving the field forward. We will also present an overview of different DNN architectures and DL frameworks. Most DL frameworks started with a single-node/single-GPU design. However, approaches to parallelize the process of DNN training are also being actively explored. The DL community has moved along different distributed training designs that exploit communication runtimes like gRPC, MPI, and NCCL. In this context, we highlight new challenges and opportunities for communication runtimes to efficiently support distributed DNN training. We also highlight some of our co-design efforts to utilize CUDA-Aware MPI for large-scale DNN training on modern GPU clusters. Finally, we also include hands-on exercises to enable the attendees gain first-hand experience of running distributed DNN training experiments on a modern GPU cluster.

http://web.cse.ohio-state.edu/~panda.2/pearc19_dl_tut.html

pdf

Monday 8:30am-12:00pm
Water Tower
Tutorial · Tutorial Half-day

Modern tools for supercomputers

Robert McLay, Lei Huang, and Si Liu (Texas Advanced Computing Center)
▾ Abstract

Powerful supercomputers have played an important role in the computational research community. However, the increasing complexity on modern systems may defer or hinder their work. A large amount of precious time and effort have been spent unnecessarily in managing the user environment, reproducing standard workflow, handling large scale I/O work, profiling and monitoring users jobs, understanding and resolving unnecessary system issues, etc. To help supercomputer users focus on their scientific and technical work and to minimize the workload for the consulting team, we designed and developed a series of powerful tools for supercomputer users. These tools are portable and effective on almost all supercomputers and are now serving thousands of supercomputer users of TACC, XSEDE, and other institutions every day.

In this tutorial, we will present and practice with supercomputer tools specifically designed for complex user environment (LMod, Sanity Tool), tools for job monitoring and profiling (Remora, TACC-Stat, etc.), tools for large-scale IO work (OOOPS, FanStore), and several other convenient tools. Attendees will learn how these tools are designed and how to use them in their daily work. Detailed hands-on exercises are prepared beforehand and will be executed mainly on the Stampede2 supercomputer at the Texas Advanced Computing Center (TACC).

pdf

Monday 8:30am-11:30am
Gold Coast
Tutorial · Tutorial Half-day

The Higher Education Campus Alliance for Advanced Visualization Tutorial and Workshop

The Higher Education Campus Alliance for Advanced Visualization Tutorial and Workshop
Emma Jane Alexander and Kyle Summerfield (University of Wyoming)
▾ Abstract

The Higher Education Campus Alliance for Advanced Visualization (THE CAAV) is presenting three speakers at PEARC19. The intention here is to address a wide range of topics, to appeal to broad range of PEARC attendees. The event will be held the morning of Monday, July 29, 2019.

Stereo 360 photography for Virtual Reality
Kyle D. Summerfield (University of Wyoming)
▾ Abstract

Title: Stereo 360 photography for Virtual Reality Duration: 1.5 hours Format: Presentation and tutorial (Participants can bring own laptops to follow along with the Unity tutorial section) Requirements: https://store.unity.com/ Free person download of Unity by participants

This presentation will focus on data capture methods (LiDAR, photogrammetry, and 360 stereo photography) and their integration into Unity for use in virtual reality. I will provide an overview of these three approaches, followed by specific focus on 360 capture and implementation. The workshop will walk through capture, stitching, VR implementation, and simple interactions with minimal scripting in C#. Emphasis will be placed on the overall workflow and project structure. At the conclusion of the prepared presentation, I will field any questions that arise, suggest ways that the interactions could be expanded, and propose use cases for the technologies that currently exist as well as those still on the horizon.

Next-Generation Cybersecurity Visualization and Analysis
Jan Springer (University of Arkansas)
▾ Abstract

Title: Next-Generation Cybersecurity Visualization and Analysis Duration: 45 minutes Format: Presentation

Over the past two to three decades visualization has been embraced in many fields. Not only that but this embrace actually includes (and prefers) advanced and, in many cases, highly interactive visualization scenarios of non-trivial and complex situations in engineering, medicine, and many other real-world applications. However, there is one area where this trend and its advances are not to be found: cybersecurity operations and its related fields.

Today, the demand for cybersecurity operators (e.g., data analysts, network engineers, thread hunters) is rapidly increasing together with ever larger amounts of data to be analyzed for patterns, threads, etc. On the other hand, people working in this area are restricted in the kind of tools they are able or allowed to use. While this is understandable in truly sensitive domains, the majority of information cybersecurity operators are dealing with is usually public non-sensitive data. But even then the restrictions are being kept and rather simplistic tools are used for very complex analysis problems.

In this talk we will explore potentially fruitful avenues for using large-data analysis and visualization strategies in the context of cybersecurity operations. The talk is based on initial work in the Department of Computer Science and the Emerging Analytics Center at the University of Arkansas at Little Rock.

Research Amplification: Two Visualization Use Cases
Shelby Hallman (James B. Hunt Jr. Library)
▾ Abstract

Title: Research Amplification: Two Visualization Use Cases Duration: 45 minutes Format: Presentation

In this session, Shelby Hallman will talk about the benefits of using publicly accessible, large scale visualization walls and immersive spaces to showcase data-driven research. In particular, two models of showcasing visualization will be discussed through the following use cases: using residencies to connect data and art, and providing an ongoing speaker series to highlight visualization based research. Attendees will come away with strategies and tips for utilizing visualization to amplify research.

Monday 8:30am-12:00pm
Columbian
Tutorial · Tutorial Half-day

Deep Dive into Microsoft's Azure for Research

T. J. Barry (Applied Cloud Systems) and Eric DeBord (Microsoft Corporation)
▾ Abstract

Whether it's a computer with more memory, a cluster with thousands of cores, a big data platform, an internet of things solution, or open-source machine learning at scale, you can achieve more using the cloud. Microsoft Azure provides an open, flexible, global platform that supports multiple programming languages, tools, and frameworks. This workshop will take us into the Azure offerings available to researchers as well as explore a few case studies on how these offerings have impacted academic research. This will be followed by two of Deployment of Microsoft's HPC Pack, and then How to use Azure CycleCloud Create, manage, operate and optimize HPC and big compute clusters of any scale.

Monday 8:30am-12:00pm
New Orleans
Workshop · Workshop Half-day

Leveraging a Research IT Maturity Model for Strategic Decision Making

Patrick Schmitz (University of California, Berkeley); Gail Krovitz and Dana Brunson (Internet2); Thomas Cheatham (University of Utah); Alex Feltus and Jill Gemmill (Clemson University); Galen Collier (Rutgers University); John Hicks (Internet2); Claire Mizumoto (University of California, San Diego); and Karen Wetzel (Educause)
▾ Abstract

Research IT (computing, data, and related infrastructure and services) is changing at an accelerating rate, while the range of scientific fields and disciplines depending on research cyberinfrastructure is expanding and becoming increasingly diverse. We present a new Maturity Model for Research IT (computing, data, and related infrastructure and services) that identifies the range of relevant approaches to supporting research IT, for use by IT practitioners, researchers, and campus leadership. Participants will apply and evaluate the model for their institutions, provide feedback to improve the model, and discuss its use for strategic decision making.

pdf

Monday 8:30am-12:00pm
McCormick
Workshop · Workshop Half-day

The ACM SIGHPC SYSPROS Symposium 2019

Jenett Tillotson and Stephen Harrell (Purdue University) and John Blaas (University of Colorado)
▾ Abstract

In order to meet the demands of high-performance computing (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 high performance computing 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 to 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. Additionally, this Workshop is affiliated with the systems professionals' chapter of the ACM SIGHPC (SIGHPC SYSPROS Virtual ACM Chapter). This session would serve as an opportunity for chapter members to meet face-to-face, discuss the chapter's yearly workshop held at SC, and continue building our community's shared knowledge base.

pdf

Monday 8:30am-12:00pm
Picasso
Workshop · Workshop Half-day

Pragmatic Science Engagement Using an Operations Center Approach

Jason Zurawski (Energy Sciences Network (ESnet), Lawrence Berkeley National Laboratory) and Jennifer Schopf (Indiana University)
▾ Abstract

Establishing a process to regularly review technical requirements (http://www.es.net/science-engagement/science-requirements-reviews/) helps to determine the current and future science communication and collaboration needs for numerous scientific communities. The purpose of these reviews is to accurately characterize the near-term, medium-term and long-term technical requirements of the science being performed. This approach brings about a network-centric understanding of the science process used by the researchers and scientists, without asking technical questions directly, and derives network requirements from that understanding.

This highly interactive session outlines a process that can be adopted by CI facilitators at the campus, regional, and national levels to advance the mission of science engagement and fully realize the investments made in networking and personnel by agencies such as the National Science Foundation and Department of Energy Office of Science.

The topics of this event focus on the intersection of research and technology, specifically use of high speed networks. Specifically:

- The purpose of science engagement (30 mins) - Overview of an outline for a case study approach to gather scientific requirements via documentation and meeting preparation (45 mins) - A live example of how to conduct an in-person review to characterize needs with a visiting scientific group (1 hour) - Participant discussion on CI facilitator techniques – sharing of challenges and lessons learned (45 mins)

Participants in National Science Foundation Campus Cyberinfrastructure Program and the cyberinfrastructure engineering community are encouraged to attend, participate, and help define a strategy that can be used to encourage growth of scientific understanding and support.

pdf

Return to Top

Monday 8:30am-5:00pm

Monday 8:30am-5:00pm
Field
Tutorial · Tutorial Full-day

Programming Modern Multicore Processors

W. Cyrus Proctor, Virginia Trueheart, Ian Wang, and John Cazes (Texas Advanced Computing Center)
▾ Abstract

Modern processors such as Intel's Scalable Xeon line, AMD's EPYC architecture, ARM's ThunderX2 design, and IBM's Power9 architecture are scaling out rather than up and increasing in complexity. Because the base frequencies for the large core count chips hover somewhere between 2-3 GHz, researchers can no longer rely on frequency scaling to increase the performance of their applications. Instead, developers must learn to take advantage of the increasing core count per processor and learn how to eke out more performance per core.

To achieve good performance on modern processors, developers must write code amenable to vectorization, be aware of memory access patterns to optimize cache usage, and understand how to balance multi-process programming (MPI) with multi-threaded programming (OpenMP). This tutorial will cover the basics of vectorization, multi-threaded programming, memory affinity, load balancing, and hybrid execution. We will also provide an overview of current HPC processors and profiling tools.

This session will include hands-on exercises that demonstrate the techniques discussed and usage of profiling tools. This tutorial is designed for experienced programmers, familiar with OpenMP and MPI, who wish to learn how to program for performance on modern architectures.

pdf

Monday 8:30am-5:00pm
Burnham
Tutorial · Tutorial Full-day

Cloud-based Virtual Clusters using Jetstream

Sanjana Sudarshan, Stephen Bird, Eric Coulter, and Jeremy Fischer (Indiana University)
▾ Abstract

Cloud computing has grown at a significant rate over the past fewyears. While cloud computing alleviates much of the hardwaremanagement challenge, many researchers and educators have difficulty embracing the aspects of cloud computing that make it unique and preferred for many activities. Among the many features cloud computing brings, elastic computing - resources on demand, maybe one of the most appealing. A purpose-built virtual machine canbe built quickly and easily as a first step. Taking the example further into scripted launching of more resources as necessary is the logical next step. The next step can use elastic computing techniques to create virtual clusters on demand, bringing compute resources into existence when needed and removing them when they are no longer necessary. This elasticity enables gateway providers and researchers to make efficient use of limited resources, while providing a resource for HPC-style jobs that isn't dependent on external HPC resources. While modest cloud-based virtual clusters won't replace traditional HPC resources for jobs that require high-speed interconnects or large core counts and high memory profiles, many smaller gateway, research, and education projects might benefit from the highly customizable and configurable, programmable cyberinfrastructure afforded by cloud computing environments. This tutorial will explore the basic methods required for interacting with elastic computing environments. It will then show a hands-on approachbto creating virtual clusters in an OpenStack environment, including the necessary steps to make the cluster elastic, to take full advantage of the cloud environment.

pdf

Monday 8:30am-5:00pm
Wrigley
Tutorial · Tutorial Full-day

Portable, Reproducible High Performance Computing In the Cloud

Joe Stubbs (TACC), Sean Cleveland (The University of Hawaii), and Julia Looney and Anagha Jamthe (TACC)
▾ Abstract

This tutorial will focus on providing attendees exposure to state-of-the-art techniques for portable, reproducible research computing, enabling them to easily transport analyses from cloud to HPC resources. We will introduce open source technologies such as Jupyter, Docker and Singularity, the emerging "serverless" computing paradigm, and how to utilize these tools within two NSF-funded cyberinfrastucture platforms, Agave and Abaco API. The approaches introduced not only increase application portability and reproducibility but also reduce or eliminate the need for investigators to maintain physical infrastructure so that more time can be spent on analysis. For the tutorial, attendees will have access to allocations on XSEDE JetStream and one or more HPC resources such as TACC's Stampede2 or Frontera.

pdf

Monday 8:30am-5:00pm
Addams
Workshop · Workshop Full-day

AI4GOOD@PEARC19: Workshop Proposal

Elizabeth Leake (STEM-Trek Nonprofit, University of Iowa) and Alana Romanella (Virginia Tech)
▾ Abstract

AI4GOOD involves a full-day of awareness, advocacy and hands-on training in basic skills needed by those who wish to employ or support artificial intelligence (AI) for accelerated research outcomes in a variety of domains. Biomedical advances, economic empowerment strategies, agricultural innovation and quality of life improvements for citizens in underserved regions will be emphasized. The day concludes with lightening talks and a panel discussion about privacy, policies, security, and ethics. Cybersecurity, IoT and legal experts will answer the question: How does AI impact the landscape for inter-institutional and international data stewardship and research compliance? Submitted by: Elizabeth Leake (STEM-Trek/U-Iowa) and Alana Romanella (Virginia Tech)

pdf

Monday 8:30am-5:00pm
Haymarket
Workshop · Workshop Full-day

CyberAmbassadors: Communications, Teamwork, Ethics and Leadership Training for Cyber-Infrastructure Professionals

Dirk Colbry and Kathleen Colbry (Michigan State University)
▾ Abstract

This workshop provides communication, teamwork and leadership training for technically-proficient CI-Professionals, with the goal of developing "CyberAmbassadors" who are prepared to expand broader engagement in multidisciplinary, computationally-intensive research. The curriculum uses interactive exercises and small-group activities to help participants build professional skills within the context of multidisciplinary computational research.

pdf

Monday 8:30am-5:00pm
Soldier Field
Workshop · Workshop Full-day

Humans in the Loop: Enabling and Facilitating Research on Cloud Computing

Craig Stewart and Brian Voss (Pervasive Technology Institute, Indiana University, Bloomington)
▾ Abstract

This workshop explores the role of humans in making cloud computing useful in research settings. Cloud computing is clearly a type of cyberinfrastructure, which, in its general definition, "consists of computing systems, data storage systems, advanced instruments and data repositories, visualization environments, and people, all linked together by software and high performance networks to improve research productivity and enable breakthroughs not otherwise possible." In this workshop we will focus on the "and people" part of cyberinfrastructure, and in particular on the role of people in supporting the use of commercial cloud resources in research.

pdf

Return to Top

Monday 10:00am-10:30am

Monday 10:00am-10:30am
Regency Foyer/Concourse Level Foyer/Crystal Foyer
Break · Break

Morning Break

Return to Top

Monday 10:00am-5:00pm

Monday 10:00am-5:00pm
Toronto/Hong Kong/Acapulco
Exhibition · Exhibition

Exhibition

Joseph Paris
▾ Abstract

The PEARC conference series works to integrate and meet the collective interests of our growing community by providing a forum for discussing challenges, opportunities, and solutions among the broad range of participants in the research computing community. The exhibitors program provides an opportunity to advance exciting new collaborations while shaping solutions through engaging with technology and service providers across both industry and academia.

Return to Top

Monday 12:00pm-1:30pm

Monday 12:00pm-1:30pm
Comiskey
Co-located Meeting · Private, Closed Meeting, Co-located Meeting

CC* Team: Great Plains Regional CyberTeam Project Meeting

Daniel Andresen (Kansas State University)

Monday 12:00pm-1:30pm
Regency CD
Lunch · Lunch

Lunch

Return to Top

Monday 1:30pm-5:00pm

Monday 1:30pm-5:00pm
New Orleans
Tutorial · Tutorial Half-day

Taking Advantage of Persistent Memory with Open Source PMDK and Intel Profiling Tools

Andy Rudoff (Intel Corporation)
▸ Biography
▾ Abstract

The recent emergence of persistent memory offers a new tier of data placement, with memory-style access and storage-style persistence. This half-day workshop covers how operating systems expose persistent memory to applications, and ways applications can fully leverage what it can do. Full program examples showing solutions using persistent memory are presented. The workshop covers the use of the Persistent Memory Development Kit, PMDK, a growing collection of libraries which have been developed for various use cases, tuned, validated to production quality, and thoroughly documented. These libraries build on the Direct Access (DAX) feature available in both Linux and Windows, which allows applications direct load/store access to persistent memory by memory-mapping files on a persistent memory aware file system. PMDK (http://pmem.io/pmdk/) will work with any persistent memory that provides the SNIA NVM Programming Model (https://www.snia.org/sites/default/files/technical_work/final/NVMProgrammingModel_v1.2.pdf). It is open source and welcomes community contributions.

In addition to PMDK, this educational session will show users how to take advantage of freely available tools from Intel – namely the Intel® VTune™ Amplifier. This premier performance profiler has new capabilities to help you optimize your persistent memory programs. This workshop will show users how to:

- Analyze systems over longer intervals. Find out which workloads can benefit from larger memory allocations and which system configuration better fits the workloads.
- Locate code that is sensitive to memory bandwidth and latency issues. Identify hot, warm, or cool data to optimize memory usage and placement.
- Identify opportunities to replace disk or SSD-based storage with faster persistent memory.

The workshop will also show users how to take advantage of the new Intel® Persistence Inspector tool. This tool finds persistence errors quickly to make software fast and reliable.

Monday 1:30pm-5:00pm
Crystal C
Tutorial · Tutorial Half-day

Building Container Images for HPC Workloads

Scott McMillan (NVIDIA)
▾ Abstract

This tutorial will describe how to build container images for HPC workloads with Docker, Singularity, and HPC Container Maker (HPCCM). Best practices including image layering, caching, multi-stage builds, and managing container image size will be covered. HPCCM will be a primary focus, describing how to use this open source tool to generate Dockerfiles and Singularity definition files from a Python recipe.

pdf

Monday 1:30pm-5:00pm
McCormick
Tutorial · Tutorial Half-day

Using the SPEC HPG Benchmarks for Better Analysis and Evaluation of Current and Future HPC Systems

Robert Henschel and Junjie Li (Indiana University) and Sunita Chandrasekaran and Mayara Gimenes (University of Delaware)
▾ Abstract

The High Performance Group (HPG) of the Standard Performance Evaluation Corporation (SPEC) develops bench- mark methodologies for High Performance Computing systems and releases production quality benchmark suites like SPEC MPI2007, SPEC OMP2012, and SPEC ACCEL. These benchmarks can evaluate all dimensions of parallelism on modern HPC systems and are used in academia and industry to conduct research in HPC systems, facilitate pro- curement, testing, and tuning of HPC systems. Since 2018, SPEC HPG offers these benchmark suites free of charge to non-commercial users. In this half-day tutorial, participants will learn how to leverage SPEC benchmarks for per- formance evaluation, tuning of system parameters, comparison of systems (e.g., for procurement), and get an outlook on its power measurement capabilities. The presenters will provide demos and hands-on guidance on how to install, compile, and run the benchmarks on HPC systems provided. The presenters will also show how the results will be in- terpreted, discuss various use cases, and present the publication process of results. An SSH-capable device is required for the hands-on sessions.

pdf

Monday 1:30pm-5:00pm
Water Tower
Tutorial · Tutorial Half-day

Practical OpenHPC: Cluster Management, HPC Applications, Containers and Cloud

Karl Schulz (University of Texas); Reese Baird (OpenHPC); Derek Simmel (Pittsburgh Supercomputing Center, Carnegie Mellon University); Eric Coulter (Indiana University); Chris Downing (Red Oak Consulting); and Nirmala Sundararajan (Dell)
▾ Abstract

Since its inception as a Linux Foundation project in 2015, OpenHPC (https:///openhpc.community) has steadily grown to provide a modern, consistent, reference collection of HPC cluster provisioning tools, together with a curated repository of common cluster management software, I/O clients, advanced computational science libraries and software development tools, container-based execution facilities, and application performance profiling tools.

Although OpenHPC enables people deploying new HPC clusters to rapidly get their clusters up and running, the OpenHPC software repository itself is a reliable, portable, integrated collection of software, libraries, tools and user environment that can be employed in containers, VMs and HPC clusters to develop and execute computational science applications.

This half-day tutorial will begin with a brief, advanced introduction to OpenHPC. We will then guide attendees through several practical, hands-on exercise modules employing an OpenHPC-based cluster and the OpenHPC software repository to explore real-world activities including: - HPC cluster management and job schedulers - Using containers to build and prototype HPC applications with OpenHPC - Running container-based applications on HPC clusters - Using Easybuild and Spack to streamline application builds - Instrumenting applications for performance measurement - Using Packer to build and use OpenHPC-ready VM images in the cloud

pdf

Monday 1:30pm-5:00pm
Gold Coast
Tutorial · Tutorial Half-day

Tools and Best Practices for Distributed Deep Learning with High Performance Computing

Weijia Xu and Zhao Zhang (Texas Advanced Computing, University of Texas) and David Walling (Texas Advanced Computing, University of Texas at Austin)
▾ Abstract

This tutorial is a practical guide on how to run distributed deep learning over distributed compute nodes effectively. Deep Learning (DL) has emerged as an effective analysis method and are adapted quickly across many scientific domains in recent years. Domain scientists are embracing DL as both a standalone data science method, as well as an effective approach to reducing dimensionality in the traditional simulation. However. due to its inherent high computational requirement, application of DL is limited by the available computational resources. Recently, we have seen the fusion of DL and high-performance computing (HPC): supercomputers show an unparalleled capacity to reduce DL training time from days to minutes; HPC techniques have been used to speed up parallel DL training. Therefore distributed deep learning has great potential to augment DL applications by leveraging existing high performance computing cluster. This tutorial consists of three sessions. First, we will give an overview of the state-of-art approaches to enabling deep learning at scale. The second session is an interactive hands-on session to help attendees running distributed deep learning with resources at the Texas Advanced Computing Center. In the last session, we will focus on the best practices to evaluate and tune up performance.

pdf

Monday 1:30pm-5:00pm
Picasso
Tutorial · Tutorial Half-day

Developing Science Gateways using Apache Airavata

Marlon Pierce, Suresh Marru, Eroma Abeysinghe, Sudhakar Pamidighantam, Marcus Christie, and Dimuthu Wannipurage (Indiana University)
▾ Abstract

This half-day tutorial will build on XSEDE14, XSEDE15, XSEDE16, and PEARC17 tutorials. Extensive tutorial material for this tutorial is available from https://cwiki.apache.org/confluence/display/AIRAVATA/PEARC+2019.

Science gateways provide science-specific user interfaces for scientific applications for end users who are unfamiliar with or need more capabilities than provided by command-line interfaces. Science gateway middleware provides the general purpose capabilities behind gateway user interfaces. In this tutorial, we present the Apache Airavata middleware for creating science gateways. The target audiences for this tutorial include a) scientific software developers, who want simplified ways to deliver their software and support larger user communities; b) educators, who want to integrate scientific software usage into their classroom without having students get bogged down in the submission mechanisms of specific resources; and c) campus computing center staff, who want to use gateways to broaden their reach beyond their traditional users and to help users make more efficient use of resources.

The format will be a mixture of presentations, demonstrations, and hands-on sessions, as described in the detailed agenda. The agenda indicates the roles of each of the organizers.

pdf

Monday 1:30pm-5:00pm
Columbian
Tutorial · Tutorial Half-day

HPC & ML on Google Cloud Platform

Wyatt Gorman (Google.com)
▾ Abstract

Introduction to HPC on Google Cloud

Customer Story: Paul Sagona @ USC - Predicting Climate Change by Sequencing Microbiomes

Hands-on HPC: Slurm Auto-Scaling Clusters on Google Cloud

Introduction to Machine Learning for HPC Engineers

Hands-on Machine Learning: Introduction to Kubeflow on Google Cloud

Monday 1:30pm-5:00pm
Atlanta
Workshop · Workshop Half-day

Second Workshop on Strategies for Enhancing HPC Education and Training (SEHET19)

Nitin Sukhija (Slippery Rock University of Pennsylvania), Scott Lathrop (University of Illinois at Urbana-Champaign), Julia Mullen (MIT Lincoln Laboratory), Kate Cahil (Ohio Supercomputer Center), Susan Mehringer (Cornell University), Dana Brunson (Oklahoma State University), and Lauren Milechin (Lincoln Laboratory Supercomputer Center (MIT))
▾ Abstract

High performance computing is becoming central for empowering scientific progress in the most fundamental research in various science and engineering, as well as society domains. It is remarkable to observe that the recent rapid advancement in the mainstream computing technology has facilitated the ability to solve complex large-scale scientific applications that perform advanced simulations of the implementation of various numerical models corresponding to numerous complex phenomena pertaining to diverse scientific fields. The inherent wide distribution, heterogeneity, and dynamism of the today's and future computing and software environments provide both challenges and opportunities for cyberinfrastructure facilitators, trainers and educators to develop, deliver, support, and prepare a diverse community of students and professionals for careers that utilize high performance computing to advance discovery.

The SEHET19 workshop is an ACM SIGHPC Education Chapter coordinated effort aimed at fostering collaborations among the practitioners from traditional and emerging fields to explore strategies to enhance computational, data-enabled and HPC educational needs. Attendees will discuss approaches for developing and deploying HPC training, as well as identifying new challenges and opportunities for keeping pace with the rapid pace of technological advances - from collaborative and online learning tools to new HPC platforms. The workshop will provide opportunities for: learning about methods for conducting effective HPC education and training; promoting collaborations among HPC educators, trainers and users; and for disseminating resources, materials, lessons learned and good/best practices.

pdf

Return to Top

Monday 3:00pm-3:30pm

Monday 3:00pm-3:30pm
Regency Foyer/Concourse Level Foyer/Crystal Foyer
Break · Break

Afternoon Break

Return to Top

Monday 5:15pm-6:15pm

Monday 5:15pm-6:15pm
Atlanta
Birds of a Feather · Birds of a Feather, Technical Tools

Ask.CI, the Q&A Site for Research Computing - Year 1 Lessons Learned, Plans for Year 2

Julie Ma (Massachussetts Green High Performance Computing Center, Northeast Cyberteam Initiative); Aaron Culich (Berkeley); Torey Battelle (Colorado School of Mines); Zoe Braiterman (Owasp Foundation); Dana Brunson (Internet2, Campus Champions); Thomas Cheatham III (University of Utah); John Goodhue (Massachusetts Green High Performance Computing Center); Christopher Hill (Massachusetts Institute of Technology); Katia Oleinik and Jacob Pessin (Boston University); Raminder Singh (Harvard University); Jack Smith (West Virginia Higher Education Policy Commission); Vanessa Sochat (Stanford University); and Scott Yockel (Harvard University)
▾ Abstract

In September, 2017, the Northeast Cyberteam began a project to build a Q&A site to optimize dissemination of knowledge for research computing. The site was launched at PEARC18, and a naming contest at the conference yielded Ask.Cyberinfrastructure.org (Ask.CI).

The goal of Ask.CI is to aggregate answers to a broad spectrum of questions that are commonly asked by the research computing community. As researchers, staff, students, and other users ask and answer questions on Ask.CI, they are creating a shared knowledge base for the larger community, and alleviating the burden of a central resource to provide support to a user base. For smaller institutions, this provides a wealth of knowledge that was not previously available in an easily searchable Q&A form. For larger institutions, this novel, self-service model frees up time for facilitators and cyberinfrastructure engineers to focus on more advanced subject matter, thereby elevating the practice.

Establishing a Q&A site of this nature requires some tenacity. In partnership with the Campus Champions, Ask.CI has gained traction, and continues to engage the broader community to firmly establish the platform as a powerful tool for research computing. In the year since launch, Ask.CI has attracted nearly 100,000 page views, hundreds of contributors, and a broad audience that spans the U.S., and parts of Europe and Asia.

At this BOF, the Ask.CI team will discuss strategies, key accomplishments, challenges, and plans for the upcoming year.

pdf

Monday 5:15pm-6:15pm
New Orleans
Birds of a Feather · Birds of a Feather, Building Community

Building a Community of Research Software Engineers

Sandra Gesing (University of Notre Dame, Purdue University); Ian Cosden (Princeton University); and Charles Ferenbaugh (Los Alamos National Laboratory)

Abstract

Software has become a major driver for research with over 90% of researchers answering surveys that they use software for their research and over 65% expressing that they even could not do their research without software. Good software can last forever, but when novices develop it, problems can arise. Often research code will become unusable at the end of the life of a project or with the departure of a single individual. The profession of Research Software Engineers (RSEs) has become internationally recognized as an increasingly critical role in creating and sustaining reliable software. We define Research Software Engineers to encompass those who regularly use expertise in programming to advance research. The goal of this BoF is to discuss the next steps for outreach and actions feasible for building/extending a US RSE community. Is maybe the FORCE11 and Research Data Alliance (RDA) model with working groups around specific topics a good way to go? Which national specifics should be considered to be as successful as the UK example? PEARC'19 is an excellent fit as venue for this BoF since it is one of the major events for the campus champions, ACI-REF facilitators and advanced research computing. Research software is a large sub-group of research computing. The report by the Campus Research Computing Consortium (CaRCC) on cyber-infrastructure professionalization emphasizes different roles in the eco-system of cyber-practitioners and the software/data facing role is one of them.

pdf

Monday 5:15pm-6:15pm
Crystal C
Birds of a Feather · Birds of a Feather, Technical Tools

ColdFront Resource Allocation Management System

Andrew Bruno and Robert DeLeon (Center for Computational Research SUNY at Buffalo), Thomas Furlani (Rosewell Park Cancer Institute), and Dori Sajdak and Mohammad Zia (Center for Computational Research SUNY at Buffalo)

Abstract

ColdFront is an open source HPC center resource allocation and management system recently developed and released by SUNY at Buffalo Center for Computational Research (CCR). ColdFront was developed to provide an easy way for HPC center staff to manage shared resources and for users to request access to those resources. It was first announced and released in conjunction with Super Computing 18 last November. In a very short period there has been a great deal of interest as judged by the number of views and clones on github and positive feedback from early adopters. Hence, there is a newly developing user community. This PEARC 19 BoF is an ideal venue to both bring together and expand this community with new users who will be exposed to ColdFront for the first time.

pdf

Monday 5:15pm-6:15pm
Water Tower
Birds of a Feather · Birds of a Feather, Technical Tools

The effective use of Jupyter notebooks by HPC centers

Carrie Ganote (National Center for Genome Analysis, Support; Indiana University); Phil Blood (PSC); Bhavya Papudeshi and Thomas Doak (National Center for Genome Analysis, Support; Indiana University); Bryan Learn (PSC); and Sheri Sanders (National Center for Genome Analysis, Support; Indiana University)

Abstract

This BOF aims to create a space to share ideas and discussion on the use of Jupyter notebooks for teaching and sharing scientific workflows on HPC systems. Major topics include experiences, pain points, success stories, and thoughts on implementation.

pdf

Monday 5:15pm-6:15pm
Gold Coast
Birds of a Feather · Birds of a Feather, Data

The Open Storage Network: Toward a National Distributed Storage Infrastructure

Melissa Cragin (San Diego Supercomputer Center, UCSD); John Goodhue (Massachusetts Green High Performance Computing Center); Jay Aikat (Renaissance Computing Institute / UNC); and Kenton McHenry (NCSA / UIUC)

Abstract

The Open Storage Network (OSN) project demonstrates a distributed, scalable storage infrastructure designed to address storage and access needs for active-use data sets that require managed, longer-term storage and/or transfer across institutions. The OSN offers scalable services with shared support and governance, and addresses cyberinfrastructure needs that are difficult for academic institutions to support individually. The purpose of this BoF session is to present the project to the broader PEARC community and gather feedback on design and future growth. The OSN needs to partner with Research Computing professionals on first and last mile solutions for scientific applications, and for future expansion of the network of storage hosting organizations.

pdf

Monday 5:15pm-6:15pm
Haymarket
Birds of a Feather · Birds of a Feather, Cybersecurity

Securing Research CI

Susan Sons (Indiana University)

Abstract

Join us for a roundtable discussion about the challenges of securing the infrastructure science depends on, lessons learned, and paths forward. Susan Sons and Mary Conley from ResearchSOC, the NSF-funded security operations center for scientific cyberinfrastructure, will facilitate the discussion. CISOs, IT professionals, researchers, and interested community members are welcome.

Scientific cyberinfrastructure has more to contend with than just commodity IT. We have sensor nets, drones, research vehicles, network-connected lab equipment, supercomputers, telescopes, and other control-system based or one-off technologies. These are often distributed across multiple organizations and networks. This BoF session provides an opportunity to connect with others who are taking on similar challenges, both technical and programmatic, in order to support science becoming ever more reproducible, trustworthy, and productive.

Monday 5:15pm-6:15pm
Columbian
Birds of a Feather · Birds of a Feather, User Needs

Service Provider Perspectives On User Needs

Lee Liming (University of Chicago) and John-Paul Navarro (Argonne National Laboratory)

Abstract

Representatives from XSEDE Service Providers (SPs) will share their perspectives on current and future user needs relative to the XSEDE cyber-infrastructure. The goal is to ensure that SP knowledge of current and future user needs is being effectively communicated and shared with XSEDE partners.

pdf

Monday 5:15pm-6:15pm
Wrigley
Birds of a Feather · Birds of a Feather, Technical Tools

Campus Research Clouds: Gathering Requirements for a Campus Cloud Toolkit

Richard Knepper (cornell Advanced Computing Center, XSEDE); Resa Reynolds (Cornell University); and Eric Coulter and Steve Bird (Indiana University)

Abstract

Research computing staff at a number of universities are finding that cloud technologies, traditionally the area of private industry, are useful in providing flexible computational resources that meet a wide variety of research computing needs. Research clouds provide different capabilities than traditional HPC resources, and require different strategies to implement, manage, and extend. Research clouds also have the oft-touted advantage that users can take their research to other clouds or burst in between interoperable clouds. This capability means that significant amounts of cloud computing capacity are available by leveraging and sharing resources.

The proposed Birds-of-a-Feather session will bring together campus cyberinfrastructure providers, cyberinfrastructure facilitators, and computational users in order to gather requirements for campus research clouds capable of providing scalable research services that meet a variety of needs.

Members of the NSF-funded Aristotle Project and Jetstream Research Cloud will facilitate a discussion on requirements for campuses installing a campus level research cloud, including hardware and software choices, implementation choices, developing usage policies and support resources, and sustainability plans. This BOF is meant to provide a venue for a range of different needs and roles for the campus research cloud to meet. The requirements that surface from this BOF will be used to determine components and documentation for an XSEDE Cyberinfrastructure Resource Integration toolkit to be made broadly available to the cyberinfrastructure community via the Research Software Portal.

pdf

Return to Top

Monday 5:30pm-6:30pm

Monday 5:30pm-6:30pm
Crystal A
Co-located Meeting · Private, Closed Meeting, Co-located Meeting

CASC - Membership meeting and reception

Lisa Arafune (CASC - Coalition for Academic Scientific Computation)

Biography

Abstract

CASC Membership meeting & reception at PEARC19; similar to PEARC18 CASC meeting & reception, July 24, 2018.

Return to Top

Monday 6:00pm-8:00pm

Monday 6:00pm-8:00pm
Regency D
Student Program · Student Program, Private, Closed Meeting

Student - Mentor Dinner

Ester Soriano (NCSA)

Biography

Abstract

The purpose of the Mentor Program is to connect the mentee to the PEARC Conference, provide the mentee with networking opportunities during the conference, and allow the mentee to learn more about Practice & Experience in Advanced Research Computing. The conference mentor's primary role will be to help guide his/her student through the conference and to share his/her personal experiences in the field of Advanced Research Computing.

The student-mentor dinner is a great opportunity to start this friendship journey for both: mentor and mentee. This event has always has been a great success in previous PEARC conferences: 2017 and 2018

Monday 6:00pm-8:00pm
Crystal Foyer
Special Event · Private, Closed Meeting, Special Event

Champion Networking Event

Marisa Brazil (Purdue University) and Dana Brunson (Internet2)

Abstract

The Champions evening networking event is an opportunity for members of the community to come together and network with other Champions. This activity is beneficial to building professional relationships and social networks outside of individual geographical areas. Time at the beginning of this activity will be utilized for administrative announcements.

Return to Top

Monday 9:00pm-11:59pm

Monday 9:00pm-11:59pm
Field
Special Event · Hackathon

SGCI Hack@PEARC19

SGCI Hackathon (sciencegateways.org)

Abstract

Join us for three days of science gateway hacking with some of the best programmers in the field!

In conjunction with PEARC19 in Chicago, SGCI is sponsoring a Hackathon to challenge up-and-coming programmers for tomorrow's science gateways. Undergraduate and graduate students will compete on gateway projects established by SGCI partners and community members.

Join us on the social media app SLACK to get the latest updates! Follow us on Twitter at @sciencegateways The competition will take place from Sunday, July 28, to Wednesday, July 31, in the Hyatt Regency Chicago where the PEARC19 conference is being held. Final awards will be announced on Thursday, August 1.

Return to Top

Tuesday 7:00am-8:00pm

Tuesday 7:00am-8:00pm
Regency Ballroom Registration Desk - Regency Foyer
Registration · Registration

Conference Registration and Information Desk

Abstract

Conference Registration and Information Desk Hours: Sunday – 12:00pm to 8:00pm Monday – 7:00am to 6:30pm Tuesday – 7:00am to 8:00pm Wednesday – 7:00am to 5:00pm Thursday – 7:00am to 12:00pm

Return to Top

Tuesday 8:00am-5:00pm

Tuesday 8:00am-5:00pm
Stetson F
Special Event · Private, Closed Meeting, Primary Room, Special Event

Advanced Computing for Social Change Institute

Kelly Gaither and Rosalia Gomez (TACC), Ruby Mendenhall (University of Illinois at Urbana-Champaign), Linda Akli (SURA), and Sue Fratkin (Fratkin Associates)

Biographies

Abstract

Advanced Computing for Social Change is a competition for undergraduate students from diverse disciplines and backgrounds who want to work collaboratively to:

Learn to apply data analysis and computational thinking to a social challenge Experience the latest tools and techniques for exploring data through visualization Expand skills in team-based problem solving Learn how to communicate ideas more effectively to the general public

To be eligible for this exciting program, you must:

Participants will be selected via an application/selection process. To be eligible, students must be full time undergraduate students at an accredited college/university; Not plan to graduate before May 2020; Have a minimum overall GPA of at least 2.5/4.0 (or equivalent); Be able to attend the full challenge, July 27-August 1; Complete the online application form before the deadline.

The costs of airfare, lodging, meals, and conference registration will be provided.

Tuesday 8:00am-5:00pm
Crystal B
Posters/Visualization · Poster, Student Poster, Vis Showcase

Posters and Visualization

Abstract

All Exhibitors, Posters, and Visualization Showcase entries will be available at this time, for viewing at your convenience.

Return to Top

Tuesday 8:15am-10:30am

Tuesday 8:15am-10:25am
Regency AB
Plenary Session · Plenary

Plenary Session A

Chair: Thomas Furlani (Roswell Park Comprehensive Cancer Center)

Welcome to PEARC19
Tom Furlani (Roswell Park Comprehensive Cancer Center)

AI for Science
Rick Stevens (University of Chicago, Argonne National Laboratory)

Biography

Abstract

In this talk, I will describe an emerging initiative at Argonne National Laboratory to advance the concept of Artificial Intelligence (AI) aimed at addressing challenge problems in science. We call this initiative "AI for Science". The basic concept is threefold: (1) to identify those scientific problems where existing AI and machine learning methods can have an immediate impact (and organize teams and efforts to realize that impact); (2) identify areas of where new AI methods are needed to meet the unique needs of science research (frame the problems, develop test cases, and outline work needed to make progress); and (3) to develop the means to automate scientific experiments, observations, and data generation to accelerate the overall scientific enterprise. Science offers plenty of hard problems to motivate and drive AI research, from complex multimodal data analysis, to integration of symbolic and data intensive methods, to coupling large-scale simulation and machine learning to drive improved training to control and accelerate simulations. A major sub-theme is the idea of working toward the automation of scientific discovery through integration of machine learning (active learning and reinforcement learning) with simulation and automated high-throughput experimental laboratories. I will provide some examples of projects underway and lay out a set of long-term driver problems.

Evolving and sustaining the ecosystem of research information technology
Dana Brunson (Internet2), Thomas Cheatham (University of Utah), Sharon Broude Geva (University of Michigan), Doug Jennewein (University of South Dakota), and Tobin Magle (University of Wisconsin - Madison)

Abstract

The needs of research and researchers on campuses is growing rapidly beyond the desktop, and also beyond the traditional high performance computing (HPC) support campus research IT groups have long provided. This includes deployment of virtual machines and their customization, adoption and support of container and orchestration technologies, support for big data, data movement, machine learning and analytics, help with migration to the cloud, and facilitating the use of accelerators, such as GPUs. Most campus research IT groups are experiencing a significant rise in the number of research groups and users requiring access to research computing and data resources. Sustaining this growth, coping with new and less experienced communities, navigating and satisfying the ever-increasing set of compliance requirements, and facilitating the use of emerging and innovative technologies, is putting a strain on campus research computing people and resources. A potential approach towards sustaining efforts—as the technologies continue to innovate and diversify—is to collaborate and share expertise, experience, leading practices, and resources among institutions. Rather than duplicate efforts on each campus, maybe we can work together to support research and researchers and/or provide better means to advocate for, support, and improve research IT on our campuses.

Realizing the needs, numerous communities have emerged at the local, regional and national levels, each with different foci, to help support research computing and data and the associated research IT professionals on our campuses. Examples include the Campus Champions (CC), the Coalition for Academic Scientific Computation (CASC), and the Campus Research Computing Consortium (CaRCC), among numerous others. Each of these work or function in different ways, but collectively they all are looking inward to define their vision and aims, at models for sustainability, and each are exploring ways to empower and help our communities.

In April 2019, a workshop was held to bring together an initial subset of participants from some of these communities to better understand the broader ecosystem of research IT. This included representatives from CaRCC, CASC, CC, Coalition for Network Information (CNI), HPC sys-pros, Big Data Hubs, Open Science Grid (OSG), Carpentries, Women in HPC, EDUCASE, Engagement and Performance Operations Center (EPOC), CI Engineers, Midscale Experimental Research Infrastructures Forum (MERIF), Global Environment for Network Investigations (GENI), SIGHPC education, HPC University, Research Data Access and Preservation Association (RDAP), Great Plains Network (GPN), The Quilt, Minority Serving Institutions / Historically Black Colleges, and the Extreme Science and Engineering Discovery Environment (XSEDE).

A key goal of the workshop was to identify key organizations and initiatives in the research computing and data ecosystem, including those not included in the initial workshop (for example Research Software Engineers, many of whom were attending a different meeting at the same time), and to identify commonalities, gaps, and areas for potential collaboration and to provide a more coherent view of the broader research computing and data ecosystem highlighting successes and challenges.

The proposed PEARC19 1.5 hour panel is to have some representatives from the Research Computing and Data Ecosystem workshop present to discuss the findings of the workshop and to discuss future directions with the larger community of interested PEARC19 participants. Moderated by Dana Brunson and with panelists from CaRCC (Thomas Cheatham) and CASC (Sharon Broude Geva), Campus Champions (Doug Jennewein), and RDAP (Tobin Magle). The panel presentations will include dissemination of the collective vision and next steps for defining and strengthening the broader ecosystem of research computing and data, with ample time put aside for discussions with the PEARC19 community.

pdf

Return to Top

Tuesday 10:00am-5:00pm

Tuesday 10:00am-5:00pm
Toronto/Hong Kong/Acapulco
Exhibition · Exhibition

Exhibition

Joseph Paris

Abstract

The PEARC conference series works to integrate and meet the collective interests of our growing community by providing a forum for discussing challenges, opportunities, and solutions among the broad range of participants in the research computing community. The exhibitors program provides an opportunity to advance exciting new collaborations while shaping solutions through engaging with technology and service providers across both industry and academia.

Return to Top

Tuesday 10:30am-11:00am

Tuesday 10:30am-11:00am
Regency Foyer/Concourse Level Foyer/Crystal Foyer
Break · Break

Morning Break

Return to Top

Tuesday 11:00am-12:30pm

Tuesday 11:00am-12:30pm
Crystal A
Workshop · Workshop Quarter-day

Python or R or Both: Tools for Your Data Analytics Workflow

Kang-Pyo Lee and Samuel Burer (University of Iowa)

Abstract

Python and R are among the most popular programming languages for data analysis. While they share many of the same important features as high-level, open-source programming languages, they each have their own strengths and weaknesses from a data-science perspective. While many researchers starting their data analysis journey feel they must choose between Python and R as their primary analytics tool, the reality is that there currently are no clear guidelines to make a decision between the two. This workshop aims to provide an in-depth comparison between Python and R for data science, thereby sparking an active discussion about which one would be the best choice given a specific goal and circumstance—or even when using both might be an appropriate choice. This workshop will benefit researchers, professionals, and students who need to choose between the two languages but also those who need to learn and use both.

pdf

Tuesday 11:00am-12:30pm
Regency AB
Paper · Facilitation, Storage, Technical Paper

Storage/Archival

Chair: Ritu Arora (TACC, University of Texas at Austin)

IceCube's Long Term Archive Software
Patrick Meade, Benedikt Riedel, and David Schultz (WIPAC)

Abstract

IceCube is a cubic kilometer neutrino detector located at the South Pole. It generates 1 TiB of raw data per day, which must be archived for possible retrieval years or decades later. Other low-level data products are also archived for easy retrieval in the event of a catastrophic data center failure. The Long Term Archive software is IceCube's answer to archiving this data across several computing sites.

pdf

Provisioning ZFS Pools On Lustre
Rick Mohr and Adam Howard (University of Tennessee)

Abstract

While Lustre's parallelism, performance, and scalability make it desirable as a storage solution for clusters, its limitations prevent it from being suitable as a general purpose storage for all of a cluster's needs. In particular, Lustre's relatively poor performance for small, random I/O as well as metadata-intensive workloads make it less suitable for use as storage for users' home areas or working directories for compilation. In this paper, an experiment of deploying a ZFS file system using files pre-allocated on a Lustre file system as the ZFS storage pools is presented. While this adds many management options like snapshots and more fine grained controls over users and quotas, we focus on examining how adding this layer can affect the performance of suboptimal Lustre workloads. Benchmarking of the Lustre file system with and without the ZFS file system is performed, and the results are presented and analyzed.

pdf

StashCache: A Distributed Caching Federation for the Open Science Grid
Derek Weitzel and Marian Zvada (University of Nebraska - Lincoln), Ilija Vukotic and Rob Gardner (University of Chicago), Brian Bockelman (Morgridge Institute for Research), Mats Rynge (Information Sciences Institute), Edgar Fajardo Hernandez (San Diego Super Computer Center), Brian Lin (University of Wisconsin - Madison), and Matyas Selmeci (University of Wisconsin -- Madison)

Abstract

Data distribution for opportunistic users is challenging as they neither own the computing resources they are using or any nearby storage. Users are motivated to use opportunistic computing to expand their data processing capacity, but they require storage and fast networking to distribute data to that processing. Since it requires significant management overhead, it is rare for resource providers to allow opportunistic access to storage. Additionally, in order to use opportunistic storage at several distributed sites, users assume the responsibility to maintain their data.

In this paper we present StashCache, a distributed caching federation that enables opportunistic users to utilize nearby opportunistic storage. StashCache is comprised of four components: data origins, redirectors, caches, and clients. StashCache has been deployed in the Open Science Grid for several years and has been used by many projects. Caches are deployed in geographically distributed locations across the U.S. and Europe. We will present the architecture of StashCache, as well as utilization information of the infrastructure. We will also present performance analysis comparing distributed HTTP Proxies vs StashCache.

pdf

Tuesday 11:00am-12:30pm
Crystal C
Paper · Technical Paper, Student Technical Paper, Software, Applications

Scientific Software Management

Chair: Lev Gorenstein (Purdue University)

Community Collections: A Framework for Openly Sharing Software Stacks Across Research Computing Centers Using Singularity and Lmod
Kevin Manalo (Johns Hopkins University), Levi Baber (Iowa State University), Ryan Bradley (Johns Hopkins University), Zhi-Qiang You (Ohio Supercomputer Center), and Nuyun Zhang (Georgia Institute of Technology)

Abstract

Automatic software build infrastructure has advanced significantly with software build tools such as EasyBuild and Spack. Similarly, scientific container tools such as Singularity have greatly eased the burden of maintaining individual software components at research computing centers. Typically, most research computing sites use tools such as Environment Modules or Lmod to present entire software hierarchies to users. These tools are designed to better serve individual sites or specific toolchains. Our aim is to complement these preexisting build tools by inverting the concept of site-specific software collections to serve multiple institutions. This shared framework relies on both Lmod and Singularity to allow individual sites to seamlessly deploy predefined collections of software components which can be further tailored to their needs.

Johns Hopkins University, Iowa State University, Ohio Supercomputer Center, and Georgia Tech have deployed an extendable framework based on Singularity and Lmod to enable "Community Collections", a software framework that allows multiple institutions to customize the deployment of Singularity software images as a module. This deployment allows for multiple institutions to trivially extend their existing Lmod based software stacks and participate in a greater effort to reduce software deployment redundancy and to allow for re-duplication of software stacks in cloud environments. These efforts are centered around deployment that allows users and systems administrators to leverage shared Singularity containers. The tool facilitates the sharing of module files and their interface to corresponding containers consisting of single-node, shared memory, GPU/CPU-based software.

pdf

Singularity: simple, secure containers for compute-driven workloads
David Godlove (Sylabs Inc) and Randy Herban (Sylabs, Inc.)

Abstract

Container technology makes it easy to create highly portable and reproducible research environments. The Singularity container platform has a unique security model allowing untrusted users to run untrusted containers safely on multi-tenant systems. Singularity has a special file format (called the Singularity Image Format or SIF) for packaging and distributing containers. This format allows for novel features like cryptographic signing and verification of containers, extreme portability, and guaranteed reproducibility. The Singularity runtime supports both integration as well as isolation, making it easy to read and write data to the host system, leverage hardware like GPUs and high-speed interconnects, and integrate with batch schedulers like Slurm, PBS, LSF, UGE, etc. This emphasis on integration is part of the overall development philosophy: Singularity is intended to be a simple, secure, feature-rich container solution. Sylabs Inc. has developed SingularityPRO, a professionally curated and supported version of the open source offering as well as the Singularity Container Services, an end-to-end cloud hosted solution for creating and distributing trusted containers. Most top HPC centers use Singularity in production, and enterprise users are rapidly adopting this solution as well.

pdf

Advanced Research Computing Containers for AI Workloads
Kyle Smathers, Marvin Newlin, and Mark DeYoung (Air Force Institute of Technology)

Abstract

Containerization has taken the software world by storm. Deployment complications, like requiring elevated (i.e. "root") permissions to run, have slowed the adoption of containers in shared advanced research computing (ARC) environments. Singularity is a containerization approach that is designed for ARC in shared high performance computing (HPC) clusters. With the creation of the Singularity, there is finally a viable scientific container solution. However very few papers have looked at the performance tradeoffs of deploying applications using a container based model. The authors are not aware of any published studies evaluating the tradeoffs of the deployment models with complex Artificial Intelligence (AI) workloads. Without detailed evaluations of the performance trade-offs scientists and engineers are unable to make an informed decision on deployment model for time sensitive training or low power inference. Furthering previous research in this area and using emerging community developed benchmarks, we examine performance trade-offs of running AI workloads in a containerized Singularity environment.

pdf

Tuesday 11:00am-12:30pm
Wrigley
Paper · Technical Paper, Student Technical Paper, Workforce, Honorable Mention Student Paper Workforce

Training/Scientific Programming Education

Chair: Asya Shklyar (Pomona College)

Integrating scientific programming in communities of practice for students in life science
Honorable Mention Student Paper Workforce
Alexa Salsbury, Anne Brown, and Justin Lemkul (Virginia Tech)

Abstract

Research in life science domains is producing larger data sets that require the use of computational approaches to understand biological phenomena. Academic institutions, industry, and other sectors in the life sciences are creating jobs that involve computation, data science, and data visualization. Therefore, there is a need for life scientists to understand and be trained in computation for this new job market. Many life science students are not taught foundational concepts of computation as a part of their curriculum. Therefore, there exists a gap in understanding when beginning to learn computer science (CS) and relate it to data-centric questions in other fields. To improve learning experiences and help train these students, this work sought to understand existing challenges that life science students face in learning scientific programming and identify routes for improvement. To do so, we evaluated three distinct learning experiences— a hands-on workshop, structured coursework, and long-term research experiences. Based on these student evaluations, we highlight the major challenges and benefits of different learning environments and provide suggestions to educators and institutions for integrating scientific programming education in life science coursework or research. Student-centered, group environments were the most successful at engaging students in computing concepts. Overall, this work provides strategies to enrich learning experiences and promote best practices in computation for life science students and engage these students in the development of in-demand skills.

pdf

Campus Champions: Building and sustaining a thriving community of practice around research computing and data
Marisa Brazil (Purdue University); Dana Brunson (Internet2); Aaron Culich (University of California, Berkeley); Lizanne DeStefano (Georgia Institute of Technology); Douglas Jennewein (University of South Dakota); Tiffany Jolley (University of Illinois Urbana-Champaign); Timothy Middelkoop (University of Missouri); Henry Neeman (University of Oklahoma); Lorna Rivera (Georgia Institute of Technology); Jack Smith (Marshall University); Julie Wernert (Indiana University Bloomington); and Lizanne DeStefano (Georgia Institute of Technology)

Abstract

The digital research landscape has changed dramatically over the years, as campuses across the nation have gained access to local research computing resources and services. The Campus Champions program, [1] founded in 2008, has also evolved with these changes to accommodate and support the diversity and growth of the research computing community that we have long supported. Our mission, to promote and facilitate the effective participation of a diverse national community of institutions in the application of advanced digital resources and services to accelerate scientific discovery and scholarly achievement, has been made possible by sustained funding from the National Science Foundation (NSF) via the eXtreme Science and Engineering Discovery Environment (XSEDE) [2] over the past ten years, and has enabled the Campus Champions program to support campus cyberinfrastructure and to foster a thriving community of practice [3] with nationwide impact.

To facilitate the continued growth and sustainability of the Campus Champions program beyond XSEDE, and to provide a deeper understanding of the Champion culture and needs, the Champion Sustainability Working Group and the XSEDE Evaluation team conducted a climate study in 2018. The recommendations provided by the climate study will inform our Leadership Team and staff on how we may further our community outreach goals and plan for the future of the program. This paper will highlight and discuss the development, implementation, key findings, and recommendations from that study.

pdf

Holistic Summer Undergraduate Research Program in High Performance Computing Research
Tyler Miller, Jung-Han Kimn, and Stephen Gent (South Dakota State University)

Abstract

This study investigated the effectiveness of a National Science Foundation-sponsored Research Experience for Undergraduates Site in High Performance Computing (HPC) at South Dakota State University. In the 2018 program, ten first- and second-year students majoring in science, technology, engineering and math (STEM) disciplines were engaged in an intensive ten-week program on the use and application of HPC in theoretical and applied research. The participants were engaged in an interdisciplinary program and were mentored by faculty and graduate students as they completed their project. Participants were provided instruction on the use of HPC, preparing posters and presentations, the Responsible Conduct of Research, and a variety of interpersonal skills. Feedback on the program was evaluated with daily journal entries prepared by the participants, Computer & Information Science & Engineering (CISE) pre- and post-survey responses, one-on-one interviews, and the Torrance Tests of Creative Thinking. Results of the 2018 cohort indicate that participants' attitudes toward their project discipline (e.g., engineering or biological sciences) are favorable. Participants' self-reported Research Skills also increased from the beginning to the end of the program. However, results suggest that their favorable attitudes decreased from the beginning to the end of the program. One hypothesis of why this occurred is that participants had unrealistic ideas about the time, effort, or duties of high-level research in math, engineering, and/or biological science disciplines. After the ten-week program, participants' views were more closely aligned with the reality of research work in the area.

pdf

Tuesday 11:00am-12:30pm
Gold Coast
Paper · Technical Paper, Machine Learning/Artificial Intelligence, ML/AI Applications, Gold Exhibitor Technical Talk, Best Paper Machine Learning

Application I

Chair: George Rudolph (Utah Valley University)

A Framework for Design Identification on Heritage Objects
Best Paper Machine Learning
Jun Zhou and Yuhang Lu (University of South Carolina); Karen Smith (S.C. Department of Natural Resources); and Colin Wilder, Song Wang, Paul Sagona, and Ben Torkian (University of South Carolina)

Abstract

A challenging problem in modern archaeology is to automatically identify fragmented heritage objects by their decorative full designs, such as the pottery sherds from Southeastern America. The difficulties of this problem lie in: 1) these pottery sherds are usually fragmented so that each sherd only covers a small portion of its underlying full design; 2) these sherds can be so highly degraded that curves may contain missing segments or become very shallow; and 3) curve patterns may overlap with each other from the making of these potteries. This paper presents a deep-learning based framework for matching a sherd with a database of known designs to find its underlying design. This framework contains three steps: 1) extracting curve pattern using an FCN-based curve pattern segmentation method from the digitized sherd's depth map, 2) matching a sherd with a non-composite (single copy of a design) pattern combining template matching algorithm with a dual-source CNN re-ranking method to find its underlying design, and 3) matching a sherd with a composite (multiple copies of a design) pattern using a Chamfer Matching based method. The framework was evaluated on a set of sherds from the heartland of the paddle-stamping tradition with a subset of known paddle-stamped designs of Pre-colonial southeastern North America. Extensive experimental results show the effectiveness of the proposed framework and algorithms.

pdf

Extracting Domain Information using Deep Learning
Amit Gupta and Weijia Xu (Texas Advanced Computing Center, University of Texas at Austin); Pankaj Jaiswal (Oregon State University); and Crispin Taylor and Jennifer Regala (ASPB)

Abstract

Across various scientific domains, digital publication of technical documents, often in the form of conference/journal article submissions, are the first accessible instance of new human knowledge in these respective fields. Synthesizing and curating this information is a slow and difficult process and often requires non-trivial human expertise. Given the ever increasing rate of these publications and the natural limitations of manual approaches, a computational solution to this problem is the paramount need of the hour. We have developed a computational tool (DIVE) to provide entity extraction and expert curation functionality for large document collections to list biological entities central to each article. However these methods depend on prior knowledge and language model often yield low precision in practice. This is caused by new unseen vocabulary and improper weighting of existing entities.

In this paper, we investigate how deep learning methods may be used to address this issue. Using the author feedback mechanism in our deployed tool we were able to create a expert user annotated gold standard dataset based on articles submitted over an entire year, that enables us to contrast several supervised machine learning methods for the entity extraction task. Our results show that DIVE's ensemble of methods (regular expression rules, keyword dictionaries, ontology files) have higher precision and recall scores than ABNER CRF based models. However, both tools in general show low precision.We further investigate using deep learning (using theNeuroNER tool) to improve the precision and recall performance of DIVE with promising early results.

pdf

Predictive Analytics using Amazon Web Services
Sanjay Padhi (Amazon.com)

Biography

Abstract

This presentation will provide an overview of AWS, advanced services such as Machine Learning, Data Analytics, and HPC for research. It will give a brief outline of studies that use predictive analytics including fully managed Artificial Intelligence (AI) services to help researchers build, train and deploy ML models in various domains including Computer Vision and Natural Language Processing (NLP). Supervised and unsupervised based learning frameworks and its implications in the fields of Scientific Computing, Medical Imaging, Cancer detection, Diabetic Retinopathy, and Voice-enabled solutions to improve management of chronic disease will be discussed. Collaborations on research with funding agencies such as the National Science Foundation and National Institutes of Health (including NIH STRIDES), as well as AWS Research Credit program will also be highlighted.

Tuesday 11:00am-12:30pm
Atlanta
Panel · Panel

Balancing Local and National Cyberinfrastructure Resources

Eric Coulter (Indiana University, XSEDE) and Richard Knepper and Resa Reynolds (Cornell University, XSEDE)

Abstract

This panel will invite cyberinfrastructure professionals from a wide range of universities to discuss their efforts to meet researchers' needs with a mix of local and national cyberinfrastructure solutions. Researchers bring a broad variety of unique requests for computational capabilities, storage, and expertise. In university research computing organizations, campus providers identify an equally broad range of solutions that meet these needs, and implement them in a number of configurations. These choices and the lessons learned from implementing solutions and working with researchers to adopt them are valuable experiences for the community at large considering balancing the costs and benefits of local hardware and infrastructure investments and leveraging national solutions such as the Open Science Grid (OSG), the Extreme Science and Engineering Discovery Environment (XSEDE), domain-specific gateways, the Department of Energy's INCITE program, and others, as well as balancing local and national resources, managing the transition of users from local to national, and deciding how to support users who make use of resources from multiple locations.

Discussion will be guided by the XSEDE National Cyberinfrastructure Resource Integration team, with an eye to fostering community among resource providers, and getting at common issues across the national CI landscape.

pdf

Tuesday 11:00am-12:30pm
New Orleans
Panel · Panel

Impact Measurement Strategies for Science Gateways: Present and Future

Michael Zentner (Purdue University, Science Gateways Community Institute); Stephen Burley (Rutgers Univeristy, RCSB Protein Data Bank / rcsb.org); Gerhard Klimeck (Purdue University, Network for Computational Nanotechnology / nanoHUB.org); Emre Brookes (UT Health, GenApp / genapp.rocks); Mark Miller (San Diego Supercomputer Center, CIPRES / phylo.org); Sabine Brunswicker (Purdue University, RCODI / rcodi.org); Dan Stanzione (University of Texas Austin, Texas Advanced Computing Center / tacc.utexas.edu); and Sharon Tettegah (University of California Santa Barbara, Center for Black Studies Research)

Abstract

For the purpose of this panel, a Science Gateway is taken to mean an online platform that enables the connection of at least three entities: i) expensive instruments, supercomputers, or databases / datasets, ii) authors of scientific codes that may consume resources or streams of data emanating from the first category and who may wish to make their codes available for others to use, and iii) vast audiences interested in using the products of the first two categories. By employing distributed computing and data infrastructures, Science Gateways facilitate the creation of significant quantities of research products in the form of new tangible and sometimes intangible digital artifacts. Such products may ultimately take the form of publications, but often may also be new standalone scientific products like datasets and scientific software. Furthermore, they also transform the process of research, for example by changing the direction of collaborations, failed experiments, advances in the "science of cyberinfrastructure" that supports research, or even databases of simulation results that may be used for a variety of purposes. They may also impact education and industry collaboration via new educational approaches. Many of these research product midpoints may never see the light of publication, and yet they contribute significantly to the enterprise of science worldwide. If and when they do reach a stage of publication, often the gateways that supported them are silent partners in such end results, making it difficult and likely inaccurate to measure impact with conventional methods. Compounding this problem is that the publication process runs much slower than the speed of thought enabled by science gateways, sometimes to the point that traditionally measured impact of a science gateway effort may only be possible near the end or after the expiration of funding for the gateway effort. While largely the scientific community acknowledges the value of such platforms often employing distributed computing and data infrastructures, measuring impact in a way that matches the speed of scientific productivity enabled by science gateways, and the acceptance of such measures by the community as a valid standard to justify funding and existence, remains unsolved.

This panel will engage the creators and operators of significant science gateways, large centers that support science gateways, and individuals with expertise in impact measurement to discuss what is currently being done, and the possibility of initiating an effort that would enable an accepted means of approaching this impact problem. Chosen gateway operator panelists can highlight the value of assets contained in the gateway, extensive analysis of user data for alternative impact, cost trade-offs of gateway development, and the value of resources that would otherwise be unavailable to researchers (see panelist list below for examples of these impact measures). Impact measurement expertise on the panel will present lessons from computational social science in analyzing the impact of gateways and future directions for such work (see panelist list).

Note to reviewers: Two possible panelists in addition to those below are checking into conflicts in time with the PEARC meeting.

Proposed panel duration: 90 minutes

Panel Participants:

Dan Stanzione: Dr. Dan Stanzione, Associate Vice President for Research at The University of Texas at Austin since 2018 and Executive Director of the Texas Advanced Computing Center (TACC) since 2014, is a nationally recognized leader in high performance computing. He is the principal investigator (PI) for a National Science Foundation (NSF) grant to acquire and deploy Frontera, which will be the fastest supercomputer at any U.S. university. Stanzione is also the PI of TACC's Stampede2 and Wrangler systems, supercomputers for high performance computing and for data-focused applications, respectively. For six years he was co-PI of CyVerse, a large-scale NSF life sciences cyberinfrastructure. Stanzione was also a co-PI for TACC's Ranger and Lonestar supercomputers, large-scale NSF systems previously deployed at UT Austin. Stanzione received his bachelor's degree in electrical engineering and his master's degree and doctorate in computer engineering from Clemson University.

Sharon Tettegah: Dr. Tettegah, Director & Professor, Center for Black Studies Research at the University of California, Santa Barbara. She is the principal investigator for a National Science Foundation grant, Coordinating Curricula and User Preferences to Increase the Participation of Women and Students of Color in Engineering. Investigators will engage in data mining of syllabi, course content, public spaces, and instructor materials to aggregate information about curricula presented in different representational forms, such as equations, images, narratives, simulations, and videos through CUE4CHANGE science gateway based on HUBzero.

Stephen K. Burley: Dr. Burley is the director of the RCSB Protein DataBank (RCSB PDB), which manages the Protein Data Bank (PDB) archive jointly with Worldwide Protein Data Bank partners in Europe (PDB in Europe; PDBe), Asia (PDB Japan; PDBj), and the BioMagResBank (BMRB). The PDB archive, established in 1971, is the single global data resource for 3D experimental structures of large biological molecules (proteins, DNA, and RNA). Since 2000, contributions to the PDB have come from more than 30,000 structural biologists worldwide. Millions of data consumers coming from every sovereign nation recognized by the United Nations access the PDB at no charge on an annual basis. Approximately ~2 million structure data files are downloaded every day from the PDB archive, without limitations on usage. The estimated replacement value of PDB data is ~$15 billion. The RCSB PDB is supported jointly by the NSF, NIH, and DOE. An independent economic impact study estimated a 15000-fold return on federal investment, excluding the impact of the PDB data on the biopharmaceutical industry. A recently published study showed that ~5,900 PDB structures contributed to 184 of 210 US FDA New Drug Approvals for the years 2010-2016.

Gerhard Klimeck: Dr. Klimeck is the director of nanoHUB.org, a science gateway established in 1998 and serving 16,000 simulation users annually and an additional 1.4 million unique consumers of other resources on nanoHUB annually. nanoHUB gathers and analyzes significant amounts of user telemetry data with novel clustering and visualization methods to demonstrate use in more than 200 classroom settings annually, and to date has impacted more than 35,000 students. nanoHUB has been cited more than 2000 times in the scientific literature, which, if it were an individual would have an h-index of 82.

Emre Brookes: Dr. Brookes is the project lead for GenApp, a framework for rapidly generating complete applications. Applications can be generated in a variety of target languages, including stand-alone GUI and web based. GenApp is being used by multiple projects to produce science gateways. SASSIE-web, serving the small angle scattering community, currently has over 600 registered users who ran over 20k jobs in 2018 and has generated over 40 publications. The goal of GenApp is to minimize the effort of scientist-developers to build and maintain applications so that they can maximize their "science" productivity. A key metric for GenApp is the amount of effort scientists can save in deploying gateways and applications, to where a gateway has been deployed in under an hour.

Mark Miller: Dr. Miller is the PI of the CIPRES science gateway, created in December 2009 from the CIPRES Portal, a highly used gateway for phylogenetic research. CIPRES supports nearly 20 parallel codes, and provides access to XSEDE resources. Since its creation, CIPRES has served over 35,000 unique users running 1.50 million jobs that have run more than 115 million core hours. Currently CIPRES supports over 25,000 monthly job submissions from over 2000 users. It has been used by over 90 instructors in an educational setting, and has supported over 5,200 publications.

Sabine Brunswicker: Dr. Brunswicker is a computational social scientist and Director of the Research Center for Open Digital Innovation at Purdue University. Her work is particularly focused on open digital innovation, describing new ways of using information technologies to organize the collective design and use of innovative digital goods on digital platforms such as science gateways. In her work, she designs and examines systems and technologies that support open digital innovation with respect to their technological and behavioral impact. She uses techniques of computational social science (agent-based modeling, network analysis, experiments) to predict individual as well as collective outcomes in open digital innovation.

Return to Top

Tuesday 11:00am-5:00pm

Tuesday 11:00am-5:00pm
Comiskey
Workshop · Workshop Full-day, Student Program

Student Program: Student Modeling Challenge

Katharine Cahill (Ohio Supercomputer Center) and Jennifer Houchins and Aaron Weeden (Shodor)

Abstract

The PEARC Student Modeling Challenge has been offered over the past several years and has been a successful activity to introduce students to the concepts and practice of computational modeling as a method for engaging in scientific analysis. During the challenge, students work in teams to develop a model of a scientific problem, choosing from a selection of problems from a number of different domains including biology, physics, and chemistry. Over the course of the day, the teams read background information about their problem, identify the questions to be considered, work with a basic starting Python code, and develop a solution by adding complexity. Finally, the teams each create a presentation of their work and results throughout the day. This gives the students hands-on experience not only using computational modeling skills but also working in teams on a technical problem and presenting their work. Most students are not explicitly introduced to computational modeling as a skill in their science or programming classes, so this is often a new and meaningful experience for them. An Introduction to Python tutorial has also been submitted and prepares students to use Python in the Student Modeling Challenge.

pdf

Tuesday 11:00am-5:00pm
Water Tower
Workshop · Workshop Three-quarter Day

Trusted CI Workshop on Trustworthy Scientific Cyberinfrastructure

Jim Basney (NCSA)

Abstract

The Trusted CI Workshop on Trustworthy Scientific Cyberinfrastructure provides an opportunity for sharing experiences, recommendations, and available resources for addressing cybersecurity challenges in research computing. Presentations by Trusted CI staff and community members will cover a broad range of cybersecurity topics, including science gateways, transition to practice, cybersecurity program development, workforce development, and community engagement (e.g., via the Trusted CI Fellows program). A selection of topics follows.

5-year Vision for an NSF Cybersecurity Ecosystem: The NSF community is large and diverse, encompassing NSF itself, its seven science directorates, over two dozen Large Facilities, and tens of thousands of smaller ephemeral projects. This community is tightly integrated with the higher education institutions and research laboratories that provide administrative homes for projects. The community also collaborates closely with communities from other federal and non-federal agencies, as well as with the international science community. The diversity of these projects' science missions, combined with the complexities of implementing cybersecurity and open science in tandem, creates a serious cybersecurity challenge. There is no off-the-shelf approach to cybersecurity for open science that the NSF community can adopt. Even Large Facilities, the largest of the NSF projects, struggle to develop tailored approaches. To address this challenge, an approach is needed to manage risks – while providing both flexibility for project-specific adaptations and access to the necessary knowledge and human resources for implementation. Hence, the Trusted CI vision is for "A NSF cybersecurity ecosystem, formed of people, practical knowledge, processes, and cyberinfrastructure, that enables the NSF community to both manage cybersecurity risks and produce trustworthy science in support of NSF's vision of a nation that is the global leader in research and innovation." See also: http://hdl.handle.net/2022/22178

A Framework for Open Science Cybersecurity Programs: Trusted CI is developing a framework for establishing and maintaining an open science cybersecurity program at any project scale and stage in a project's life cycle. Such a framework is useful even for projects having significant compliance requirements (e.g., FISMA, HIPAA, NIST SP 800-171) in that it provides a starting point for evolving a cybersecurity program rather than hundreds of pages dense with unprioritized requirements. Trusted CI's framework is built around four pillars: Mission Alignment, Governance, Resources, and Controls. Like the pillars supporting any structure, all are vital and required for an efficient and effective cybersecurity program.

Cybersecurity for Science Gateways: Science gateways provide the science community with access to scientific data and resources. Trusted CI partners with the Science Gateways Community Institute (SGCI) to provide cybersecurity training and consultation for science gateways at all phases of development. We will present the risks to science gateways and how to mitigate them.

Transition to Practice: Trusted CI's Cybersecurity Research Transition To Practice (TTP) program is open to all NSF funded cybersecurity researchers and practitioners. We will discuss how we can work together to identify the cybersecurity gaps in scientific CI, and leverage cybersecurity research to fill those gaps. Information on Research for TTP from the academic community including AI/ML for information security and IOT/CPS research will be discussed, as well as the process we can enable to bring research and practitioners together to leverage the research in support of scientific cyberinfrastructure.

Fellows Program: Trusted CI launched the Open Science Cybersecurity Fellows Program in 2019. The program identifies members of the scientific community (e.g., an IT professional working with a science project), empowers them with basic knowledge of cybersecurity and the understanding of Trusted CI's services, and has them serve as cybersecurity liaisons to their respective community. They then assist members of the community with basic cybersecurity challenges and connect them with Trusted CI for advanced challenges. The members of the fellows program will present their activities in a panel format at the workshop.

Lightning Talks and Open Discussion: Based on experience from our PEARC17 workshop and Cybersecurity Summit meetings, it is valuable to end the day with broad community participation in the form of lightning talks and open discussion. Community members are invited to give 5 minute presentations about experiences, recommendations, and available resources for addressing cybersecurity challenges in research computing. Following these presentations, members of the Trusted CI leadership team facilitate a discussion about takeaways from the topics covered in the workshop and open questions about topics not covered.

About Trusted CI: Trusted CI, the NSF Cybersecurity Center of Excellence, is comprised of cybersecurity experts who have spent decades working with science and engineering communities and have an established track record of usable, high-quality solutions suited to the needs of those communities. The team draws from best operational practices and includes leaders in the research and development of new methodologies and high-quality implementations. Trusted CI addresses the challenge for meeting the needs of individual cyberinfrastructure projects through deep engagements, training, and dissemination of experiences. At the same time, Trusted CI advances the state of cybersecurity practice across the community by analyzing gaps in cybersecurity technology to provide guidance to researchers and developers, addressing the application of software assessment to complicated cyberinfrastructure software stacks, and fostering the transition of cybersecurity research to practice.

pdf

Return to Top

Tuesday 12:30pm-1:30pm

Tuesday 12:30pm-1:30pm
Regency CD
Lunch · Lunch

Lunch

Return to Top

Tuesday 1:30pm-3:00pm

Tuesday 1:30pm-3:00pm
Regency AB
Paper · Facilitation, Technical Paper, Student Technical Paper, Cluster Management

Managing Clusters I

Chair: Kent Milfeld (University of Texas, Texas Advanced Computing Center)

Managing a Heterogeneous Cluster
Kyle Hutson, Dan Andresen, Adam Tygart, and Dave Turner (Kansas State University)

Abstract

Most HPC clusters are purchased with a large quantity of identical hardware, which is maintained through its lifecycle and then another HPC cluster takes its place. However, some clusters, like ours, are maintained by frequently adding new hardware, which is then integrated into the system. Over the years, the cluster has grown to include 339 compute nodes with 8220 cores from 6 vendors, spanning 4 generations of CPUs; 7 network technologies from 6 switch vendors (1GbE-100Gb OPA); 102 GPUs (3 different GPU models); 28 storage nodes (3.15 PB raw storage); and 7 virtualization nodes hosting 65 VMs.

Having such a diverse system has significant advantages, although the management is more difficult. This paper outlines our strategy of managing this very heterogeneous and complex system. Topics covered include software optimization, consistency of operating system updates, identity management, resource prioritization, network infrastructure, storage, and management of non-compute-intensive resources. Our combination of open source and internally developed software used to manage this cluster are a model to other heterogeneous systems and to smaller clusters which have not expanded because of management worries.

pdf

A Multi-Environment HPC-Scale Puppet Infrastructure for Compliance and Systems Automation
Raychel Benson, Edward Munsell, Nicholas Bertrand, Michael Baynton, Evan Bollig, and Jeffrey McDonald (University of Minnesota, Minnesota Supercomputing Institute)

Abstract

Unifying configuration management across High Performance Compute (HPC) and supporting infrastructure systems is a challenge commonly faced by research computing centers. System operators no longer have the luxury to manage one-off cases manually; automation is essential. Centralized tooling and site-wide configuration promises efficiencies, but the substantial differences between, for example, compute nodes and networking gear can dissuade operators from attempting to manage their entire fleet with a single tool. To solve this, the Minnesota Supercomputing Institute (MSI) at the University of Minnesota assembled an innovative collection of utilities for Institute-wide management of systems with Puppet 5 as the centerpiece. The innovative solution emerged from the need for more than a dozen system operators/administrators to coordinate changes across a medium- to large-scale data-center with over 1500 nodes. A centralized base configuration ensures all systems are inline with University security policies and other compliance needs, while the Puppet infrastructure additionally enables per-cluster or even per-node customizations as needed.

This document presents the architecture of MSI's orchestration and management infrastructure, as well as the workflow followed by operators to provision disparate systems with Puppet and ensure that quality, accountability, and compliance requirements are met. The resulting system is actively used in the day-to-day management of the Institute, including persistent infrastructure services, cluster head- and compute-nodes, and even networking equipment. The workflow described herein is recommended for similar research computing institutions, including those with as few as five nodes and two operators.

pdf

Expressing and Managing Network Policies for Emerging HPC Systems
Sergio Rivera, James Griffioen, Zongming Fei, and Jane Hayes (University of Kentucky)

Abstract

▾ Abstract

Traditional high performance computing (HPC) centers that operate a single large supercomputer cluster have not required sophisticated mechanisms to manage and enforce network policies. Recently, HPC centers have expanded to support a wide range of computational infrastructure, such as OpenStack-based private clouds and Ceph object stores, each with its own unique characteristics and network security requirements. Network security policies are becoming more complex and harder to manage. To address the challenge, this paper explores ways to define and manage the new network policies required by emerging HPC systems. As the first step, we identify the new types of policies that are required and the technical capabilities needed to support them. We present example policies and discuss ways to implement those policies using emerging programmable networks and intent-based networks. We describe our initial work toward automatically converting human readable network policies into network configurations and programmable network controllers that implement those policies using business rule management systems.

pdf

Tuesday 1:30pm-3:00pm
Crystal C
Paper · Storage, Technical Paper, Student Technical Paper, Applications, Gold Exhibitor Technical Talk

Applications-Based Storage

Chair: Scott Yockel (Harvard University)

Greyfish: An Out-of-the-Box, Reusable, Portable Cloud Storage Service
Carlos Redondo (The University of Texas at Austin) and Ritu Arora (Texas Advanced Computing Center)
▾ Abstract

A scalable storage system is an integral requirement for supporting large-scale cloud computing jobs. The raw space on storage systems is made usable with the help of a software layer which is typically called a filesystem (e.g., Google's Cloud Filestore). In this paper, we present the design and implementation of an open-source and free cloud-based filesystem named as "Greyfish" that can be installed on the Virtual Machines (VMs) hosted on different cloud computing systems, such as Jetstream and Chameleon. Greyfish helps in: (1) storing files and directories for different user-accounts in a shared space on the cloud, (2) managing file-access permissions, and (3) purging files when needed. It is currently being used in the implementation of the Gateway-In-A-Box (GIB) project. A simplified version of Greyfish, known as Reef, is already in production in the BOINC@TACC project. Science gateway developers will find Greyfish useful for creating local filesystems that can be mounted in containers. By doing so, they can independently do quick installations of self-contained software solutions in development and test environments while mounting the filesystems on large-scale storage platforms in the production environments only.

pdf

Exploring Untrusted Distributed Storage for High Performance Computing
Austin Smith, Justin Riley, Muneeba Syed, Milan Kupcevic, Paul Edmon, and Scott Yockel (Harvard University)
▾ Abstract

High performance computing systems are typically built with high-throughput and infrastructural uniformity in mind, but generally do not easily accommodate diverse data security requirements on a single cluster. Rather than fracturing that infrastructure by building many network isolated storage "islands" to secure each dataset covered by an individual data use agreement, we explore using the Ceph distributed storage system with client-side encryption to provision secure storage from a single, untrusted data lake.

pdf

Architecting Parallel File Systems for Modern Data Challenges
Sam Spencer (DDN)
▾ Abstract

Managing data at scale is a continuously evolving challenge. Today's desire to retain, analyze and productize increasingly diverse and growing data sets strains the capabilities of legacy architectures. Learn about recent advances in parallel file system architectures that DDN has brought to market are designed to address the challenges introduced by the recent focus on machine learning and artificial intelligence, and how these innovations are making data faster to access, easier to manage and more efficient to store.

Tuesday 1:30pm-3:00pm
Wrigley
Paper · Technical Paper, Applications, Science Gateways

Science Gateways I

Chair: Carol Song (Purdue University)

Leveraging Elasticsearch to Improve Data Discoverability in Science Gateways
Jake Rosenberg, Josue Coronel, Joseph Meiring, Sarah Gray, and Tracy Brown (Texas Advanced Computing Center)
▾ Abstract

The Texas Advanced Computing Center (TACC) maintains an expanding portfolio of science gateway projects aimed at connecting researchers in diverse domains with TACC's high-performance computing (HPC) resources. Recent additions include gateways for the NeuroNex Technology Hub, Planet Texas 2050, and the UT Research Cyberinfrastructure. These projects all share the goal of facilitating the curation, processing, and sharing of data by users. By providing a web interface for data curation and submission of jobs to HPC systems, we are able to provide advanced computing capabilities to users without specialist knowledge of HPC architectures. We also provide new features, such as interactive applications, which are unavailable to users with only command-line access.

As the volume of data shared using these gateways grows, a key architectural challenge and request from our users is to make this data quickly and efficiently discoverable from the user interface. As the volume of data grows and with some projects generating tens of millions of files, the discoverability of data has often been a hindrance to the science goals of a project, especially in a collaborative environment. We describe our efforts to enhance data discoverability at TACC using Elasticsearch, a search engine and document store which provides the capability to analyze and search file metadata.

pdf

Implementing a Flexible, Fault Tolerant Job Management System for Science Gateways
Dimuthu Wannipurage, Suresh Marru, Marlon Pierce, and Eroma Abeysinghe (Science Gateways Research Center, Pervasive Technology Institute, Indiana University, Bloomington, Indiana); Gourav Shenoy (IBM Watson Health, Cambridge, Massachusetts); Ajinkya Dhamnaskar (Paypal, San Jose, California); Sudhakar Pamidighantam and Marcus Christie (Science Gateways Research Center, Pervasive Technology Institute, Indiana University, Bloomington, Indiana); and Lahiru Jayathilaka (University of Moratuwa)
▾ Abstract

This paper summarizes our experiences evaluating and deploying a new task execution management system within the open source Apache Airavata framework for science gateways. We base our choices on our operational requirements and experiences running Airavata software as a multi-tenanted production service for multiple gateway clients. Our considerations include integrating semi-independent components, making major upgrades to those components while retaining the system's overall functionality, and choosing between integrating third party and in-house developed components. While we focus on Apache Airavata as the platform for evaluation, our results should be of general interest. After considering the options of extensions to our previous, in-house job management system using Apache Kafka or replacing it with Kubernetes, we ultimately chose Apache Helix, primarily for its ability to execute multiple tasks coupled into directed acyclic graphs. We have integrated this approach into Apache Airavata and have tested extensively over several months with many thousands of jobs, both from our internal throughput testing and operational tests with early adopter science gateway clients. The new system has proven to be at least as reliable as the previous system with the advantages that we now have simplified maintenance, do not need to support an in-house system that required extensive developer training to modify, and can support more sophisticated job execution scenarios.

pdf

Web of Trust Tool for Gateway User Vetting
Kenneth Yoshimoto (San Diego Supercomputer Center); Nicholas Carnevale (Yale University); and Subhashini Sivagnanam, Amitava Majumdar, and Mark Miller (San Diego Supercomputer Center)
▾ Abstract

The Neuroscience Gateway (NSG) is a science gateway that provides neuroscience researchers with tools on National Science Foundation high performance computing (HPC) platforms. The capability of NSG to accept not just static data for processing, but also user code to be compiled and executed, requires a user vetting process. In order to provide an automated and distrubted mechanism for this process, a web of trust tool, consisting of Gnu Privacy Guard (GPG) and helper scripts in newLISP and Expect were developed.

pdf

Tuesday 1:30pm-3:00pm
Gold Coast
Paper · Technical Paper, Machine Learning/Artificial Intelligence, ML/AI Applications

Application II

Chair: Amit Gupta (Texas Advanced Computing Center)

Developing a Meta Framework for Key-Value Memory Networks on HPC Clusters
Choonhan Youn (UC San Diego), Arghya Das (University of Wisconsin Platteville), and Seungwon Yang and Joohyun Kim (Louisiana State University)
▾ Abstract

We propose a novel framework, DARE-MetaQA, whose underlying architecture is versatile to support various learning models for large-scale Question Answering (QA), in particular aimed at the biomedical domain, for example, backed with PubMed as an information resource for finding facts for answering. Our framework addresses inherent challenges with the existing model of the Key-Value Memory Networks in two specific aspects: (1) prediction performance and (2) computational scalability. For the first aspect, an extension of algorithmic repertoire leveraging ensemble and meta-learning is proposed, and the second aspect is tackeld with the choice of effective distributed parallel techniques. The resulting meta framework is highly advantageous for developing a flexible learning model, desirable for dealing with the complicated nature of machine reasoning. In this work, we focus on the development side towards an optimal implementation for a specific computing environment of shared-nothing multi-node systems. Specifically, the two high-end cluster systems (Comet and Bridges) were selected, both of which are equipped with multiple GPUs for each node and support the use of the container technology for HPC, Singularity, We highlight our main implementation decisions and achievements during the development, and discuss theoretical and technical underpinnings for the current framework architecture and the future work. The development has been supported by the XSEDE ECSS program.

pdf

VAStream: A Visual Analytics System for Fast Data Streams
Raju Gottumukkala, Satya Katragadda, and Siva Venna (University of Louisiana at Lafayette)
▾ Abstract

Processing high-volume, high-velocity data streams is an important big data problem in many sciences, engineering, and technology domains. There are many open-source distributed stream processing and cloud platforms that offer low-latency stream processing at scale, but the visualization and user-interaction components of these systems are limited to visualizing the outcome of stream processing results. Visual analysis represents a new form of analysis where the user has more control and interactive capabilities either to dynamically change the visualization, analytics or data management processes. VAStream provides an environment for big data stream processing along with interactive visualization capabilities. The system environment consists of hardware and software modules to optimize streaming data workflow (that includes data ingest, pre-processing, analytics, visualization, and collaboration components). The system environment is evaluated for two real-time streaming applications. The real-time event detection using social media streams uses text data arriving from sources such as Twitter to detect emerging events of interest. The real-time river sensor network analysis project uses unsupervised classification methods to classify sensor network streams arriving from the US river network to detect water quality problems. We discuss implementation details and provide performance comparison results of various individual stream processing operations for both stream processing applications.

pdf

Automated Deep Learning Analysis of Purple Martin Videos Depicting Incubation and Provisioning
Heather Williams, L. Shawn Matott, and Robert DeLeon (State University of New York at Buffalo)
▾ Abstract

Deep learning models have been developed to analyze automatically video clips of Purple Martin nesting behavior. Two separate models have been constructed, one for incubation and one for provisioning. The incubation model is a simple two class model that analyzes the videos to determine if an adult Purple Martin is incubating the eggs/young nestlings or not. The model is a Keras/Tensor Flow convolutional neural network (CNN) trained with ~12 thousand still images and achieves a validation data set accuracy of ~99.5% on still images. A comparison of the results of the automated video analysis with sample validation videos viewed manually shows good agreement; the model approaches human accuracy. Some conclusions from the incubation analysis will be discussed. The provisioning analysis requires a much more complex 3 class model which must distinguish between zero, one parent or both parents on the nest. With training sets including ~26 thousand images the CNN model demonstrates a validation set accuracy of ~99% on the still images. However, the actual the video analysis still presents difficulties. Several different CNN models have been tried but results were similar. The best results to date on analyzing the videos for provisioning events have been ~88% accuracy with ~10% false positives. A discussion of the conclusions from the provisioning model and model analysis will be presented.

pdf

Tuesday 1:30pm-3:00pm
Atlanta
Panel · Panel

Community Engagement at Scale: NSF Centers of Expertise

Ruth Marinshaw (Stanford University); Daniel Crawford (MoISSI); Ewa Deelman (CI CoE Pilot); Jennifer Schopf (EPOC); Von Welch (ResearchSOC, Trusted CI); Nancy Wilkins-Diehr (SGCI); and Frank Wuerthwein (OSG)
▾ Abstract

Summary: This panel brings together the leaders of centers of expertise serving the CI and NSF communities to present what they wish everyone knew about their respective area and to explore the challenges and lessons learned with the cross-cutting topic of community engagement at scale.

Moderator: Ruth Marinshaw, CTO for Research Computing at Stanford University and leader of the Stanford Research Computing Center

Time: 90 minutes preferred, 60 minutes if requested.

What do science gateways, network performance, molecular science, cybersecurity operations, distributed computing and cyberinfrastructure development have in common? They are all domains that are the focus of centers of expertise funded by NSF to serve the community. They bring together teams with expertise in technical domains and community engagement with the goal of doing community engagement at scale broadly across the science community. This panel brings together the leaders of these centers to present what they wish everyone knew about their respective area and to discuss the challenges and lessons learned with the cross-cutting topic of community engagement at scale.

Benefits of the panel include: >Disseminating knowledge about community engagement to the community and to other and future projects. >Capture and advance our shared knowledge about community engagement. >Build the community's awareness of the centers of expertise and their value. >Entertaining discussion around differences of perspective from the centers.

The agenda for the panel will be:

1. Each speaker presents for <5 minutes on (1) what they wish everyone knew about their respective area, and (2) what are the biggest surprises they have learned in community engagement?

2. A moderated discussion which:

>Explores the similarities and differences in the presented surprises. >Explores what relationships should exist at the international, national level (with other centers, big CI projects, and NSF itself), with regional networks, and at the state and institutional levels. >Explores the needs of the target communities and how they can be met by the centers. >Explores the challenges and opportunities of developing CI workforce.

Differentiation from similar offerings: This is a unique opportunity to present the community-driven efforts around developing centers that aim to enhance the national cyberinfrastructure.

pdf

Return to Top

Tuesday 1:30pm-5:00pm

Tuesday 1:30pm-5:00pm
New Orleans
Tutorial · Tutorial Half-day

Tutorial on Floating-Point Analysis Tools

Ignacio Laguna (Lawerence Livermore National Laboratory); Michael Bentley, Ian Briggs, and Ganesh Gopalakrishnan (University of Utah); Harshitha Menon (Lawerence Livermore National Laboratory); and Cindy Rubio Gonzalez (UC Davis)
▾ Abstract

Scientific software is central to the practice of research computing. While scientific software is widely used in several science and engineering disciplines to simulate real-world phenomena, developing accurate and reliable scientific software is notoriously difficult. One of the most serious difficulties come from dealing with floating-point arithmetic to perform numerical computations. Round-off errors occur and accumulate at all levels of computation, and compiler optimizations and low precision arithmetic can significantly affect the final computational results. With accelerators such as GPUs dominating high-performance computing systems, computational scientists are faced with even bigger challenges, given that ensuring numerical reproducibility in these systems pose a very difficult problem. This tutorial will demonstrate tools that are available today to analyze floating-point scientific software. We focus on tools that allow programmers to get insight about how different aspects of floating-point arithmetic affect their code and how to fix potential bugs. The tools presented in the tutorial will allow programmers to understand how compiler optimizations affect floating-point computations, detect hidden floating-point exceptions on GPUs, reduce floating-point precision to obtain performance speedups, and understand the sensitivity of different regions of the code to floating-point rounding errors.

pdf

Tuesday 1:30pm-5:00pm
Crystal A
Tutorial · Tutorial Half-day

Accelerating Data Science Workflows with RAPIDS

Jonathan Bentz and Zahra Ronaghi (NVIDIA)
▾ Abstract

The open source RAPIDS project allows data scientists to GPU-accelerate their data science and data analytics applications from beginning to end, creating possibilities for drastic performance gains and techniques not available through traditional CPU-only workflows. Learn how to GPU-accelerate your data science applications by:

* Utilizing key RAPIDS libraries like cuDF (GPU-enabled Pandas-like dataframes) and cuML (GPU-accelerated machine learning algorithms) * Learning techniques and approaches to end-to-end data science, made possible by rapid iteration cycles created by GPU acceleration * Understanding key differences between CPU-driven and GPU-driven data science, including API specifics and best practices for refactoring

Upon completion, you'll be able to refactor existing CPU-only data science workloads to run much faster on GPUs and write accelerated data science workflows from scratch.

pdf

Return to Top

Tuesday 2:00pm-4:00pm

Tuesday 2:00pm-4:00pm
Field
Co-located Meeting · Private, Closed Meeting, Co-located Meeting

Open OnDemand Quarterly Meeting

Alan Chalker (Ohio Supercomputer Center)
▾ Abstract

We'd be potentially interested in hosting a special event in conjunction with PEARC19. Specifically, it would be an invitation only meeting ~2 hours in duration for the Open OnDemand project team. We have quarterly project team meetings, and since most of the team will be at PEARC it makes sense to try to time it around the event. We're pretty flexible as to when specifically it would happen, but obviously wouldn't want to overlap with any of the major PEARC program elements (especially since Tom is on the team). Having it be during a meal would be fine and we have plenty of project budget to pay any expenses that would be incurred. I'd expect there to be ~10 onsite attendees, and we'd want to setup a WebEx to allow other team members to connect in remotely if needed.

Return to Top

Tuesday 3:00pm-3:30pm

Tuesday 3:00pm-3:30pm
Regency Foyer/Concourse Level Foyer/Crystal Foyer
Break · Break

Afternoon Break

Return to Top

Tuesday 3:30pm-5:00pm

Tuesday 3:30pm-5:00pm
Crystal C
Paper · Data, Technical Paper, Student Technical Paper, Applications, Workflows, Phil Andrews Award, Best Paper Applications, Best Student Paper Applications

Data Management Workflows

Chair: Tony Arick (Mississippi State University)

Shuffler: A Large Scale Data Management Tool for Machine Learning in Computer Vision
Best Student Paper Applications
Evgeny Toropov (Carnegie Mellon University); Paola Buitrago (Pittsburgh Supercomputing Center, Carnegie Mellon University); and José Moura (Carnegie Mellon University)
▾ Abstract

In the academic research community, datasets in the computer vision (CV) field are primarily static. Once a dataset has been accepted as a benchmark for a CV task, researchers will not alter it in order to make their results reproducible. At the same time, when exploring new tasks and new applications, datasets tend to be an ever-changing entity. A practitioner may combine existing public datasets, filter images or objects in them, change annotations or add new ones to fit a task in hand, visualize sample images, or perhaps output statistics in the form of text or plots. In fact, the dataset changes as the practitioner experiments with both the data and algorithms. Considering that ML and deep learning call for large volumes of data, it is no surprise that the resulting data and software management associated with research dealing with alive datasets can become quite complex. At the moment of this publication, there is no flexible, publicly available instrument that would facilitate manipulating image data and their annotations throughout a ML pipeline. In this work, we present Shuffler, an open source tool that makes it easy to manage large CV datasets. It stores annotations in a relational, human-readable database. Shuffler defines over 40 data handling operations with annotations that are commonly useful in supervised learning applied to CV and supports some of the most well-known CV datasets. Finally, it is easily extensible, making the addition of new operations and datasets a fast and easy to accomplish task.

pdf

Integrity Protection for Scientific Workflow Data: Motivation and Initial Experiences
Phil Andrews Award, Best Paper Applications
Mats Rynge, Karan Vahi, and Ewa Deelman (University of Southern California Information Sciences Institute); Anirban Mandal (RENCI - University of North Carolina, Chapel Hill); Omkar Bhide, Randy Heiland, Von Welch, and Raquel Hill (Indiana University); Ilya Baldin (RENCI - University of North Carolina, Chapel Hill); and William Poehlman and F. Alex Feltus (Clemson University)
▾ Abstract

With the continued rise of scientific computing and the enormous increases in the size of data being processed, scientists must consider whether the processes for transmitting and storing data sufficiently assure the integrity of the scientific data. When integrity is not preserved, computations can fail and result in increased computational cost due to reruns, or worse, results can be corrupted in a manner not apparent to the scientist and produce invalid science results. Technologies such as TCP checksums, encrypted transfers, checksum validation, RAID and erasure coding provide integrity assurances at different levels, but they may not scale to large data sizes and may not cover a workflow from end-to-end, leaving gaps in which data corruption can occur undetected. In this paper we explore an approach of assuring data integrity - considering either malicious or accidental corruption - for workflow executions orchestrated by the Pegasus Workflow Management System. To validate our approach, we introduce Chaos Jungle - a toolkit providing an environment for validating integrity verification mechanisms by allowing researchers to introduce a variety of integrity errors during data transfers and storage. In addition to controlled experiments with Chaos Jungle, we provide analysis of integrity errors that we encountered when running production workflows.

pdf

Introducing the Open Science Chain – Protecting Integrity and Provenance of Research Data
Subhashini Sivagnanam (SDSC/UCSD) and Viswanath Nandigam and Kai Lin (SDSC)
▾ Abstract

Data sharing is an integral component of research and academic publications, allowing for independent verification of results. Researchers have the ability to extend and build upon prior research when they are able to efficiently access, validate, and verify the data referenced in publications. Despite the well known benefits of making research data more open, data withholding rates have remained constant. Some disincentives to sharing research data include lack of credit, and fear of misrepresentation of data in the absence of context and provenance. While there are several research data sharing repositories that focus on making research data available, there are no cyberinfrastructure platforms that enable researchers to efficiently validate the authenticity of datasets, track the provenance, view the lineage of the data and verify ownership information. In this paper, we introduce and provide an overview of the NSF funded Open Science Chain, a cyberinfrastructure platform built using blockchain technologies that securely stores metadata and verification information about research data and tracks changes to that data in an auditable manner in order to address issues related to reproducibility and accountability in scientific research.

pdf

Tuesday 3:30pm-5:00pm
Regency AB
Paper · Facilitation, Technical Paper, Containers/VMs

Containers/Virtual Machines

Chair: Randy Herban (Sylabs, Inc)

Performant Container Support for HPC Applications
Yinzhi Wang, R. Todd Evans, and Lei Huang (Texas Advanced Computing Center)
▾ Abstract

The demand for the ability to easily customize, reproduce and migrate applications and workflows has been steadily increasing among the HPC community as software environments and applications grow in complexity. Lightweight containers that are suitable for HPC applications at scale are considered to be a viable approach to meet this demand. Previous studies have addressed the performance aspects of most existing containers using microbenchmarks and revealed the performance overheads of the best implementations to be small. However, the feasibility of providing containerized, real-world HPC applications on HPC systems, and the impact on overall application performance at scale, has not yet been explored. Here we present a basic feasibility and performance study using the Singularity container. We evaluate what is required to enable container images to utilize the high-speed fabric present on most HPC systems and explore their performance by comparing real-world applications run both within a container and in the absence of a container (natively). The results indicate lightweight Singularity images are a promising approach to the HPC communities demands for not only customizability, reproducibility and portability, but also performance.

pdf

Making it More Secure: The Technical and Social Challenges of Expanding the Functionality of an Existing HPC Cluster to Meet University and Federal Data Security Requirements
Jason Christopher (UC Berkeley), Gary Jung (Lawrene Berkeley National Laboratory), and Christopher Doane (UC Berkeley)
▾ Abstract

As High Performance Computing and Big Data analytics become more commonplace, we see researchers applying these tools in new areas. Indeed, in the past few years, we've seen the use of HPC in diverse areas such as archeology, public policy, and digital humanities. So it comes as no surprise that many life science researchers are now approaching us to use large scale computation and data analytics on their sensitive data sets, such as de-identified patient or genomics data, for the purposes of scientific inquiry. At UC Berkeley, this has become a pressing issue, as existing faculty need a place to do research on sensitive data. And we knew of at least one instance where it affected the campus' ability to recruit a new faculty member. Clearly something had to be done!

An informal survey informed us that most other institutions built a new dedicated system to support their sensitive data research, including identified and HIPPA data. This paper is a case study of how we met this need by using a methodology to apply our campus cybersecurity framework, with the help of our institution's cybersecurity team, to convert our traditional production HPC cluster, with over 2000 users across 100 research groups, and our VM service offering to also support this type of research. Our efforts show that is is not only possible, but also it is a practical alternative to take this approach instead of building an all new environment.

pdf

Developing and Operating Edge Services on Federated Infrastructure Using MiniSLATE
Joe Breen (University of Utah), Robert Gardner (University of Chicago), Shawn McKee (University of Michigan), Jason Stidd (University of Utah), Christopher Weaver (University of Chicago), Benjamin Kulbertis (University of Utah), Lincoln Bryant (University of Chicago), Jiahui Chen and Emerson Ford (University of Utah), Rose Pierce and Benedikt Riedel (University of Chicago), Mitchell Steinman and Luan Truong (University of Utah), Jeremy Van and Iljia Vukotic (University of Chicago), Gage Glupker (University of Michigan), and Skyler Griffith (University of Utah)
▾ Abstract

Modern software development workflow patterns often involve the use of a developer's local machine as the first platform for testing code. SLATE mimics this paradigm with an implementation of a light-weight version, called MiniSLATE, that runs completely contained on the developer's local machine or scales to larger machines (laptop, virtual machine, or another physical server). MiniSLATE resolves many development environment issues by providing an isolated and local configuration for the developer. Application developers are able to download MiniSLATE which provides a fully orchestrated set of containers on top of a production SLATE platform, complete with central information service, API server, and a local Kubernetes cluster. This approach mitigates the overhead of a hypervisor but still provides the requisite isolated environment. They are able to create the environment, iterate, destroy it, and repeat at will. A local MiniSLATE environment also allows the developer to explore the packaging of the edge service within a constrained security context in order to validate its full functionality within limited permissions. As a result, developers are able to test the functionality of their application with the complete complement of SLATE components local to their development environment without the overhead of building a cluster or virtual machine, registering a cluster, interacting with the production SLATE platform, etc.

pdf

Tuesday 3:30pm-5:00pm
Wrigley
Paper · Technical Paper, Student Technical Paper, Applications, Science Gateways

Science Gateways II

Chair: Erik Gough (Purdue University)

A Science Gateway for Atomic and Molecular Physics
Sudhakar Pamidighantam (Indiana University); Barry Schneider (NIST); Klaus Bartschat (Drake University); Igor Bray (Curtin University); Armin Scrinzi (Ludwig-Maximilians Universit¨at,); Fernando Martin (Universidad Aut´onoma de Madrid); Jonathan Tennyson (University College London); Jimena Garfinkiel (School of Physical Sciences, The Open University); Oleg Zatsarinny (Drake University); and Markus Klinker (Universidad Autónoma de Madrid,)
▾ Abstract

This paper describes the creation of a new Atomic and Molecular Physics science gateway (AMPGateway). The gateway is designed to bring together a subset of the AMP community to work collectively to make their codes available and easier to use by the partners as well as others. By necessity, a project such as this requires the developers to work on issues of portability, documentation, ease of input, as well as making sure the codes can run on a variety of architectures. Here we describe our efforts to build this AMP gateway and future directions. The paper is divided into four major sections. In the introduction we provide a history of how and why the project got started and our decision to go to XSEDE for support for the gateway. In the second section, we present some information on the AMP codes. Section three is devoted to the details of the deployment. In section four we discuss general community building. A brief description of the theoretical and computational tools embodied in the codes is given. In the first stage of our efforts the software suites are deployed as independent applications with application specific input interfaces. The community building has started and a few other community software suites identified to be included in a second phase. The interoperability among the software suites will be addressed as a follow-on. The AMPGateway uses multi tenanted Apache Airavata middleware framework served by the SciGaP hosting services for sustained operation.

Expanding the Computational Anatomy Gateway from clinical imaging to basic neuroscience research
Daniel Tward, Anthony Kolasny, and Fatima Kahn (Johns Hopkins University, Center for Imaging Science); Juan Troncoso (Johns Hopkins University School Of Medicine, Department of Pathology); and Michael Miller (Johns Hopkins University, Department of Biomedical Engineering)
▾ Abstract

The Computational Anatomy Gateway, powered largely by the Comet (San Diego Supercomputer Center) and Stampede (Texas Advanced Computing Center) clusters through XSEDE, provides software as a service tools for atlas based analysis of human brain magnetic resonance images. This includes deformable registration, automatic labeling of tissue types, and morphometric analysis. Our goal is to extend these services to the broader neuroscience community, accommodating multiple model organisms and imaging modalities, as well as low quality or missing data. We developed a new approach to multimodality registration: by predicting one modality from another, we can replace ad hoc image similarity metrics (such as mutual information or normalized cross correlation) with a log likelihood under a noise model. This statistical approach enables us to account for missing data using the Expectation Maximization algorithm. For portability and scalability we have implemented this algorithm in tensorflow. For accessibility we have compiled and many working examples for multiple model organisms, imaging systems, and missing tissue or image anomaly situations. These examples are made easily usable in the form of Jupyter notebooks, and made publicly available through github. This framework will significantly reduce the barrier to entry for basic neuroscientists, enabling the community to benefit from atlas based computational image analysis techniques.

pdf

Reproducible Hydrological Modeling with CyberGIS-Jupyter: A Case Study on SUMMA
Fangzheng Lyu, Dandong Yin, and Anand Padmanabhan (University of Illinois at Urbana-Champaign); Young Don Choi and Jon Goodall (University of Virginia); Anthony Castronova (CUAHSI); David Tarboton (Utah State University); and Shaowen Wang (University of Illinois at Urbana-Champaign)
▾ Abstract

CyberGIS-Jupyter is a cyberGIS framework for achieving data- intensive, reproducible, and scalable geospatial analytics using Jupyter Notebook based on advanced cyberinfrastructure. As a cutting-edge hydrological modeling framework, the Structure for Unifying Multiple Modeling Alternative (SUMMA) functions as a unified approach to process-based modeling. The purpose of this research is to investigate the feasibility of coupling CyberGIS- Jupyter with SUMMA to realize reproducible hydrological modeling. CyberGIS-Jupyter is employed to systematically integrate advanced cyberinfrastructure, including two high- performance computers – Virtual ROGER and XSEDE Comet, data management, and execution and visualization of SUMMA- based modeling. By taking advantage of CyberGIS-Jupyter, users can easily tune different parameters for a SUMMA model and submit computational jobs for executing the model on HPC resources without having to possess in-depth technical knowledge about cyberGIS or cyberinfrastructure. Computational experiments demonstrate that the integration of CyberGIS-Jupyter and SUMMA achieves a high-performance and easy-to-use implementation for reproducible SUMMA-based hydrological modeling.

pdf

Tuesday 3:30pm-5:00pm
Gold Coast
Paper · Technical Paper, Student Technical Paper, Machine Learning/Artificial Intelligence, Image Analysis

Image Analysis and Recognition

Chair: David Walling (Texas Advanced Computing, University of Texas at Austin)

Towards a Framework for Validating Machine Learning Results in Medical Imaging
William Monroe, Thomas Anthony, Murat Tanik, and Frank Skidmore (UAB)
▾ Abstract

In the medical imaging domain, non-linear warping has enabled pixel by pixel mapping of one image dataset to a reference dataset. This co-registration of data allows for robust, pixel-wise, statistical maps to be developed in the domain, leading to new insights re- garding disease mechanisms [18]. Deep learning technologies have given way to some impressive discoveries. In some applications, deep learning algorithms have surpassed the abilities of human im- age readers to classify data. As long as endpoints are clearly defined, and the input data volume is large enough, deep learning networks can often converge and reach prediction, classification, and segmen- tation with success rates as high or higher than human operators [11]. However, machine learning, and deep learning algorithms are complex and interpretability is not always a straightforward byproduct of the classification performed. Visualization techniques have been developed to add a layer of interpretability. The work presented here compares a simplified machine learning workflow for medical imaging to a statistical map from a previous study to validate that the machine learning model used does indeed focus its attention on known important regions.

pdf

Facial Expression Recognition: Utilizing Digital Image Processing, Deep Learning, and High-Performance Computing
Francisco Reveriano (University of Houston - Clear Lake)
▾ Abstract

The purpose of this project was to analyze which image pre-processing technique was most beneficial in improving the performance of Facial Expression Recognition through Deep Learning and High-Performance Computing. Contrary to our expectations, the results obtained in this work showed that deep learning does not significantly benefit from various commonly used image pre-processing techniques such as resizing, smooth, or edge detection. As generally expected, the results showed that an increase in accuracy was obtained by simply increasing the size of the training dataset. This study proceeds to show that the increase in training data size was easily handled by the High-Performance Computing (HPC) cluster provided by the Pittsburg Supercomputing Centered through XSEDE.

pdf

In and out of the nucleus - CNN-based segmentation of cell nuclei from images of a translocating biosensor
Tarek Zikry, Katarzyna Kedziora, Michael Kosorok, and Jeremy Purvis (UNC-Chapel Hill)
▾ Abstract

This study demonstrates application of convolutional neural networks (CNNs) for the analysis of a unique image analysis problem in fluorescence microscopy. We employed the U-Net CNN architecture and trained a model to segment nuclear regions in images of a translocating biosensor—which alternates between the nucleus and cytoplasm—without the need for a constant nuclear marker. The model provided high-quality segmentation results that allowed us to accurately quantify the extent of cyclin-dependent kinase activity in a population of cells. We envision that the development of this kind of analysis tools will enable biologists to design live-cell fluorescence imaging experiments without the need for providing a constant marker for a subcellular region of interest. As a consequence, they will be free to increase the number of biosensors measured in single cells or reduce the phototoxicity of cellular imaging.

pdf

Tuesday 3:30pm-5:00pm
Atlanta
Panel · Panel

Professionalization and Career Arcs for Roles Supporting Compute- and Data-intensive Research

Patrick Schmitz (University of California, Berkeley); Scott Yockel (Harvard University); Stuart Geiger (University of California, Berkeley); Joanna Thielen (Oakland University); and Kerk Kee (Chapman University)
▾ Abstract

Research across all domains is becoming more data- and computationally intensive, and more and more domains are adopting computational and data-centric methods. This change is increasing the importance of roles for consulting, instruction, and training of researchers in the effective use of the associated infrastructure and tools, roles for supporting the infrastructure itself, and roles for discovery and management of the datasets generated and used in research. There is a related rise in the roles of data scientists, both independently and as part of research teams. All of this has led to discussion of how library, data science, and Research IT support roles should be integrated into academic organizational models to build effective teams and enterprises.

This panel brings together four projects that are investigating these issues, and are developing models for supporting and advancing professionals who purse these various careers. The panelists will summarize the respective projects, and we will explore the lessons learned for different roles. Attendees will engage the panelists, and learn how they can leverage the lessons both as individuals and at their institutions.

pdf

Return to Top

Tuesday 4:00pm-6:00pm

Tuesday 4:00pm-6:00pm
McCormick
Co-located Meeting · Private, Closed Meeting, Co-located Meeting

PEARC SC Meeting

Leslie Froeschl (University of Illinois at Urbana-Champaign)

Return to Top

Tuesday 4:30pm-6:30pm

Tuesday 4:30pm-6:30pm
Regency D
Student Program · Student Program

Student/Exhibitor Speed Networking

Return to Top

Tuesday 5:15pm-6:15pm

Tuesday 5:15pm-6:15pm
Field
Co-located Meeting · Private, Closed Meeting, Co-located Meeting

PEARC20 Program Committee Meeting

GWEN JACOBS (University of Hawaii)
▾ Abstract

Meeting of the PEARC20 planning committee. Expecting about 20 committee members for an hour meeting.

Tuesday 5:15pm-6:15pm
Atlanta
Birds of a Feather · Birds of a Feather, Building Community

What the Heck is a Microaggression?

Jackie Milhans (Northwestern University); Marisa Brazil (Purdue University, Campus Champions); and Gladys Andino Bautista (Purdue University)
▾ Abstract

In this Birds of a Feather session, the effects of microaggressions on the research computing and data workforce will be discussed. Microaggressions are brief and common verbal, nonverbal, and environmental slights and indignities, whether they are intentional or not, towards marginalized groups. These acts can have major impact on the productivity of teams and individuals, as they can lead to lack of participation, engagement, and performance.

It is widely reported that diverse and inclusive teams are more innovative, productive, and effective. Inclusivity is also important in the retention of employees . The purpose of this discussion is to make the PEARC community aware that recruiting diverse professionals into the field is not enough. While unaware of an existing study on the demographics of the research computing and data workforce, studies of the higher education IT workforce show demographics as largely male, white, and older. In order to attract and retain the diverse workforce needed to provide researchers important tools and skills, the research computing and data community members at all levels need to create an inclusive community. Developing and retaining skilled individuals is required to support cloud-based and containerized workflows, growing numbers of end users, and emerging and growing computational research areas (such as digital humanities, genomics, and the social sciences). This Birds of a Feather session will be focused on one key area of providing an inclusive environment so that we can develop a workforce that we need to support research successfully.

pdf

Tuesday 5:15pm-6:15pm
New Orleans
Birds of a Feather · Birds of a Feather, Data

Management, Movement, and Lifecycle of High-Volume Research Data

Jennifer Hartman (Northwestern University)
▾ Abstract

Track: Facilitation of Advanced Research Computing

This particular Birds of a Feather session will be focused on bringing together individuals from research universities that are looking to discuss data storage infrastructure expansion, data workflow management, data lifecycle policies, moving data storage to the cloud, creating hybrid data storage options, and issues involving secure data movement to different tiers or types of storage locations. The goal is to facilitate new ideas around managing research data, discuss challenges related to growing research data volume, and provide a forum for members of the research IT community to network with each other and share resources that may benefit all members of the community.

pdf

Tuesday 5:15pm-6:15pm
Crystal A
Birds of a Feather · Birds of a Feather, Building Community

Northeast Cyberteam Program - a Workforce Development Strategy for Research Computing at Small and Medium Sized Institutions

John Goodhue (Massachusetts Green High Performance Computing Center); Julie Ma (Massachussetts Green High Performance Computing Center, Northeast Cyberteam Initiative); Adrian Del Maestro and Stephen Everse (University of Vermont); Bruce Segee (University of Maine); Scott Valcourt (University of New Hampshire); Sia Najafi (Worcester Polytech Institute); and Ralph Zottola (University of Alabama)
▾ Abstract

The Northeast Cyberteam is an NSF-funded program that aims to make regional and national cyberinfrastructure more accessible to researchers at small and mid-sized institutions that lack the critical mass to support these resources on campus. The program is starting its third year of operation in Maine, New Hampshire, Vermont and Massachusetts. We are taking a two-pronged approach:

* Build a regional pool of research computing facilitators (RCFs), and a process to share them across institutional boundaries.

* Develop web-based tools and information to enable self-service learning by RCFs and end users.

Central to the Cyberteam Program is the idea of giving interested students hands on RCF experience while advancing research/education projects that need assistance. Toward that end, we have launched 28 projects over the past two years that pair student facilitators with experienced mentors to work with researchers and educators at small and medium sized institutions who have specific near term needs.

The program relies heavily on the Northeast Cyberteam Portal, which serves three purposes: (1) manage project workflows and capture results; (2) serve as a front door to the self-service learning resources; and (3) encapsulate the knowledge and experience that we are gaining, with the goal of making it possible to replicate the methodology in other regions.

At this BOF, we will describe the Northeast Cyberteam process; demonstrate the tools that support it; and share plans for future enhancements. We are especially interested in engaging others who might have interest in replicating the model.

pdf

Tuesday 5:15pm-6:15pm
Crystal C
Birds of a Feather · Birds of a Feather, Technical Tools

Open OnDemand User Group Meeting

Alan Chalker, David Hudak, Eric Franz, Morgan Rodgers, and Doug Johnson (Ohio Supercomputer Center); Steve Gallo (University at Buffalo); and Bob Settlage (Virginia Tech)
▾ Abstract

Open OnDemand is an NSF-funded open-source HPC portal based on OSC's original OnDemand portal. The goal of Open OnDemand is to provide an easy way for system administrators to provide web access to their HPC resources, including, but not limited to: -Plugin-free web experience -Easy file management -Command-line shell access -Job management and monitoring across different batch servers and resource managers -Graphical desktop environments and desktop applications

This BoF is meant to be an open discussion to guide the future roadmap for Open OnDemand in the near term, by getting feedback from the community on the prioritization of the various tasks planned for the next few years.

As there are many people that attended this BOF at PEARC'18 and spoke highly of it, the session leaders intend to replicate the same BOF format (with appropriate updates regarding what has been done in the past year and the roadmap for the newly started phase 2 NSF award) and anticipate more attendees than were present at the PEARC'18 BOF. A report summarizing the status of current installations and additional feature requests will be generated and distributed to anyone that has expressed interest in Open OnDemand.

Tuesday 5:15pm-6:15pm
Water Tower
Birds of a Feather · Birds of a Feather, Technical Tools

Supporting Federation for Large-Scale Collaborations

Martial Michel (Data Machines Corp.), Craig Lee (The Aerospace Corporation), Robert Bohn (NIST), Khalil Yazdi (Yazdi and Associates), and John Lowe (Indiana University)
▾ Abstract

Cloud federation standardization efforts continue to evolve. The NIST/IEEE Joint WG on Federated Cloud has been continuing their work, and the Open Research Cloud Alliance has been actively developing consensus across the scientific community stakeholders. ORCA's goal is to support international science collaborations. NIST's goal is to "define" federation, build a conceptual model, map-out federation deployment models, and identify needed areas of federation-specific standardization. The IEEE goal is to push these federation standards through the standardization process.

Jetstream is one of the first NSF-funded open, interactive production systems targeted at the national science and engineering research communities to provide on-demand computing and storage anytime, anywhere. By design and in practice, Jetstream is programmable cyberinfrastructure in that users can reconfigure the system to suit their needs.

Keystone is the OpenStack Identity Service: it provides API client authentication, service discovery, and distributed multi-tenant authorization. Keystone was designed from the ground up to be amenable to multiple styles of backend; it is organized as a group of internal services exposed on one or many endpoints.

This BoF will provide a discussion around cloud federation, their use case and implementation, and discuss possible federation deployment and governance models that embody the key concepts and design principles.

pdf

Tuesday 5:15pm-6:15pm
Gold Coast
Birds of a Feather · Birds of a Feather, Building Community

Xpert Network

Rudolf Eigenmann (University of Delaware)
▾ Abstract

This BOF aims to bring together: -Researchers developing and using computational and data-intensive (CDI) applications, -Those who assist these researchers with expertise in CDI technology, and -Developers of tools that support the development and use of CDI applications. The presenters will report on the progress of an NSF-supported project with these goals and will create discussions among the BOF participants. The goal is to create synergy among these groups and, in this way, begin to build a community and a record of best practices and tools for those who conduct CDI research and assist these researchers with expertise and tools. The BOF follows a series of online and in-person meetings pursuing these goals.

pdf

Tuesday 5:15pm-6:15pm
Wrigley
Birds of a Feather · Birds of a Feather, Resource Access

The Role of SSH in Campus and National Computing Services

Jim Basney (NCSA) and John-Paul Navarrro (Univ. of Chicago/ANL)
▾ Abstract

For many years, SSH was the only interface to large-scale high-performance and high-throughput computing systems. Recent XSEDE systems--and many campus systems--are now offering web browser interfaces to their services, including CyVerse (for cloud computing), Open OnDemand (for HPC computing), and Jupyter (for interactive data analysis and education). In this session, XSEDE and SP representatives and users of both XSEDE and campus computing services will share experiences with new ways to access large-scale computing systems and outline where SSH is still necessary. The goal is to clarify where SSH fits in context with newer computing interfaces.

Return to Top

Tuesday 6:30pm-8:30pm

Tuesday 6:30pm-8:30pm
Crystal Foyer and Crystal B
Posters · Reception, Poster, Student Poster

Poster Reception

Approaches to High Resolution Network Telemetry & Analytics with Machine Learning In Support of High Performance Computing
Corey Eichelberger (National Center for Supercomputing Applications)
▾ Abstract

Network measurement and analytics are key to the overall operation, planning and root-cause analysis of issues within network infrastructure. The increased line rate (100Gb/s and beyond) of network connections along with the explosion in data transfer needs for scientific datasets has changed the methodology of effectively monitoring network infrastructure. Network measurement utilizing 5 minute polling intervals with binary threshold based alerting has proven to be unreliable for accurate measurement and alerting of critical network systems. This presentation will discuss the ongoing efforts at NCSA to gather high resolution (< 10s collection interval) network telemetry data utilizing SNMP and streaming telemetry with machine learning being utilized to analyze and generate alerts on the data being collected.

A Brief Introduction to GAIL (Guaranteed Automatic Integration Library) Version 2.3
Sou-Cheng Choi and Jagadeeswaran Rathinavel (Illinois Institute of Technology Chicago), Xin Tong (University of Illinois at Chicago), and Kan Zhang and Fred Hickernell (Illinois Institute of Technology Chicago)
▾ Abstract

Function approximation, optimization, and integration are among the most fundamental mathematical problems. They are especially challenging when the functions in consideration fluctuate wildly in certain parts of the domain, or if the domain is high dimensional. Development of efficient and effective numerical algorithms for approximating the solutions to these problems that are based on rigorous mathematical frameworks, (statistical) error bounds, advanced sampling strategies, and careful complexity analysis remain active research areas. The Guaranteed Automatic Integration Library (GAIL) is our multi-year research efforts directed towards addressing the aforementioned overarching theme. Briefly, GAIL is a free, open-source MATLAB software package with eight main algorithms based on over a dozen peer-reviewed publications, involving three faculty and nine doctoral students in Illinois Tech as well as their international collaborators. The theories behind our algorithms consider functions in cones, rather than balls. GAIL algorithms sample data points of a given function in an automatic, adaptive manner. Our algorithms are iterative by design, with error bounds computed in each iteration. The solvers stop at the first iteration where user-given tolerance is met. They are proven to have optimal costs. We consistently employ good industrial software development practices for GAIL such as unit tests and integration tests, searchable online documentation, version control systems (GitHub), and agile methodology (using Trello). Our reproducible research have demonstrated that GAIL could outperform more established algorithms in various ways.

Improving Science Gateways usage reporting for XSEDE
Amit Chourasia and Scott Sakai (SDSC, UCSD); Michael Shapiro (NCSA, UIUC); and Steven Gallo (University of Buffalo)
▾ Abstract

Science Gateways have gained wide use by research community by providing a rich and easy to use web interface to access and utilize complex cyberinfrastructure. Science Gateways are consuming an increasing proportion of computational capacity provided by XSEDE. A typical approach employed by Science Gateways is to use a single community account with a compute allocation to process compute jobs on behalf of their users. The computation usage for Science Gateways is compiled from batch job submission systems and reported by the XSEDE service providers. However, this reporting does not capture and provide any information about the user who actually initiated the computation, as the batch systems do not have record this information. To overcome this reporting limitation, Science Gateways utilize a separate pipeline to submit job-specific attributes to XSEDE, which is then later co-joined with batch system information submitted by the Service Providers to create detailed usage reports. In this article we describe improvements to the Gateway attribute reporting system, which better serves the needs of the growing number of Science Gateways by providing their operators a simpler and versatile interface to easily report a richer set of attributes and ultimately publish this information via XDMoD that enables various stake holders to assess resource utilization and their potential impact.

pdf

Data Commons to Support University-Wide Cross Discipline Research
Morgan Rodgers, Eric Franz, Alan Chalker, and David Hudak (Ohio Supercomputer Center) and Christopher Stewart, Raghu Machiraju, Steve Chang, and Cathie Smith (Ohio State University)
▾ Abstract

The Translational Data Analytics Institute (TDAI) at the Ohio State University (OSU) aims to promote cross discipline work in order to improve data science and analytics in many areas of university-wide research for use in solving real world problems including areas such as smart city, smart agriculture, and biomedical research. As a result of their interest in promoting transdisciplinary work, TDAI decided that they needed to grapple with the problems of where do researchers get their data and where do they store it for the long term. In particular they were interested in improving the discoverability of and access to data sets that have already been produced by other researchers. To support their mission TDAI envisioned the Data Commons as a website where researchers upload research data sets, curators link data sets together, and researchers visualize and analyze data sets in a data processing environment. Development of this data portal has gone through several phases: requirements gathering, a "ground up" implementation by a local contractor, transfer to the Ohio Supercomputer Center (OSC) for additional development and hosting, and an initial release. Requirements gathering for version 2 is in progress.

pdf

Experiences from scaling scale Science Gateway operations
Suresh Marru (Indiana University, Science Gateways Research Center Pervasive Technology Institute) and Marlon Pierce, Eroma Abeysinghe, Sudhakar Pamidighantam, Marcus Christie, and Dimuthu Wannipurage (Indiana University)
▾ Abstract

Science Gateways provide a crucial user-centric and science-centric point of entry to the collection of computing, storage, and software that are used to support science and commonly referred to as cyberinfrastructure (CI). Over the past two decades, Science Gateways have dramatically increased cyberinfrastructure usage and accessibility for scientists and educators around the world. Gateways provide a federating bridge over cyberinfrastructure spanning campus resources, National Resources likes XSEDE, commercial clouds, and international computing resources.

Science Gateways develop and operate a domain-specific presentation layer as well as a generic gateway middleware layer that supports a common set of required functionalities. Repeating this development process for each new gateway is inefficient and wasteful of resources. In this paper summarizes our experiences, architectural and deployment choices made in within the open source and open community-based Apache Airavata framework to create a robust, sustainable Science Gateway Platform (SciGaP) which minimizes the net operating cost of Science Gateways. for science gateways.

pdf

HPC Ecosystems Project: Facilitating Advanced Research Computing in Africa
Bryan Johnston (CHPC)
▾ Abstract

The HPC Ecosystems Project is an initiative of the Centre for High Performance Computing in South Africa (CHPC) and is responsible for the distribution and re-deployment of decommissioned HPC hardware as mid-tier systems to research institutions both locally within South Africa and regionally across Africa.

Primarily focusing on Human Capital Development for HPC-related skills (both for administration and usage) and the deployment of HPC capacity in Africa, the Project runs regular on-site workshops to develop the growing cohort of HPC System Administrators, typically co-located with an HPC site deployment. Square Kilometre Array (SKA) partner sites (part of the SKA-readiness programme) and local South African sites (Universities and research institutes) are the primary recipients of this training, although other nations outside of this scope have been and continue to be invited to participate (self-funded). The Project spends varying amounts of time providing HPC cluster administration support, training and workshops / lectures on general parallel computing concepts.

The project aims to enable self-sufficient, independent, functional and successful HPC sites while also producing skilled administrators. Currently the HPC Ecosystems community has representation from more than 100 people across 16 nations and 3 continents.

HPC Equipment has been repurposed from leading HPC Centers such as The Texas Advanced Computing Center (TACC - Ranger and Stampede equipment), Cambridge University, and the CHPC.

The most recent deployment of equipment is TACC's former Stampede HPC, to countries across Africa.

pdf

Enabling rich data sharing for Science Gateways via the SeedMe2 platform
Amit Chourasia and David Nadeau (SDSC, UCSD); Jiaping Luo (UCSD); Tony Chen and Mark Miller (SDSC, UCSD); and Emre Brookes (UT Health San Antonio)
▾ Abstract

Science Gateways provide an easily accessible and powerful computing environment for researchers. These are built around specific software tools that are heavily used by large research communities in specific domains. Science Gateways have been catering to a growing need of researchers for computational tools, however their usage model is single user-centric. With a growing user base, the need for collaborative capabilities in Gateways has been emerging. One such need is a the ability to share data/results to others. Tools for data management and data sharing have not been addressed by many gateways. In this poster we will describe and discuss our effort to provide a provide rich environment for data organization and sharing via integrating the SeedMe2 platform with two Science Gateways: CIPRES and GenApp.

Detailed description of this work in pap134, that was recommended for poster. The enclosed poster will be updated to include Gateway information.

pdf

Trusted CI, the NSF Cybersecurity Center of Excellence
Von Welch (Indiana University) and Jim Basney (NCSA)
▾ Abstract

Trusted CI, the NSF Cybersecurity Center of Excellence, is comprised of cybersecurity experts who have spent decades working with science and engineering communities and have an established track record of usable, high-quality solutions suited to the needs of those communities. The team draws from best operational practices and includes leaders in the research and development of new methodologies and high-quality implementations.

Trusted CI addresses the challenge for meeting the needs of individual cyberinfrastructure projects through deep engagements, training, and dissemination of experiences. At the same time, Trusted CI advances the state of cybersecurity practice across the community by analyzing gaps in cybersecurity technology to provide guidance to researchers and developers, addressing the application of software assessment to complicated cyberinfrastructure software stacks, and fostering broadly the transition of cybersecurity research to practice.

Trusted CI's mission is to provide the NSF community a coherent understanding of cybersecurity's role in producing trustworthy science and the information and know-how required to achieve and maintain effective cybersecurity programs.

This mission is accomplished through one-on-one engagements with projects to address their specific challenges; education, outreach, and training to raise the state of security practice across the scientific enterprise; and leadership on bringing the best and most relevant cybersecurity research to bear on the NSF cyberinfrastructure research community.

For more information about what Trusted CI does, how it can help your project, and the advances it is making in cybersecurity and resources for cybersecurity professionals, please visit our poster, our exhibit table, and trustedci.org.

Interdependent Networked Community Resilience Modeling Environment (IN-CORE)
Jong Lee, Christopher Navarro, Nathan Tolbert, Yong Wook Kim, Vara Veera Gowtham Naraharisetty, Chen Wang, Michal Ondrejcek, Diego Rivera, Shannon Bradley, and Mark Fredricksen (National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign)

Abstract

The National Institute of Standards and Technology (NIST) funded the multi-university five-year Center of Excellence for Risk-Based Community Resilience Planning (CoE), headquartered at Colorado State University, to develop the measurement science to support community resilience assessment. Measurement science is implemented on a platform called Interdependent Networked Community Resilience Modeling Environment (IN-CORE). On IN-CORE, users can run scientific analyses that model the impact of natural hazards and resiliency against the impact on communities. Initial implementation of IN-CORE, called version 1, adopted the Ergo Platform, a multi-hazard risk assessment platform originally developed by NCSA with MAE center. The platform is built upon the Eclipse Rich Client Platform (RCP), a plug-in based architecture for building standalone applications. With the initial implementation, the CoE identified needs of a modeling environment which is easy for researchers to use and develop their own models and data. In addition, potential need of web applications was identified. In order to meet the needs, IN-CORE version 2 has been designed and developed. Different from version 1, the new implementation is a web-enabled architecture. The platform is built on a Kubernetes cluster with Docker container technology. On the cluster, customized JupyterHub, python library of scientific analyses, web services, and light-weight web applications are implemented. This paper will present the architecture and implementation of IN-CORE version 2. Note that it will be released as an open source project in latter half of 2019.

pdf

Leveraging Public Cloud Services for CLIA-Certified Personalized Medicine Pipelines
Evan Bollig, Christine Henzler, Ham Lam, Sarah Munro, and Rebecca LaRue (University of Minnesota, Minnesota Supercomputing Institute); Getiria Onsongo (Macalester College); Sophia Yohe and Andrew Nelson (University of Minnesota, Department of Laboratory Medicine and Pathology); Matthew Bower and Matthew Schomaker (Fairview, Molecular Diagnostic Lab); and Bharat Thyagarajan (University of Minnesota, Department of Laboratory Medicine and Pathology)
Abstract

Dating back to 2012, a joint effort between Fairview Hospital's Molecular Diagnostic Lab and the Minnesota Supercomputing Institute introduced the Next-Generation Sequencing Diagnostic Pipeline (NGSDP), a novel cloud-based analysis pipeline that was validated for clinical use to be compliant with the Clinical Laboratory Improvement Amendments (CLIA) . Since then, the partnership has grown the portfolio of active pipelines to seven, and expanded the available gene panels to over 6700 genes. Likewise, the design of the backing cloud infrastructure has been significantly overhauled for efficiency, scalability, automation, and fault tolerance.

This document introduces our revised cloud-native, scheduled infrastructure solution for containerized applications. The cloud-native infrastructure, deployed on Amazon Web Services, has two core components: a compute cluster of virtual machines, and a simple queue-based job scheduler. The cluster runs analysis pipelines packaged as Docker containers; each with their own resource requirements. The scheduler, akin to products for traditional HPC, allows batch submission of patient samples with a persistent job queue.

Our solution has been successfully running as CLIA-validated pipelines since 2016. Two variants of the solution are presented to address the need for such architecture in both the traditional public cloud space, as well as in Amazon's GovCloud where fewer cloud services are available to handle sensitive or controlled-access data. Furthermore, the innovative solution is driven by a traditional research computing environment; details of which are presented with emphasis on security, monitoring, and user workflow.

pdf

Software-Enhanced Teaching and Visualization Capabilities of an Ultra-High-Resolution Video Wall
Ramses Van Zon and Marcelo Ponce (University of Toronto, SciNet HPC Consortium)
Abstract

This paper presents a modular approach to enhance the capabilities and features of a visualization and teaching room using software. This approach was applied to a room with a large, high resolution (7680x4320 pixels), tiled screen of 13x7.5 feet as its main display, and with a variety of audio and video inputs, connected over a network. Many of the techniques described are possible because of a software-enhanced setup, utilizing existing hardware and a collection of mostly open-source tools, allowing to perform collaborative, high-resolution visualizations as well as broadcasting and recording workshops and lectures. The software approach is flexible and allows one to add functionality without changing the hardware.

pdf

ResearchSOC: Cybersecurity for the Nation's Research
Todd Stone, Von Welch, and Susan Sons (Indiana University); Richard Biever (Duke University); and Tom Davis and Inna Kouper (Indiana University)
Abstract

Research projects are prime targets for cyberattacks, and those attacks can prevent once-in-a-lifetime scientific discoveries. Such was the case in late 2017, when astrophysicists at Australia's Zadko Telescope found themselves unable to move their telescope to measure gravitational waves from colliding neutron stars because their systems had been hacked. Like the Zadko facility, many research projects are large, highly collaborative, and use diverse infrastructure. Additionally, these projects have requirements to seamlessly share and access data across institutional and geographic boundaries. And, unfortunately, these facilities face an increasing array of cyber threats to the integrity of their data, the availability of their instrumentation, and the accessibility of their systems. Moreover, research projects also share the condition of constrained resources to devote to cybersecurity. The ResearchSOC is an innovative, NSF-funded new model for providing key operational services and expertise to address these unique cybersecurity concerns. The ResearchSOC has three NSF large facilities as early adopter clients: the National Radio Astronomy Observatory (NRAO), the Gemini Observatory, and UNAVCO and the National Earth Science Geodetic Facility (GAGE). ResearchSOC is a collaborative effort of three institutions: Indiana University, Duke University, and the Pittsburgh Supercomputing Center. This new and unique model of for research cybersecurity aims to make scientific computing resilient to cyberattacks and capable of supporting trustworthy, productive research. The ResearchSOC does this by providing the cybersecurity services, expertise, training, and information sharing specifically designed for research and education projects. The poster outlines the rationale, organization, collaboration, services, benefits, and use cases for the ResearchSOC.

Application of GPU-Accelerated Molecular Dynamics Simulations of G-Quadruplexes with the Drude Polarizable Force Field
Alexa Salsbury, Justin Lemkul, and Anne Brown (Virginia Tech)
Abstract

Integrating computational tools in life science research is beneficial due to their low cost and ability to provide unique insights into biological and chemical phenomena. Molecular dynamics (MD) simulations is one such tool that can further characterize atomic-scale behaviors of biomolecules. MD simulations rely on force fields (FFs) to compute the potential energy of a given system of atoms, as such the quality of the chosen force field can greatly affect simulation results. The recent development of polarizable FFs is a significant advancement in MD that provides new insight on complex biomolecular systems. Our research group performs MD simulations of DNA G-Quadruplexes (GQs), which are fundamental to maintaining genomic stability and are implicated in cancer and neurodegenerative disease. Previous studies suggest that polarization is important for modeling GQs, so we use both the CHARMM36 (C36) nonpolarizable and Drude-2017 polarizable FFs to contextualize the role of polarization in the simulations. The Drude FF assigns negatively charged particles, called "Drude oscillators," to all non-hydrogen atoms in the system to allow for the quantification of induced electronic polarization. Adding these auxiliary charged particles to the system increases computational burden and the time it takes to produce trajectories, which is a primary shortcoming of polarizable FFs. Therefore, accelerating computational performance with GPU-equipped clusters is important to produce results in a practical time frame. Our work has utilized GPU- and CPU-based clusters to produce simulations with the C36 and Drude FFs, allowing us to compare cluster hardware and expand our understanding of nucleic acid dynamics.

Bridging Local Cyberinfrastructure and XSEDE with CyberGIS-Jupyter
Anand Padmanabhan, Dandong Yin, Fangzheng Lyu, and Shaowen Wang (University of Illinois at Urbana-Champaign)
Abstract

The fabric of national and international cyberinfrastructure ecosystems for scientific discovery and innovation can be viewed as distributed computing environments composed of powerful supercomputers, various cloud computing resources, and numerous local cyberinfrastructure (including both cloud and HPC) resources. However, extensive computational work conducted by academic researchers is often siloed in one of these environments. Science Gateways, by simplifying access to advanced cyberinfrastructure resources, have made significant progress on connecting these silos by enabling researchers in many fields to access advanced cyberinfrastructure through web browsers. This research bridges between national and local cyberinfrastructure resources through: (a) horizontal scaling of CyberGIS-Jupyter between the cloud resources provided by JetStream on the Extreme Science and Engineering Discovery Environment (XSEDE) and a VMWare-based cloud environment on Virtual ROGER, a local cyberinfrastructure resource hosted by the CyberGIS Center for Advanced Digital and Spatial Studies at the University of Illinois at Urbana-Champaign campus; and (b) enabling the submission of computationally intensive models to the batch systems of both Virtual ROGER and Comet. We have developed a mechanism to provide access to a scalable JupyterHub platform together with cutting-edge cyberGIS software and hardware called CyberGIS-Jupyter, which allows seamless access to high-performance computing (HPC) resources while shielding the complexity of managing cyberinfrastructure access from users. The user-friendly environment provided by CyberGIS-Jupyter along with computational scalability achieved through this research provides a powerful environment for conducting collaborative and reproducible research at scale with seamless access to advanced cyberinfrastructure at both local and national levels.

pdf

A Lightweight Framework for Research Data Management
Dimitar Nikolov and Esen Tuna (Indiana University)
Abstract

We describe a framework for managing live research data involving two major components. First, a system for the scalable scheduling and execution of automated policies for moving, organizing, and archiving data. Second, a system for managing metadata to facilitate curation and discovery with minimal change to existing workflows. Our approach is guided by four main principles: 1) to be non-invasive and to allow for easy integration into existing workflows and computing environments; 2) to be built on established, cloud-aware, open-source tools; 3) to be easily extensible and configurable, and thus, adaptable to different academic disciplines; and 4) to integrate with and take advantage of infrastructure and services available on academic campuses and research computing environments. These principles give our solution a well-defined place along the spectrum of research data management software such as sophisticated electronic lab notebooks and science gateways. Our lightweight and flexible data management framework provides for curation and preservation of research data within a lab, department or university cyberinfrastructure.

pdf

CHURP: A lightweight CLI framework to enable novice users to analyze sequencing datasets in parallel
Joshua Baller, Thomas Kono, Adam Herman, and Ying Zhang (University of Minnesota)
Abstract

Progressive decreases in the cost of DNA sequencing have contributed to a decades-long exponential increase in the production of new sequencing datasets. The processing of these datasets has in turn led biology, a field that has traditionally relied on local 'lab' servers to address its computational needs, to become increasingly reliant on High Performance Computing (HPC) resources. Though many operations on sequencing datasets are trivially parallelizable on multiple levels, the lack of an HPC tradition in biological research has hampered fully parallelized deployments.

Here we present a lightweight flexible framework for performing parallelized processing of raw gene expression data. The framework uses a Python3 based frontend for specifying analysis options, data paths and reference datasets. This frontend sanitizes and resolves the options, providing verbose error checking before writing a human readable configuration file and basic scripts for batch submission. The submission scripts leverage the scheduler to implement a scatter-gather approach, submitting potentially hundreds of individual jobs via a job array, each small enough to take advantage of backfill in a high contention HPC environment. The gather component is handled through a script submitted with an 'after-okay' dependency.

pdf

How the Science Gateways Community Institute Supports Those Who Are Creating Websites to Access Shared Resources
Katherine Lawrence and Nayiri Mullinix (University of Michigan), Maytal Dahan (TACC), Linda Hayden (Elizabeth City State University), Marlon Pierce (Indiana University), Nancy Wilkins-Diehr (San Diego Supercomputer Center), and Michael Zentner (Purdue University)
Abstract

The Science Gateways Community Institute (SGCI) serves the gateway community with free, NSF-funded resources, services, experts, and ideas. Our community shares expertise, technologies, and best practices for creating and sustaining gateways. Gateways offer Web-based access to shared resources, including data, software, computing services, instruments, and educational materials. Gateways are also known by other names, including portals, virtual research environments or virtual labs, eScience, and research cyberinfrastructure.

pdf

SSR – A Searchable, Sharable, and Reproducible Imaging Informatics Environment for Biomedical Image Storage, Annotation, and Analysis
Christian Suloway, Tianyi Miao, and Yanling Liu (Advanced Biomedical Computational Science, Frederick National Laboratory for Cancer Research)
Abstract

Application of machine learning (ML) and artificial intelligence to biomedical images has huge potential to improve treatment of cancer and other diseases. Classification and segmentation of histopathology and radiology images is critical to diagnosing cancer in clinical settings and understanding cancer in research. Yet one great barrier to applying ML to this problem is the quality and quantity of annotated data for training models. Deriving informatics from images requires quantification and thus segmentation, a time consuming and laborious process for subject-matter experts. We have created the Searchable, Sharable, and Reproducible (SSR) Imaging Informatics environment, a research computing platform, for managing, annotating, and analyzing biomedical images through an intuitive and accessible web interface to overcome this obstacle. Integrated annotation and analysis workflows in SSR allow for updates to manual annotations with model annotation predictions in situ. SSR is based on the established, open-source data platform Girder (Kitware) with data models stored in MongoDB and files on local file systems, HDFS, and Amazon S3. In-browser, JavaScript-based gigapixel image and volume viewing and pixel-level annotation tools synchronize with servers through RESTful APIs. Distributed analysis of the annotated data is managed by Celery task queues capable of launching Docker container instances for ease of portability on high-performance computing clusters and utilization of current imaging processing and ML libraries with GPU optimizations. SSR provides easy to use, web-based pixel annotation tools combined with integrated ML image segmentation to greatly enhance production and management of training data for ML models.

Funded by NCI Contract No. HHSN261200800001E

Towards an open and reusable computing infrastructure for Minnesota Earth observation data
Jeffery Thompson (Minnesota Supercomputing Institute, University of Minnesota) and David Porter (Minnesota Supercomputing Institute, Univesity of Minnesota)
Abstract

Earth observation satellites (EOS) have become indispensable for studying Earth system processes. While the challenges of extracting information from EOS data in general are often discussed in relation to the Big Data paradigms of variability and volume, an underappreciated facit of EOS data is their multipurpose nature, particularly optical imagery obtained from satellites like Landsat and Sentinel-2. Because they record observations across a number of discrete portions of the electromagnetic spectrum from all objects within their field of view, even a single image has relevance to any number of different domain science research questions. The value of the 40+ year observational record from Landsat missions in particular represents an underexplored resource for many domain science questions and the effective utilization of this record requires new tools and open HPC computing environments.

In this paper, we discuss efforts at the Minnesota Supercomputing Institute to establish a framework that enable EOS data reuse and exploration. We describe the tools and cyberinfrastructure developed to obtain Landsat imagery for Minnesota. These tools document data provenance, allowing for fully reproducible data products. In addition, we discuss the prototype python programming environment and workflows that enables Landsat data exploration and reuse and facilitate computation within MSI's HPC environment. Future needs and vision for creating an Open Data Cube environment for Minnesota are also discussed.

Parallel Integro-Differential MHD FEM FMM Solver
Keith Brauss (Francis Marion University)
Abstract

We numerically approximate the solution to the resistive magnetohydrodynamics (MHD) equations in a velocity-current, integro-differential formulation. The linearized and discretized MHD system is solved using the finite element method (FEM) and the freely available, open-source, differential equation software library deal.ii (capable of scaling to over hundreds-of-thousands of processors). The MHD system's magnetic field integral equation is approximated by incorporating a parallel, multi-level fast multipole method (MLFMM) into the deal.ii open-source library. The resulting MHD algebraic system is solved using the Krylov subspace method GMRES and a Schur-complement preconditioner. Examples are simulated to assess parallelization and convergence rates.

The Digital Object Architecture and Enhanced Robust Persistent Identification of Data
Rob Quick (Indiana Unversity), Larry Lannom (CNRI), and Marina Krenz and Yu Luo (Indiana Unversity)
Abstract

The expansion of the research community's ability to collect and store data has grown much more rapidly than its ability to catalog, make accessible, and make use of data. Recent initiatives in Open Science and Open Data have attempted to address the problems of making data discoverable, accessible and re-usable at internet scales. The Enhanced Robust Persistent Identification of Data (E-RPID) project's goal is to address these deficiencies and enable options for data interoperability and reusability in the current research data landscape by utilizing Persistent Identifiers (PIDs) and a kernel of state information available with PID resolution. To do this requires integrating a set of preexisting software systems along with a small set of newly developed software solutions. The combination of these software components and the core principles of making data FAIR (findable, accessible, interoperable and reusable) will allow us to use Persistent Identifiers to create an end-to-end fabric capable of realizing the Digital Object Architecture for researchers. This paper will introduce the audience to the concepts of the Digital Object Architecture, describe the software services necessary to enable this architecture, introduce the existing E-RPID testbed that is available for experimental usage on the Jetstream cloud environment, and describe the diverse set of use cases already using E-RPID to enhance their data accessibility, interoperability and reusability.

pdf

Characterizing Performance Improvement of GPUs
Dodi Heryadi (University of Notre Dame) and Scott Hampton (University Of Notre Dame)
Abstract

Evaluating supercomputer performance through benchmarking is an integral part of an HPC center operation. It can be used to guide the selection of a system during hardware acquisition, to validate that the selected hardware performs as expected, and to ensure that the system performs reasonably throughout its lifetime. In this work we quantify the performance improvement of our recent GPU architectures, GeForce RTX 2080 Ti, Titan V, and Tesla V100 over an older architecture, GeForce GTX 1080 Ti. We perform a workload analysis to select a suite of applications to benchmark. Using the benchmark results, a single performance number is then assigned to each GPU system under test. Taking hardware price into consideration, our results show consumer-grade GeForce RTX 2080 Ti as the most cost competitive GPU platform.

pdf

A Comprehensive Campus-based Approach to Address the Opportunities and Challenges Posed by Data Intensive Research
James Wilgenbusch and Joshua Baller (University of Minnesota, Minnesota Supercomputing Institute); Carla Bates (University of Minnesota, Psychology); Evan Bollig (University of Minnesota, Minnesota Supercomputing Institute); Lisa Johnston (University of Minnesota); and Claudia Neuhauser (University of Houston)
Abstract

In 2007 Jim Gray distinguished data intensive science as a fourth paradigm of scientific exploration. This declaration was made with a sense of both optimism and caution: optimism in anticipation of the novel breakthroughs and insights that these new sources and volumes of data would likely bring and caution because, as Gray described at the time, "the tools for capturing data both at the mega-scale and at the milli-scale are just dreadful." Over a decade later and we can confirm that the enthusiasm and apprehension expressed by Gray and others were warranted. Not only have the horizons of scientific discovery been expanded by the vast and diverse quantities of new data, but these data have also challenged our ability to sustain research cyberinfrastructure at universities in new and unexpected ways. These challenges are not strictly technical, but also involve social, cultural, legal, and financial components that can impede and sometimes completely inhibit our ability to conduct meaningful research. In this paper we will describe a comprehensive framework used by the University of Minnesota (UMN) to better address current and emerging challenges and opportunities brought by data intensive research and the process that our University used to develop this framework. Elements of this framework and the process used to develop it could be applied at other research institutions. While approaches may differ slightly, addressing these challenges within our respective universities is critical, and perhaps a prerequisite to building and sustaining partnerships among providers of advanced research computing.

pdf

Beyond Campus Bridging: A retrospective of Cyberinfrastructure Integration Efforts
John Coulter (Indiana University, XSEDE); Richard Knepper, Resa Reynolds, and Jodie Sprouse (Cornell University, XSEDE); and Stephen Bird (Indiana University, XSEDE)
Abstract

Over the past seven years, the XSEDE Cyberin- frastructure Resource Integration (XCRI) Team has occupied a unique space in the national research computing ecosystem. Aiming to close the loop between large-scale resources of na- tional supercomputing centers and campuses working to provide researchers with local resources that will help them scale up their analyses to larger systems. XCRI continues to offer a specialized set of tools and services that support campuses and research computing centers exploring and implementing their own resources. These services assist in the management of research cyberinfrastructure systems, provide a basis for future funding requests, and attract researchers who will leverage local and national systems to empower novel and innovative research. XCRI has helped shepherd several schools on the path to research computing, some from nothing more than a pile of hardware, and continues to provide services running the gamut from debugging new hardware to training administrators of new compute resources. Thanks to a broad mandate to explore and promulgate best practices and extend capabilities of campus cyberinfrastructure, XCRI has been successful in offering aid to resource-constrained institutions at several different levels. This paper describes the mission and activity of the XCRI team, its progress made in helping campuses adapt national- scale cyberinfrastructure software and best practices for their own environment to increase access to and availability of com- putational resources. This paper will also discuss organizational practices that have helped the team pursue their mission.

pdf

Employing Augmented Reality for Cybersecurity Operations in High Performance Computing Environments
Nitin Sukhija and Cory Haser (Slippery Rock University of Pennsylvania) and Elizabeth Bautista (Lawrence Berkeley National Laboratory)
Abstract

With the advent of myriad different technologies in the modern computing era, the future of high performance computing (HPC) data centers comprising heterogeneous compute servers and networks involves a significantly greater degree of cyber-threats than we are observing currently. As a result, conducting cybersecurity operations at the HPC data centers encompasses a highly cognitive workload environment, where information is rapidly generated, is exceedingly dynamic and is vastly unpredictable. An important challenge faced by cybersecurity specialists is managing the extreme levels of stress and utilizing emerging technologies for monitoring, assessing and defending multiple cyber-attacks in real-time. In this paper, we present an augmented reality based framework for storing, analyzing, visualizing and augmenting data at the HPC centers. This framework is proposed with the goal of improving the data presentation, forecasting, analysis and decision making, thus reducing the time to insight and learning the HPC security operations to combat increasing cyber physical threats.

pdf

Accelerated Computing with GenASiS
Reuben Budiardja (Oak Ridge National Laboratory, University of Tennessee)
Abstract

GenASiS (General Astrophysical Simulation System) is framework to facilitate simulation of astrophysical phenomena on leading capability supercomputers. We describe recent efforts to enable GenASIS to use accelerated computing technology such as GPUs. Using OpenMP 4.5 target directive, we demonstrate simplified access to GPU programming capable of significant speedups.

Augmenting High Performance Computing Support staff with University Student Workers
Joel Frahm, Shelley Knuth, Andrew Monaghan, and Daniel Trahan (University of Colorado)
Abstract

The University of Colorado Boulder Research Computing (CURC) supports High Performance and other research focused computational resources in partnership with Colorado State University and the Rocky Mountain Advanced Computing Consortium (RMACC). The University of Colorado Boulder Research Computing (CURC) supports High Performance and other research focused computational resources in partnership with Colorado State University and the Rocky Mountain Advanced Computing Consortium (RMACC).

With CURC receiving numerous help requests (support cases) for common, well-documented issues and questions regarding regular scheduled maintenance or other downtime, an effort was made train them to be able to handle some of the caseload, reducing demands on full-time staff.

It was determined that at a manageable scale (2-3 students) it does not seem realistic to have consistent low-latency help with common support requests during the normal workday. Student schedules also proved to require a higher level of flexibility than that of the full-time staff, requiring more time off or schedule changes to account for exams, study sessions and other demands of student life. However, there are considerable long-term benefits to developing and maintaining productive student employees in support roles. Students will continue to augment the CURC workforce and we will continue to develop effective training and utilization strategies going forward.

A Multi-Layer Classifier for Hyperspectral Images
Luis Cueva Parra (University of North Georgia) and Semih Dinc and Babak Rahbarinia (Auburn University at Montgomery)
Abstract

Different from regular RGB images that only store red, green, and blue band values for each pixel, hyperspectral images are rich with information from the large portion of the spectrum, storing numerous spectral band values within each pixel. An efficient, two-layer region de- tection framework for hyperspectral images is introduced in this paper. The proposed framework aims to automatically identify various regions within a hyperspectral image by providing a classification for each pixel of the image, associating them to distinct regions. The first layer of the system includes two new classifiers, and is responsible for generating probability scores as the "new feature set" of the original dataset. The second layer works as an ensemble classifier and combines the newly gen- erated features to estimate the region of the sample. Experimental results show that the proposed system can produce accurate classifications with an average area under the ROC curve of 0.98 over all regions. This result indicates the higher accuracy of the proposed system compared to some other well-known classifiers. Also, the speedup of the proposed classifier performs very well compared with more sophisticated methods.

pdf

Automated network provisioning and orchestration for vendor agnostic networks
Kapil Agrawal (National Center for Supercomputing Applications)
Abstract

Manual provisioning and configuration management for network devices in a multi-vendor network environment can be a tedious and error prone job. When running network systems at scale, copying and pasting configurations all over the place is certainly not an option specially in a multi-vendor network. This session/poster will highlight how network engineers can abstract their network infrastructure as a code using a data modeling strategy and use network automation tools to auto provision devices in a vendor agnostic manner. This session will also talk about the recent efforts NCSA has been making towards vendor neutral Zero touch provisioning and challenges facing towards that

InterACTWEL Science Gateway for Adaptation Planning in Food-Energy-Water Sectors of Local Communities
Meghna Babbar-Sebens and Samuel Rivera (Oregon State University); Eroma Abeysinghe, Suresh Marru, Marlon Pierce, and John Eric Coulter (Indiana University); Majid Farahani (Oregon State University); Dimuthu Wannipurage (Indiana University); Akila Ishani (Consultant); and Marcus Christie (Indiana University)
Abstract

Since their inception in mid 2000s, adoption of Science Gateways as interfaces and conduits for digital infrastructure needed in science and engineering research and education has significantly increased. This trend has also driven changes in the types of services and resources that are now being expected from the Science Gateways by a growing group of diverse end users. In this poster, we present a novel Science Gateway, InterACTWEL (Interactive Adaptation and Collaboration Tool for managing Water, Energy and Land), which serves as a research cyberinfrastructure as well as an applied decision support system for adaptive natural resources management in interdependent food, energy, and water sectors. End users of this gateway include not only interdisciplinary technical and social science researchers, but also public and private sectoral stakeholders. The gateway is a collaboration between Oregon State University and Science Gateways Research Center, Pervasive Technology Institute at Indiana University, and is addressing challenges and solutions related to computational services, visualization techniques, advanced software applications, collaboration capabilities, cyber security and privacy, and data repositories unique to food-energy-water sectors and their stakeholders.

pdf

Community Clusters or the Cloud: Continuing cost assessment of on-premises and cloud HPC in academia
Preston Smith, Stephen Harrell, Alexander Younts, and Xiao Zhu (Purdue University)
Abstract

Understanding total cost of ownership (TCO) for HPC systems on a site-specific range of applications is a critical component of any RFP or capacity plan at research computing centers. With the increasing industry use of the HPC in the cloud, understanding and comparing costs of on-premises HPC resources and cloud-based options is more important than ever. In this paper we present methods to determine site-specific applications mix, calculating TCO for on-premises and cloud HPC resources, and the comparison of the two. Additionally we present a cost-analysis of Purdue's Community Cluster program with current cloud offerings.

pdf

XSEDE Integration with ORCID for Research Resources
David Hart (NCAR), Ester Soriano (NCSA), Carrie Arnold (TACC/UT Austin), Steven Peckins (NCSA), Rob Light and Burt Cubbison (PSC), Josh Berger (NCAR), and Michael Shapiro (NCSA)
Abstract

Many research infrastructure providers face the challenge of identifying the science output of their user communities. The XSEDE Resource Allocations Service (XRAS) has recently completed integration with the ORCID ecosystem, allowing XSEDE to post allocation award information to the ORCID profiles of its users. Through these ORCID-posted awards, publishers will be able to link users' papers to the XSEDE-allocated projects that supported the research. To support the integration, XSEDE staff enhanced the user portal and XRAS and performed some novel administrative steps. Other XRAS client organizations can take advantage of these features to connect their allocation awards to their users' ORCID profiles.

pdf

Shared Memory High-Throughput Computing with Apache Arrow (TM)
Geoffrey Lentner (Purdue University)
Abstract

As the barriers to entry in scientific computing have lowered with languages like Python and their libraries, the demand for ever more sophisticated and capable frameworks that provide almost turn-key functionality has grown proportionally (see keras [1]).

For researchers who use the pilot job [2], many-task [3] design pattern for high-throughput computing [4] on modern systems, the Apache Arrow project (Apache Software Foundation, 2019) with its now included Plasma in-memory object store provides a high-level interface for sharing data structures between processes in a way that requires no serialization or copying of that data.

This poster outlines a common scenario in which a constraint is induced by the way memory is managed. The direct sharing of memory by a common middle data layer allows for accelerated workflows that can operate at a more rapid cadence. A real-world example is provided; the goal is to raise awareness in facilitators at other major research computing centers who work with users to architect similar such high-throughput data pipelines.

pdf

Breaking Resource Discovery Barriers
John-Paul Navarro (University of Chicago) and Amy Hovious (University of Illlinois at Urbana-Champaign)
Abstract

Researchers, educators, and students rely on online information technology resources and services, including software, computing/storage/instrument resources, online and offline training, and consulting and technical services, to name a few. Discovering technology resources across organizational boundaries can be challenging because organizations generally only facilitate discovery of their own resources, because general purpose search engines are too generic, and because resource information may reside in database and not be indexed by general purpose search engines. Starting in 2017, University of Illinois Technology Services' Research IT Team set out to create a single point of entry via the Research IT Portal to identify, gain access to use and access to support for Research IT resources and services. One of the biggest hurdles faced during the design process was how to make the wide variety of resources and services discoverable by the Research IT Portal.Recognizing that the XSEDE program had a similar challenge to enable resources, software and services discovery and that it had developed Information Services to enable aggregation, indexing, and programmatic cross-institutional discovery, the University of Illinois and XSEDE decided to collaborate and build a joint solution. This paper presents a new approach developed through a collaboration between the University of Illinois Technology Services and the XSEDE Cyberinfrastructure Integration teams to enable discovery and use of technology resources across departments, campuses, and institutional boundaries to facilitate research and educational activities.

pdf

Navigating an ocean of computational environments: understanding researcher needs in selecting and building virtual machine and container images
Amy Neeser (UC Berkeley) and Anna Sackmann and Steve Masover (University of California, Berkeley)
Abstract

Many research computing services are seeing rapid growth in the use of virtual machine images (VM images) and container images to support research computation, applications, and workflows (Belmann, 2015). These images are used across a range of compute models. Researcher interest in VMs and containers is driven by a wide range of considerations, including need to provision and run VM images when using cloud resources; portability of environments to multiple computational resources or platforms on which researchers work; efficiency in reproducing and sharing environments that are proven and configured for a researcher's workflow; and citability in descriptions of research methodology. Researchers can be significantly supported by a maintenance regime that minimizes time and attention they must pay OS and application maintenance tasks, and by mechanisms that ease discovery of images best suited to their need. From a researcher's perspective, an initial step in their research workflow is the discovery and tailored configuration of reliable, preferably vetted images, from which to select an appropriate fully-functional or substrate image for their computational research (by "substrate image" we mean an image to which project-specific software can be added with reasonable ease).

To begin to understand how researchers undertake the discovery of such images, a small group of UC Berkeley researchers were interviewed to elicit an overview of their computational research needs and practices, as well as past or prospective characteristics or features they do or expect to look for when choosing or augmenting a VM or container in which to run computational research.

pdf

Performance Characterization of Single and Multi GPU Training of U-Net Architecture for Medical Image Segmentation Tasks
Trupesh Patel, Sandeep Bodduluri, Thomas Anthony, William Monroe, Pravinkumar Kandhare, John-Paul Robinson, Arie Nakhmani, Chengcui Zhang, Surya Bhatt, and Purushotham Bangalore (University of Alabama at Birmingham)
Abstract

ABSTRACT: The rapid advancements in high performance computing hardware and corresponding rise in deep convolutional neural network (CNN) architectures have led to state-of-the-art results in several biomedical image segmentation tasks. Recently, U-Net, a modified fully convolutional network, has become the state-of-the-art in various two-dimensional and three-dimensional semantic (pixel-level) segmentation tasks related to medicine. U-Net has achieved most success in terms of dealing with datasets where the well annotated ground truth is beyond reach. However, there has not been a detailed analysis on the impact of computing configurations, input data type, preprocessing, and data augmentation techniques on U-Nets training speed and end-to-end computation pipeline. In this paper, we trained U-Net on multiple configurations of GPU (single vs. multi-node) hardware in terms of memory and time efficiency towards finding optimal set up for U-Net segmentation tasks.

pdf

Improving a Genetic Algorithm for Route Planning Using Parallelism with Speculative Execution
David Mathias and Samantha Foley (University of Wisconsin - La Crosse)
Abstract

A typical United Parcel Service (UPS) route consists of approximately 100 stops. The distance traveled to visit those 100 stops is determined by the order in which they are visited and the paths traveled between consecutive stops. This is an example of a problem we call route planning. The reward for finding short routes is substantial. UPS estimates that reducing each of their 55,000 North American routes by only 1 mile per day results in annual savings of up to $50,000,000, in fuel costs alone. We approach the route planning problem with a novel, parallel genetic algorithm. We demonstrate the utility of parallelism for route planning by showing decreasing run times and solving much more complex problem instances that a non-parallel implementation of the same algorithm. By introducing speculative execution to the algorithm, we further, significantly improve the results.

pdf

Research Computing at a Business University
Jason Wells (Bentley University) and J. Eric Coulter (Indiana University)
Abstract

Research Computing demands are expanding beyond traditional disciplines due to the proliferation of data in all walks of life. At Bentley University ("Bentley"), a business university in the Boston area primarily funded through tuition, this expansion has been most readily seen in our Accounting, Economics, Mathematics, and Natural Sciences departments (non-traditional Research Computing disciplines). The result has been a small effort to build a research computing capability at this small New England university. This poster will serve as an overview of the steps taken to build such an effort at a business university, the revelations we have had along the way, and our resulting plans. It is hoped that this poster can thus: (1) serve as a template for other efforts in the non-Traditional discipline and Long Tail of Science communities; (2) become a starting point for a wider discussion about the unique needs of these communities; and (3) simply document what worked, at least at this business university.

pdf

Empowering Women in HPC careers through participation in the WINS apprenticeship
Maria Kalyvaki (South Dakota State University) and Dori Sajdak (University of Buffalo)
Abstract

The purpose of this poster is to explore the benefits for young professional women in male-dominated careers participating in HPC apprenticeships, focusing on those participating in the WINS apprenticeship model. Women are particularly underrepresented in the Information Technology field, an issue of concern both in terms of increased economic participation of women, and gender equality more broadly. It is known that the foundations for transition from education and training to employment are usually established during the school years. During those years, men and women have reinforced notions of what is possible and not possible for them. Unfortunately, gendered stereotypes and perceptions around certain career options are still reinforced within schools and create barriers to widening young women's participation in a range of careers, particularly in fields traditionally dominated by males. The poster discusses the outcomes of participating in the WINS HPC apprenticeship, and the experiences of the past WINS awardees.

Improving regional cyberinfrastructure services through collaboration: Cyberteam for the Rocky Mountain Advanced Computing Consortium
Andrew Monaghan (University of Colorado Boulder); Brett Milash (University of Utah); Tobin Magle (University of Wisconsin Madison); Thomas Hauser (University of Colorado Boulder); Thomas Cheatham (University of Utah); Patrick Burns (Colorado State University); Shelley Knuth (University of Colorado Boulder); Anita Orendt (University of Utah); Suzi White and Dawn Paschal (Colorado State University); H.J. Siegel (Department of Electrical and Computer Engineering, Colorado State University); Timothy Kaiser (Academic Network and Computing Resources, Colorado State University); and Andrew Johnson and Becky Yeager (University of Colorado Boulder)
Abstract

Advances in the volume, diversity, and complexity of research data and associated workflows requires enhanced capabilities to access, secure, reuse, process, analyze, understand, curate, share, and preserve data. To address this need at the regional level the University of Colorado, Colorado State University, and the University of Utah formed a "Cyberteam" in 2017 to provide cyberinfrastructure (CI) support to researchers at institutions in the Rocky Mountain Advanced Computing Consortium (RMACC) encompassing states across the Intermountain West. The Cyberteam is comprised of CI professionals across the three institutions who collaborate closely, sharing expertise and resources. Since its establishment, the Cyberteam has worked to broaden accessibility and options for computing, storage, and data publishing for RMACC researchers; enhance training on data- and workflow-oriented topics; improve engagement with researchers using CI; and better understand user needs and challenges. One key accomplishment has been the development of a series of focus group and survey instruments to achieve better understanding the CI needs and challenges of researchers across a diverse spectrum of disciplines. This paper provides an overview of the RMACC Cyberteam's objectives, accomplishments, challenges, and future direction.

pdf

Practical Cyber-Security Solutions for the Science DMZ
Peter Hawrylak (The University of Tulsa), George Louthan (University of Oklahoma), and John Hale and Mauricio Papa (The University of Tulsa)
Abstract

The Science DMZ provides high-bandwidth and high-throughput private connectivity to or among computing resources. However, the performance comes at a cost: there is minimal security on the Science DMZ. This is by design, as conventional enterprise network security controls introduce cost, latency and bottlenecks. Data intensive applications, particularly those using geographically distributed resources, rely on the performance of the Science DMZ. With a particular lens toward multi-institution DMZs such as the OneOklahoma Friction Free Network, this paper analyzes performance and security requirements, introduces a threat model specific to the Science DMZ, and delivers recommendations for controls to mitigate these threats in the context of the identified requirements. In particular, we apply a risk-based view to the tradeoffs between performance and security, considering impacts both to the science goals of the DMZ's multiple research stakeholders and the more conventional enterprise security posture of the facilities it connects.

pdf

Slicing and Dicing OpenHPC Infrastructure: Virtual Clusters in OpenStack
Satrio Husodo, Jacob Chappell, Vikram Gazula, Lowell Pike, and James Griffioen (University of Kentucky)
Abstract

The OpenHPC system provides a set of tools for setting up an HPC cluster. It has been gaining wide-spread adoption and support. It is best-suited for native bare-metal installations across homogeneous nodes; however, some groups have attempted to integrated OpenHPC with OpenStack Ironic to get the benefits of cloud elasticity and bursting, which essentially uses Ironic to create the same homogeneous bare-metal environments used in native OpenHPC installs.

Here, we explore the use of OpenStack VMs (not Ironic bare-metal nodes) as a platform to build a virtual cluster that runs OpenHPC tools. VMs allow administrators more flexibility in defining the characteristics of a cluster's nodes. This is important as systems obtain larger and larger nodes whose resources may not be used effectively unless they can be virtualized (sliced and diced) to the specific needs of applications.

We discuss how the head node/VM can be used to provide the base image for the compute VMs through PXE-booting. As a result, the head node supports the necessary tools needed to build the compute VM image and provision the compute VMs, as opposed to pre-building images outside the optimized OpenHPC workflow using OpenStack's image mechanisms.

Using our approach, we can build virtual OpenHPC clusters dynamically, allowing users to submit jobs using Slurm just like a bare-metal implementation. Our approach allows users to create virtual clusters on demand without adding complexity in setting up OS provisioning as it uses the conventional OpenHPC packages and allows non-HPC workloads to co-exist on the same hardware.

pdf

A Policy-Driven License Manager for Controlling Access to Arbitrary Applications
Adam Howard and Gary Rogers (University of Tennessee)
Abstract

Seeking to fill a gap in licensed software access control solutions, a license management solution based on administrator-controllable policies is presented. This solution, Lisense, is a client-server style license manager that can be applied to any application binary without the need to alter source code or recompile the application. While Lisense provides several builtin policies like controlling number of concurrent users running a wrapped application simultaneously, it also allows definition of new, external policies via a simple Python API, allowing a high level of flexibility to the administrator in designing and developing access control systems beyond those already available. In this paper, we discuss the implementation of Lisense, as well as an example deployment, and place it in context of other access control solutions available.

pdf

The USD Science Gateway: A Bridge Between Research and Advanced Computing
Adison Kleinsasser and Douglas Jennewein (University of South Dakota); Joseph Madison (Mayo Clinic); and Eroma Abeysinghe, Sudhakar Pamidighantam, Marcus Christie, Suresh Marru, and Marlon Pierce (Indiana University)
Abstract

Science Gateways are virtual environments that accelerate scientific discovery by enabling scientific communities to more easily and effectively utilize distributed computing and data resources. Successful Science Gateways provide access to sophisticated and powerful resources, while shielding their users from the underlying complexities. Here we present updated work completed by the University of South Dakota (USD) Research Computing Group in conjunction with the Science Gateways Community Institute (SGCI) [1] and Science Gateways Research Center at Indiana University to set up a Science Gateway to access USD's high-performance computing resources. We also introduce improvements to the system since the previous presentation of our work. These resources are now available to both faculty and students and allow ease of access and use of USD's distributed computing and data resources. The implementation of this gateway project has been multifaceted and has included placement of federated user login, user facilitation and outreach, and integration of USD's cyberinfrastructure resources. We present this project as an example for other research computing groups so that they may learn from our successes and the challenges that we have overcome in providing this user resource. Additionally, this project serves to exemplify the importance of creating a broad user base of research computing infrastructure through the development of alternative user interfaces such as Science Gateways.

pdf

Scaling R Shiny Apps to Multiple Concurrent Users in a Secured HPC Environment Using Open OnDemand
Eric Franz (Ohio Supercomputer Center); Venkat Gadepalli, Maciej Pietrzak, Hatice Ozer, and Amy Webb (Biomedical Informatics Shared Resource, The Ohio State University); and Morgan Rodgers, Trey Dockendorf, Alan Chalker, Doug Johnson, and David Hudak (Ohio Supercomputer Center)
Abstract

Open OnDemand is an open source project to provide web based access to HPC resources (https://openondemand.org). OnDemand's "Interactive Apps" support launching web applications like Jupyter and RStudio on cluster compute nodes, proxying HTTP requests to those apps, and secure user separation of both app and data access enforced at the OS level. This paper describes leveraging OnDemand to enable the deployment of R Shiny applications that can be launched in OnDemand's interactive HPC environment by multiple analysts at the same time to visualize different datasets. Dataset and app authorization is managed using network file system (NFS) file access lists (ACLs). A custom OnDemand app was built to enable admins to manage who has access to specific apps and what datasets can be loaded by those apps. By using OnDemand to facilitate the use of an HPC batch scheduler as the process manager and OnDemand as the proxy we can support multiple analysts launching Shiny apps to access their data in a scalable and secure manner without the need for managing a separate dedicated installation of RStudio or Shiny Server.

pdf

SciTokens: Demonstrating Capability-Based Access to Remote Scientific Data using HTCondor
Alex Withers (NCSA), Brian Bockelman (Morgridge Institute for Research), Derek Weitzel (University of Nebraska-Lincoln), Duncan Brown (Syracuse University), Jason Patton (University of Wisconsin-Madison), Jeff Gaynor and Jim Basney (NCSA), Todd Tannenbaum (University of Wisconsin-Madison), You Alex Gao (University of Illinois), and Zach Miller (University of Wisconsin-Madison)
Abstract

The management of security credentials (e.g., passwords, secret keys) for computational science workflows is a burden for scientists and information security officers. Problems with credentials (e.g., expiration, privilege mismatch) cause workflows to fail to fetch needed input data or store valuable scientific results, distracting scientists from their research by requiring them to diagnose the problems, re-run their computations, and wait longer for their results. SciTokens introduces a capabilities-based authorization infrastructure for distributed scientific computing, to help scientists manage their security credentials more reliably and securely. SciTokens uses IETF-standard OAuth JSON Web Tokens for capability-based secure access to remote scientific data. These access tokens convey the specific authorizations needed by the workflows, rather than general-purpose authentication impersonation credentials, to address the risks of scientific workflows running on distributed infrastructure including NSF resources (e.g., LIGO Data Grid, Open Science Grid, XSEDE) and public clouds (e.g., Amazon Web Services, Google Cloud, Microsoft Azure). By improving the interoperability and security of scientific workflows, SciTokens 1) enables use of distributed computing for scientific domains that require greater data protection and 2) enables use of more widely distributed computing resources by reducing the risk of credential abuse on remote systems.

In this extended abstract, we present the results over the past year of our open source implementation of the SciTokens model and its deployment in the Open Science Grid, including new OAuth support added in the HTCondor 8.8 release series.

pdf

LSU Computational System Biology Gateway for Education
Eroma Abeysinghe (Indiana University), Michal Brylinski (Louisiana State University), Marcus Christie (indiana University), and Suresh Marru and Marlon Pierce (Indiana University)
Abstract

Science gateways are a mechanism for delivering scientific software as a service, especially when the software requires high performance computing (HPC) resources to run effectively. The existence of a science gateway eliminates the user's need to learn to work with HPC systems and to manage software installations and updates. With well-designed user interfaces, users can more quickly become effective users of scientific applications and can manage information needed for replicating, modifying, and sharing results. All of these efficiency gains enable users to focus more on their research. In addition, science gateways are being identified as an effective educational tool, a tool to be used in classroom environments as a method to get students quickly into research on domain specific questions. In the absence of a science gateway, students are likely to need a considerable time to learn to work with HPC systems, and any time spent on such will reduce their time on the actual science. This poster presents how the Louisiana State University (LSU) gateway for the Computational System Biology Group (CSBG) - (www.brylinski.org) was updated and improved to be a classroom teaching tool. This work makes extensive use of Apache Airavata's group management capabilities.

pdf

From BigDog to BigDawg: Transitioning an HPC Cluster for Sustainability
Chet Langin (Southern Illinois University Carbondale)
Abstract

This paper relates the experiences of managing the transition of a high performance computing (HPC) cluster from Rocks, SGE, and Cisco to OpenHPC, SLURM, and Dell. This transition was made because of sustainability issues related to security, the software, and the hardware. The BigDog HPC cluster was placed in production at Southern Illinois University Carbondale (SIUC) December, 2015, with 40 Cisco servers containing 800 CPU cores running Rocks management software and the SGE job scheduler. Sustainability issues of BigDog were encountered, especially related to keeping software updated for security reasons, and the decision was made to update the cluster for sustainability. This was done by replacing some Cisco servers with Dell hardware, and replacing Rocks with OpenHPC, and replacing SGE with SLURM. The name of the cluster was changed to BigDawg in keeping with the southern pronunciation and the spelling of the names of the university mascots Grey Dawg and Brown Dawg.

pdf

Visualization and Interpretation of Cluster Usage and User Datasets
Chitrak Gupta, Lee Reynolds, Jason Yalim, and Gil Speyer (Arizona State University)
Abstract

In the effort to make data-driven decisions to support end users, research computing teams gather various usage statistics about compute resource requests, runtimes and software loads. Auxiliary user data sources, such as department and domain information can be integrated to make inferences about shared features across the customer base.

This project explores various methods of visualizing and processing cluster workloads and user data. Datasets reporting the usage of the Agave cluster at Arizona State University are mined and interrogated, employing statistical and graphical algorithms for analysis.

This work complements the PEARC "Facilitation of Advanced Research Computing" technical track with an effort to better understand and support the customers of research computing systems.

The Distant Reader - Tool for Reading
Eric Lease Morgan (University of Notre Dame) and Eroma Abeysinghe, Sudhkar Pamidighantam, Eric Coulter, Suresh Marru, and Marlon Pierce (Indiana University)
Abstract

The Distant Reader science gateway can be used to automatically create and analyze text corpora at a scale of thousands of user-supplied documents. These processing steps are deployed on a dynamic virtual cluster deployed on XSEDE's Jetstream academic cloud computing resource and are accessed through a Web interface. The science gateway uses Apache Airavata middleware to manage the interactions between the Web interface and the virtual clusters. The gateway leverages the Science Gateway Platform as a service (SciGaP) infrastructure at Indiana University, which provides user authentication, authorization, and identity management as well as access to the Distant Reader tools. The Distant Reader is designed to assist in the process of using & understanding corpora -- reading.

pdf

Geo-Grid: Interactive Visualization of Geological Datasets
Nitin Sukhija, Sonny Sevin, Jordan Walters, and Michael Zieg (Slippery Rock University of Pennsylvania)
Abstract

As the emergent technology trends are progressing towards extreme-scale HPC coupled with data intensive analytics, the challenge of storing and analyzing data is significantly greater than we are observing currently. Given the vast volume of data generated by various simulation and experiments, the process of successfully analyzing and validating diverse sets of research phenomena mandates execution of various complex queries, the correlation of the disparate data retrieved followed by generating visualizations facilitating meaningful insights of both the simulated processes and the experimental phenomena. This process of data discovery, analysis, and visualization requires execution of several phases by the researchers. Therefore, managing and generating data visualizations is a major bottleneck in the research discovery process that degrades not only the ability to mine the data but also limits the utilization of scientific data sets. An important challenge in exploring geospatial data is visualizing different scales ranging from micrometers to kilometers. The purpose of this study is to develop and implement a framework for storing, analyzing and visualizing geological datasets at different scales, with the goal of improving the scientific discovery process and reducing the time to insight in the geological domain along with hiding the intricacies of the underlying low-level complex hardware platforms from the end-user.

pdf

Streamlining Simulation and Data Transfer in a Heterogeneous Environment
Matthew Gardner, Ping Luo, Matthew Johnson, and Hamid Toliyat (Texas A&M University)
Abstract

Our research uses ANSYS Maxwell for finite element analysis of electromagnetic devices. Our studies use heavily parameterized models to analyze numerous design scenarios and operating points, resulting in hundreds and thousands of simulation cases. Consequently, we designed and implemented a novel workflow to automate the large number of simulations and minimize the amount of human involvement required in the process. The workflow integrates simulation and data transfer seamlessly in a heterogeneous working environment that consists of Windows workstations and Linux clusters. The two data transfer schemes – polling and autonomous data transfer – can also be applied to other workflows where data need to be transferred without user intervention in cloud computing.

pdf

Deep Learning Enabled Predicting Modeling of Mortality of Diabetes Mellitus Patients
Ian Wittler, Xinlian Liu, and Aijuan Dong (Hood College)
Abstract

Diabetes mellitus (DM) is a major public health concern that requires continuing medical care. It is also a leading cause of other serious health complications associated with longer hospital stay and increased mortality rates. The fluctuation of blood glucose levels is easy to monitor. Physicians manage patients' blood glucose to prevent or slow the progress of diabetes. In this paper, the MIMIC-III data set is used to develop and train multiple models that aim to predict the mortality of DM patients. Our deep learning model of convolutional neural network produced a 0.885 AUC score, above all baseline models we constructed, which include decision trees, random forests, and fully connected neural networks. The inputs for each model were comprised of admission type, age, Elixhausercomorbidity score, blood glucose measurements, and blood glucose range. The results obtained from these models are valuable for physicians, patients, and insurance companies. By analyzing the features that drive these models, care management for diabetic patients in an ICU setting can be improved resulting in lowered mortality rate.

pdf

Growing Researchers' Computational Skills to Meet Future Needs: A New Zealand Case Study
Georgina Rae (New Zealand eScience Infrastructure), M. Fabiana Kubke (University of Auckland), and Nick Jones (New Zealand eScience Infrastructure)
Abstract

Training is integral to how the New Zealand eScience Infrastructure (NeSI) supports New Zealand researchers with the work being driven by a clearly defined Training Strategy.

In late 2017, NeSI worked with Fabiana Kubke to review our achievements in training with a view to developing our future training strategy.

Our approach has been to quantify the outputs and estimate the impact of NeSI's training activities in proportion to overall sector capacity, and to make some inferences on the possible extent of need. We have done this through the following steps:

- Estimate workforce capacity for Higher Education Institutes and Crown Research Institutes - Map the use of NeSI infrastructure against these workforce estimates - Identify the reach and gaps in the current NeSI training initiatives

Having reviewed our efforts so far, we are now scoping the training that will be required to truly make a difference to the levels of digital literacy in New Zealand.

This poster will present our findings from these pieces of work.

Solving the Traveling Salesmen Problem with IBM Q Quantum Computers
Alfonso Guevara (UHCL, XSEDE EMPOWER) and Nicholas Kindla (University of Houston-Clear Lake)
Abstract

Quantum Computing (QC) is an exotic type of computational resource that offers a possibility of better solutions to problems where classical computers have failed to find. Such type of problems are the notorious NP-Complete problems which center around optimization and complex decision-making algorithms. A good example of this type problems is the Traveling Salesman Problems (TSP) which its intend is to find the shortest optimal path across cities and returning to the starting city. In this project we measure the efficiency of IBM quantum computers and their capability of finding an optimal path for the TSP against a regular classical computer.

Recommender Systems for Gene Expression Recovery in scRNA-seq Data
Elnaz Mirzaei Mehrabad (UNIVERSITY OF UTAH, Huntsman Cancer Institute); Aditya Bhaskara (UNIVERSITY OF UTAH); and Benjamin Spike (Huntsman Cancer Institute, UNIVERSITY OF UTAH)
Abstract

Single cell RNA sequencing (scRNA-seq) is a powerful gene expression profiling technique, presently revolutionizing the study of complex cellular systems and responses in the biological sciences. However, scRNA-seq approaches currently suffer from sub-optimal target recovery leading to many false negatives. The resulting inflation of null readings adds noise to data visualization and may confound its interpretation. Since cells represent coherent phenotypes defined by conserved molecular circuitries, and since these are encoded in multi-gene expression patterns, information about one node in the data is predicted to be embedded in other nodes of the data. Under this hypothesis, several approaches have been proposed to impute missing values by extracting information from non-zero values in the data set. As recommender systems have been widely used to make predictions from sparse data matrices, we hypothesized they could provide effective means to recover missing values in RNA-seq data. We applied machine learning (NMF) and collaborative filtering (KNN) recommender systems to scRNA-seq data to impute missing values. We compared these approaches to existing imputation approaches and found they produce significantly better results, and potentially uncover hidden features in the data.

pdf

Creating a Cluster Oriented Workflow for Users of NetLogo
Ben Burnett (Northeast Cyberteam, University of Massachusetts Dartmouth); Abigail Waters (Northeast Cyberteam, Suffolk University); Ben Levy (Fitchburg State University); and Julie Ma (MGHPCC, Northeast Cyberteam Initiative)
Abstract

As part of the Northeast Cyberteam Initiative we created a workflow for using the simulation software NetLogo in a Linux cluster environment that uses the SLURM workload manager. This workflow was oriented towards a user that was making the initial transition from a Windows desktop work environment to the Linux cluster environment. We explore tools like MobaXterm and Ask.CI that were used to assist in this transition and their effectiveness in providing the support necessary to create the confidence and self reliance necessary for future research endeavors.

Ranking Supercomputers using SPEC HPG Benchmarks
Mayara Gimenes, Sunita Chandrasekaran, and Rudolf Eigenmann (University of Delaware)
Abstract

In this poster, we present a framework to rank supercomputers using a benchmark suite, ACCEL v1.2 that consists of several mini-applications. These mini-applications uses programming standard models including OpenMP, OpenACC and OpenCL. The benchmark suite is developed by SPEC High Performance Group. The published results contain hardware and software details that we scrape through using web scraping tools in order to collect specific attributes. In addition, we build a database of these results and also create a GUI that consists of different types of filters to satisfy the user's needs. These filters also help perform a more refined search. Our focus is on the different types of metrics the user can use to sort and filter as a given HPC. We aim to improve this framework in due course of time to classify and rank HPC systems. This framework is the first step for future ranking of large scale supercomputers that will be evaluated using real-world applications.

Community Outreach and the Discovery HPC Cluster: An Analysis of User Profiles and Growth
Jelena Karapetrović, Edin Šabić, and Diana Dugas (New Mexico State University)
Abstract

High Performance Computing (HPC) clusters can serve as a catalyst for university research by providing powerful computational resources . Equilibrium between available computing resources and user demand is a prerequisite for an economical and efficient HPC cluster. While use of the New Mexico State University HPC cluster, "Discovery", is growing, there are still many users that do not know that the cluster exists. Moreover, researchers are often hesitant to start using Discovery, as learning how to use the command line or job scheduler can be intimidating. Inefficient or erroneous cluster operation is often rooted in a lack of user training, potentially dawdling staff due to troubleshooting and therefore further deterring users from using the HPC. The Discovery team raises campus awareness through outreach and offers timely guidance to prospective users of the system. Campus outreach aims to educate and attract new HPC users through a variety of methods. These include: advertising Discovery directly to faculty and graduate students, teaching undergraduate students how to use the HPC as part of their coursework, promoting the system at university technology demonstrations, and hosting campus workshops. The goal of the research is to assess impact of outreach activities on Discovery use. Further, the Discovery team seeks to better understand prospective users, both in terms of their areas of expertise and previous HPC knowledge. The research findings hold implications for shaping future outreach efforts and increasing campus HPC utilization.

pdf

Creating Effective STEM Outreach Experiences for High School Students: Development of a Molecular Modeling Workshop Series
Amanda Sharp and Anne Brown (Virginia Tech)
Abstract

STEM-based careers are growing exponentially, which drives an increased need for students to pursue degrees in this field. While student interest in the STEM field is increasing, studies support that high school students are not prepared for college-level science and mathematics courses. Coding and computation skills are in-demand and require further training and experience. This motivates investigation into effective methods for high school level STEM development and outreach activities that integrate these domains. To meet these STEM and workforce-based needs, we have created the summer experience, Experience in Molecule Modeling and Informatics (EMMI), to combine interests in STEM and computation into a hands-on, applied environment. It is necessary to combine the experimental learning theory (ELT) and problem-based learning (PBL) methods to provide strategic and reflective programming. We hypothesize that utilizing students personalized system for learning will be enhanced in scientific and computer science education by employing ELT and PBL techniques. Framework and assessment designs will be shared, as well as processes used to introduce coding, computational biology, and bioinformatics, to high school students. The significance of this study will show how learning in a research-based environment employing computational techniques can enrich learning about STEM topics. Results will be used to develop a better outreach program for high school students based on experimental and PBL techniques. This allows for outreach assessment and strategic PBL design for a more meaningful program, in addition to providing students with domain-specific projects that will enhance their skills in data and computation-heavy processes.

Parallel Framework for Data-Intensive Computing with XSEDE
Ranjini Subramanian and Hui Zhang (University of Louisville)
Abstract

With the increase in data-driven analytics, the demand for high performing computing resources has risen. There are many high-performance computing centers providing cyberinfrastructure (CI) for academic research. However, there exists access barriers in bringing these resources to a broad range of users. Users who are new to data analytics field are not yet equipped to take advantage of the tools offered by CI. In this project, we propose a framework to lower the access barriers that exist in bringing the high-performance computing resources to users that do not have the training to utilize the capability of CI. The framework uses divide-and-conquer (DC) paradigm for data-intensive computing tasks. It consists of three layers – user interface (UI), parallel scripts generator (PSG) and underlying cyberinfrastructure (CI). The goal of the framework is to provide a user-friendly method for parallelizing data-intensive computing tasks with minimal user intervention. Some of the key design goals are usability, scalability and reproducibility. The users can focus on their problem and leave the parallelization details to the framework. This paper will outline the rationale behind this framework, its detailed implementation and demonstrate its usage with practical use cases.

pdf

A Collaborative System Design to provide Dynamic Decision Planning and assess Trust among FEW Actors
Suleyman Uslu and Shreyansh Mohnot (Indiana University - Purdue University Indianapolis)
Abstract

In sectors that are closely related to resources such as food, water, and energy, decisions have a crucial impact on the sustainability of these resources in the environment. Moreover, those decisions involve multiple actors which sometimes construct a hierarchy and become more complex as the number of actors and parameters increases. We propose a model of the interactive dynamic decision process as a multidisciplinary optimization problem and achieve optimal decisions using sequential optimization algorithm. Furthermore, we propose an application of a measurement theory-based trust framework where actors provide feedback to proposed solutions and help to find hidden or previously unknown optimization parameters or consequences of solutions.

pdf

Performance Considerations for Containers on San Diego Supercomputer Center's Comet
Philip Brown (University of San Diego, California; San Diego Supercomputer Center) and Mahidhar Tatineni (San Diego Supercomputer Center)
Abstract

Containers are an increasingly popular addition to the modern High Performance Computing workflow and are widely used on XSEDE computing resources. They allow users of traditional HPC resources more freedom with OS environments and software stacks while operating in non-privileged userspace. A common concern of researchers is getting the most computation out of their allocations. This poster explores the comparative performance analysis of common HPC application benchmarks (e.g. AMBER, GROMACS) and popular Machine Learning software (e.g. TensorFlow, PyTorch) in containers and on bare metal through regular queues. In addition, performance differences between runs via precompiled containers (such as from Singularity or Docker hub) vs compiling with Comet hardware specific optimization flags are analyzed. Often publicly available precompiled images may not have specific optimization flags for the CPUs and GPUs on a particular machine. Others like NVIDIA's NGC may handle GPU architecture optimizations but not have extensive CPU specific optimizations. In this study special attention is paid to using different libraries (with and without architecture specific optimization) and compilers to understand performance impacts and develop best practices. Methods for validation of these results are additionally presented.

Leveraging XDMoD Job Statistics Data to Predict Performance
Marian Passmore and Adrian Martinez (University of California, San Diego) and Mahidhar Tatineni (San Diego Supercomputer Center)
Abstract

A collection of HPC resources, XSEDE was responsible for 24.8 million jobs in 2018, making up 7.8 billion CPU hours and using an average of 72 processors per job. User allocations on XSEDE High Performance Computing resources are made based on code performance and based on projections for the planned simulations in a project. Metrics on XSEDE jobs are collected via an open-source tool called XDMoD leveraging the job performance module SUPReMM and available to aid users in analysis of their performance. The SUPReMM data collection is extensive with CPU (various performance counters), memory, IO (block device statistics, Lustre, NFS filesystems info), and system info collected at 10 minute intervals. Typically users do not delve into details of their jobs (which can number in the thousands). The goal of this work is to propose and demonstrate the potential use of this data for the purpose of predicting job performance, and analyze the feasibility of predicting overall CPU utilization of a job based on a fraction of the temporal data from the early section of a run. Potentially this can flag slow running jobs in the early stages and make resource usage more efficient. The analysis is performed for a dataset of Comet jobs longer than 1 hour wall clock time (to provide enough time samples). Given the extensive list of parameters, a PCA analysis is performed for feature reduction. Several regression approaches, deep neural network based models are considered. TensorFlow and sklearn machine learning frameworks are used in the work.

ADCIRC Analysis of Super Typhoon Haiyan in Leyte Gulf
Jack Gaither (University of Arkansas), Clint Dawson (The University of Texas at Austin), and Nilo Espinoza (University of Guam)
Abstract

Typhoon Haiyan made landfall as a Category 5 storm on the eastern coast of the Philippines on November 8, 2013. Characterized by very low central pressures, fast forward moving speed, and sustained wind speeds higher than 150 miles per hour, Haiyan was one of the most devastating Typhoons to make landfall in the Philippines, displacing 4.1 million people, killing more than 6000 and leaving 1800 unaccounted for. The combination of the unique sea floor topography in Leyte Gulf and the Philippines Trench paired with such a powerful storm resulted in widespread destruction of the areas surrounding Leyte bay and more generally, in the path of Haiyan. The Leyte Gulf is defined by primarily shallow water that quickly drops off at the open ocean boundary into the Philippines Trench. This combination of underwater features paired with the unique characteristics of Haiyan allowed for extensive inundation of the surrounding area. This poster details the simulation of storm surge and tidal interactions of Typhoon Haiyan in the Island chain surrounding Leyte Gulf. All computations were done on Lonestar using the Advanced Circulation (ADCIRC) unstructured grid hydrodynamic model.

Facilitating high performance computing for the sciences: a case study in machine learning for materials science
Michael Butler and Bruce Segee (University of Maine, Northeast Cyberteam) and Liping Yu (University of Maine)
Abstract

The Northeast Cyberteam Initiative is a National Science Foundation funded program aimed at creating a pool of Research Computing Facilitators (RCFs) to help scientists at small and mid-sized universities in the northeast migrate their computational work to high performance facilities in order to more efficiently conduct their research. At the University of Maine a Northeast Cybteream RCF assisted a computational materials scientist in parallelizing and optimizing their single computer machine learning (ML) algorithms for deployment in a high performance computing (HPC) cluster.

Traditionally, the physical properties of solids are calculated from solving the many-body Schrödinger Equation using their atomic compositions as inputs. However, these first-principles methods are computationally demanding and the results cannot be generalized to predict the properties of novel materials without performing additional time-intensive computations. This research project aimed to develop novel material descriptors and ML models to accurately predict the properties of materials in a fraction of the runtime of traditional methods. The development of such descriptors and models is a key first step in high-throughput discovery of novel energy materials.

This research project used large datasets (~128k molecules/materials) and involved significant computation time. It therefore benefited greatly from being moved to the University of Maine HPC cluster. Through this collaboration the RCF was able to dramatically speed up material descriptor generation and ML model training, which in turn allowed for faster generation and evaluation of descriptors and models. This case study provides an excellent example of how RCFs can help accelerate computationally demanding scientific research.

Dynamic metabolic modeling based on in vivo, time series measurements
Yue Wu, Michael Judge, Arthur Edison, Jonathan Arnold, and Heinz-Bernd Schuttler (UGA)
Abstract

Multiple modeling approaches of metabolic networks have been developed for Neurospora crassa. However, common methods often assume steady state and fixed network topologies, such as Flux Balance Analysis (FBA). In order to further analyse metabolism, dynamic modeling with network topology identification is needed. Identifying kinetic parameters and topology has been difficult task. Metabolic time-series measurements can provide crucial constraints on these parameters, but the conventional approach is labour-intensive and subject to biological and technical variability. In this project, using continuous in vivo monitoring of metabolism by NMR (CIVM-NMR), we obtained time-series metabolic measurements on single N. crassa cultures, which eliminates extraction variance. Based on these time-series measurements, we built a model for the central metabolic pathway (~50 reactions), including, inter alia, glycolysis, fermentation and TCA cycle. Reactions in the network are modeled by ODE (Ordinary Differential Equation). A Metropolis-Hastings based MCMC (Markov chain Monte Carlo) optimization process is implemented to cycle. Reactions in the network are modeled by ODE (Ordinary Differential Equation). A Metropolis–Hastings based MCMC (Markov chain Monte Carlo) optimization process is implemented to explore the parameter space and find a large sample of models consistent with the measurements. Evaluation of existing allosteric regulation is tested by fixed topology models, and new regulation will be estimated by variable topology models. Different regulation functions (power equation and Hill equation) are compared for numerical performance and fitting accuracy. Parallel tempering will also be used to speed-up optimization. ODE simulation is CPU intensive and takes most of the time in optimization. Parameter space exploration by MCMC takes tens thousands of sweeps. MPI resources from GACRC and XSEDE are currently used for model identification in this project.

Automatic recognition of frog vocalizations using Jetstream
Eliza Foran (Indiana University), Tenecious Underwood (Livingstone College), Evan Suggs (University of Tennessee at Chattanooga), and Sheri Sanders and Winona Snapp-Childs (Indiana University)
Abstract

Recording animal calls and vocalizations is a time-honored data collection method in various fields of biological and environmental science. In the past, the only method available for analyzing such recordings involved extensive training of human experts. Now, however, machine learning techniques have made automatic recognition of such vocalizations possible. Automatic recognition of animal calls and vocalizations is desirable on two fronts – it reduces the burden of (at least initial) data analysis and serves as a piece that supports nonintrusive environmental monitoring. Here, we outline a proof-of-concept workflow that will make the quest from collecting data to understanding data more attainable for researchers. We simulate this data collection process by collecting animal (frog) calls using recording devices and Raspberry Pi's, then feed this data to a database and virtual machine hosted on XSEDE resources (i.e. Jetstream and Wrangler). We then show how database pulling, machine learning, and visualization works on Jetstream.

Porting CPU Agent-Based Modelling Applications to GPU
Vincent Li and Gil Speyer (Arizona State University)
Abstract

Agent-based modelling (ABM) employs concurrent, communicating agents to model emergent, complex behaviors. Netlogo, a widely used CPU application, incorporates many ABM models. Graphical processing unit (GPU) implementations of ABMs have demonstrated significant acceleration due to the massive data-parallel computation, custom random-number generators, and the GPGPU "scatter" communication model. This project focuses on converting Netlogo models to run on one such GPU-based ABM application, FLAME GPU. This effort will facilitate researchers within domains that employ Netlogo to access accelerated computation.

This application of advanced research computing software and applications should dovetail with the purpose of PEARC, exploiting the performance made feasible on GPUs by empowering state-of-the-art ABM simulations.

Spatial Data Decomposition and Load Balancing on HPC Platforms
Jie Yang, Anmol Paudel, and Satish Puri (Marquette University)
Abstract

We are in the era of Spatial Big Data. Due to the developments of topographic techniques, clear satellite imagery, and various means for collecting information, geospatial datasets are growing in volume, complexity and heterogeneity. For example, OpenStreetMap data for the whole world is about 1 TB and NASA world climate datasets are about 17 TB. Spatial data volume and variety makes spatial computations both data-intensive and compute-intensive. Due to the irregular distribution of spatial data, domain decomposition becomes challenging. In this work, we present spatial data partitioning technique that takes into account spatial join cost. In addition, we present spatial join computation using Asynchronous Dynamic Load Balancing (ADLB) library. ADLB is a software library designed to help rapidly build scalable parallel programs using MPI. We evaluated the performance of ADLB-based MPI-GIS implementation. In our existing work, spatial data movement cost from ADLB server to worker MPI processes limited the scalability of MPI-GIS.

pdf

Generative Adversarial Network for Lupus Diagnostics and Medical Applications
Pradeep Periasamy and Dr. Vetria L Byrd (Purdue University)
Abstract

The recent boom of Machine Learning frameworks like Generative Adversarial Networks (GAN), Deep Convolution Generative Adversarial Networks (DCGAN) and the development of high-performance computing for big data analysis has the potential to be highly beneficial in many domains and fittingly in the early detection of chronic diseases. The clinical heterogeneity of one such chronic autoimmune disease like Systemic Lupus Erythematosus (SLE), commonly referred to as Lupus, makes it difficult for medical diagnostics. This research employs unsupervised deep learning mechanisms to identify clinical manifestations of lupus from publicly available anonymous pictures of persons who present with cutaneous lesions like the butterfly rash, commonly seen in patients diagnosed with Lupus. We demonstrate the use of artificially generated butterfly rash images generated from GAN to train the discriminator model that differentiates Lupus from its other counter skin diseases using a Neural Network Classifier, as a use-case example. The expected outcomes are to help reduce the time in detection and treatment by gathering insights from its huge heterogeneous data clusters.

pdf

Flexible Partitioning of Scientific Workflows Using the JX Workflow Language
Tim Shaffer, Nathaniel Kremer-Herman, and Douglas Thain (University of Notre Dame)
Abstract

Scientific workflows are typically expressed as a graph of logical tasks, each one representing a single program along with its input and output files. A conventional workflow manager transforms each logical task into a discrete batch job and submits it to an underlying execution system. However, converting every logical task into one batch job is not necessarily the most efficient partitioning of a workflow. By grouping multiple logical tasks into a single batch job, we may decrease data transfer, increase system utilization, and reduce the execution time of a workflow. This paper presents JX (JSON eXtended), a declarative language that can express complex workloads as an assembly of sub-graphs that can be partitioned in flexible ways. We present a case study of using JX to represent complex workflows for the Lifemapper biodiversity project. We evaluate partitioning approaches across several computing environments, including HTCondor at the University of Notre Dame, TACC Stampede2, and SDSC Comet, and show that a coarse partitioning results in faster turnaround times, reduced data transfer, and lower master utilization across all three systems.

pdf

Towards Run Time Estimation of the Gaussian Chemistry Code for SEAGrid Science Gateway
Angel Beltre, Shehtab Zaman, and Kenneth Chiu (Binghamton University); Sudhakar Pamidighantam (Indiana University); and Xingye Qiao and Madhusudhan Govindaraju (Binghamton University)
Abstract

Accurate estimation of the run time of computational codes has a number of significant advantages for scientific computing. It is required information for optimal resource allocation, improving turnaround times and utilization of science gateways. Furthermore, it allows users to better plan and schedule their research, streamlining workflows and improving the overall productivity of cyberinfrastructure. Predicting run time is challenging, however. The inputs to scientific codes can be complex and high dimensional.Their relationship to the run time may be highly non-linear, and, in the most general case is completely arbitrary and thus unpredictable (i.e., simply a random mapping from inputs to run time). Most codes are not so arbitrary, however, and there has been significant prior research on predicting the run time of applications and workloads. Such predictions are generally application-specific, however. In this paper, we focus on the Gaussian computational chemistry code. We characterize a data set of runs from the SEAGrid science gateway with a number of different studies. We also explore a number of different potential regression methods and present promising future directions.

pdf

A Virtual Reality Approach to Pediatric Conflict De-escalation and Anger Management
Heyi Liu (Baylor School)
Abstract

In recent years, anger management has received considerable attention in the school setting. Youth suffering from elevated anger often have related health problems, typically cardiovascular diseases; they also tend to report more physical aggression, negative verbal responses, drug use and negative consequences of their anger. While many forms of treatment currently exist, they often require a counselors' assistance in providing suitable de-escalation methods. For public schools, most states have difficulties allocating funding to establish a large enough school counseling program; therefore, anger management services are not commonly accessible in schools. Virtual Reality creates an immersive experience through computer-generated environments. First popularized as a game enhancement, Virtual Reality has contributed to academia by recreating a desired atmosphere, notably within psychological research. This study investigates the potential of a virtual reality de-escalation mechanism to increase its accessibility and affordability in a school setting. The program aims to achieve an immediate de-escalation for the youth through a calm room with minimal adult assistance and prepare them to re-enter the real world by creating a solution to a common anger trigger which can hopefully eventually reduce their dependence on the technology.

pdf

A Virtual Reality Based Environment for Immersive, Hands-On Learning
Atherton Mook (Baylor School)
Abstract

Within the virtual reality (VR) world, the process of creating a usable model is relatively intuitive and easy. The goal of this work was to create a learning tool for teachers to use in their classrooms that allows the students to take apart and put back together different models of the teacher's choice. This virtual model construction/deconstruction could be an affordable, more feasible alternative when hands-on learning is not possible. The product could save schools money and allow the students to explore a much more diverse range of material.

pdf

Not you average day at the beach: a virtual exploration of mans influence on the natural world
Zachary Baptiste (Chaminade University)
Abstract

The 2 mile long Ala Wai Canal was designed and artificially created in the 1920's to drain the area that would become the extremely popular tourist destination, Waikiki beach on the island of O'ahu in Hawaii. The canal runs between the ecologically important Ko'olau Mountain Range down to the densely populated Waikiki beach. Dramatic elevation changes combined with high rainfall poses safety risks due to flash flooding. The Ala Wai canal is also extremely polluted and poses health risks due to disease and chemical exposure.

A'ohe hua o ka mai'a i ka lā ho'okahi: is a native Hawaiian saying that means "no task is too big when done together by all". Pollution is a huge global problem whose solution beckons every individual's participation every moment of the day. In order to engage and empower an individual, you must first make the problem relative to them. Storytelling narratives have historically been an important way knowledge is passed on in traditional Hawaiian culture. This poster accompanies the Visualization Showcase virtual reality perspective of the Ala Wai Canal network stations pollutant levels and water flow over an Arc-GIS framework integrated with a Hawaiian storytelling perspective to provide a culturally-informed water map of the Ala Wai area. This poster describes the overall process in the creation of the VR experience: from data science to storytelling.

Data Curation and Analysis for Semiconductor Device Research
Robert Kaufman, Patrick Su, Zhe Yang, Steve Konstanty, Todd Nicholson, John Dallesasse, and Klara Nahrstedt (University of Illinois at Urbana-Champaign)
Abstract

The process of fabricating and characterizing a semiconductor device can be tedious, in part because of the copious amount of variables that contribute to the final outcome. Proper documentation of process parameters, physical device progress, and preliminary data is crucial to fine-tuning the fabrication recipe and creating a quality device. A secure database system designed with this documentation in mind greatly improves the fabrication workflow. By allowing the database to connect with a networked instance of Jupyter Notebook, the processing of intermediate data is further streamlined. Future improvements planned for the system include utilizing the advanced computational methods such as image classification and object detection to identify connections between process parameters and results that are difficult for an individual to identify.

Virtual Exploration in a Socially Aware Context: An exploration of Waller Creek
Fidelia Hernandez (Highland College)
Abstract

Waller Creek is an urban waterway and Austin landmark that runs from just north of the University of Texas (UT) campus through the eastern edge of downtown into Lady Bird Lake. Its proximity to the city center presents conservation challenges. By collecting water conductivity data, Dr. Mary Poteet from UT Geological Sciences is studying the creek's water quality with the goal identifying locations indicating influx that shouldn't be there eg. from city drinking water, sanitary sewers or other pollution sources. An important aspect of her work is communicating her research results to city officials and legislature to inform conservation efforts and also help locate leaks in the city's plumbing. To more effectively communicate the results of her water surveys, she requires many layers of map data to be visualized in context. This allows the pollutant sources to be identified, and also traced to outflows so that the effects on the entire watershed can be assessed at once. This poster provides supplemental information to accompany the Visualization Showcase submission that utilizes virtual reality to present an overlay of maps showing Poteet's collection locations and results with watershed topography. This poster highlights considerations made in designing the virtual exploration of the region in order to promote an awareness of human sources of pollution.

10-Week Internship at Lawrence Berkeley National Laboratory
Edsel Norwood (Elizabeth City State University)
Abstract

During the summer of 2018 I took part in a 10-week undergraduate internship at Lawrence Berkeley National Laboratory (LBL). The department in which I was working under was the Data Science and Technology department, directed by Dr. Deborah Agarwal. During the course of the internship I was tasked with the development of user tools for a service by the name of Environmental Systems Science Data Infrastructure for a Virtual Ecosystem (ESSDIVE). This included working on multiple different project during my 10 weeks at LBL.

Data Driven Urban Data Mining
Tangee Beverly (Elizabeth City State University)
Abstract

Urban planning seeks to use big data to build smart cities. Urban data in NewMexico was crawled from creditable sources on-line and by requesting. Twocase studies were conducted in this research for the purpose of urban plan-ning and address two questions in our research. For case study 1, all the datawas collected fromNew Mexico Environmental Public Health Tracking, such as,Asthma, COPD, and Ozone, which are the major factors focused in this researchand case study II requested data from theNew Mexico Department of Trans-portation. Case study I, processed the data using programming language ofpython by implementing imported library,pandas. The data was inputted intoscatter plots to find correlation between Ozone and Asthma, Ozone and COPD.However, the plots did not show much promising correlation, so Pearson corre-lation and K-means clusters are the analysis tools used to test whether there is acorrelation between these factors. Case study II used python imported librariespandas, sklearn, and seaborn to visualize the data spatially and by time. Trendsfound through visualization was used as a basis for preprocessing of data andthen running the apriori algorithm to determine highest frequent combinationsof factors. Future study will involve crawling of more factors to include in spa-tial analysis to help better understand the trends and peaks seen in the data.Analysis of the factors will provide invaluable information to the Departmentof Transportation and first responders in the area.

SENSELET: Sensory Network Infrastructure for Scientific Lab Environments
Zhe Yang, Patrick Su, Robert Kaufman, Steven Konstanty, John Dallesasse, and Klara Nahrstedt (University of Illinois at Urbana-Champaign)
Abstract

Materials and devices are sensitive to environments. Environmental conditions can cause issues in yield, performance, and assembly, and consequently must be carefully monitored and maintained. Due to limited budgets, universities and government laboratories cannot commit to expensive state-of-the-art monitoring systems. However, low- cost wireless sensors provide a non-invasive solution of tracking and monitoring key environmental data in real-time via cloud-based networks. The acquired data can not only be utilized to alert research scientists and laboratory administrators of environmental conditions and safety hazards, but also can be used for large data analysis.

Dynamic Parallelization
Karthik Gurunathan and Kaustubh Kartikey (PES Institute of Technology)
Abstract

This idea involves building a unified powerful super compiler that can execute programs written in any language. The idea is to preempt decisions to be taken during run time. Such a decision can only be taken if a powerful machine learning model is built, that has the ability to assign the correct accelerator or processing unit to the correct segment of code. This process is called dynamic parallelization. Dynamic parallelization provides the compiler with ultimate power, thereby aiding in realizing every computer scientist's ultimate dream, to develop one super-language that can cater to any utility and is easy to interpret.

The Challenges of Representing Geographical Information in Virtual Reality
Jennie Youn (Emory University)
Abstract

AR/VR technology (headsets and the game engine Unity3D) is ideal for providing a geographical context for data display and visualization. The metadata for the geographical location of the sensor data needs to be maintained to accurately display the results in context. There currently isn't an easy way for environmental scientists to get this data into the Unity3D game engine for viewing. This poster addresses the shortcomings of the current applications/scripts for processing geographical data for VR exploration and efforts to mitigate the problems by 1) developing a library of python/pandas scripts for data processing sensor data from direct as well as online databases and the c# scripts for getting those into the Unity3D visualization application 2) development of the workflow for the pipeline from data collection to visualization 3) development of a compelling VR experience to engage the general audience to the environmental issues of Waller Creek and call them to action.

Reducing Faulty Jobs by Job Submission Verifier in Grid Engine
Misha Ahmadian, Eric Rees, Yu Zhuang, and Yong Chen (Texas Tech University)
Abstract

Grid Engine is a Distributed Resource Manager (DRM), that manages the resources of distributed systems (such as Grid, HPC, or Cloud systems) and executes designated jobs which have requested to occupy or consume those resources. Grid Engine applies scheduling policies to allocate resources for jobs while simultaneously attempting to maintain optimal utilization of all machines in the distributed system. However, due to the complexity of Grid Engine's job submission commands and complicated resource management policies, the number of faulty job submissions in datacenters increases with the number of jobs being submitted. To combat the increase in faulty jobs, Grid Engine allows administrators to design and implement Job Submission Verifiers (JSV) to verify jobs before they enter into Grid Engine. In this paper, we will discuss a Job Submission Verifier that was designed and implemented for Univa Grid Engine, a commercial version of Grid Engine, at the High Performance Computing Center of Texas Tech University. Our newly developed JSV communicates with Univa Grid Engine (UGE) components to verify whether a submitted job should be accepted as is, modified then accepted, or rejected due to improper requests for resources. It had a substantial positive impact on reducing the number of faulty jobs submitted to UGE by far. For instance, it corrected 28.6% of job submissions and rejected 0.3% of jobs from September 2018 to February 2019, that may otherwise lead to long or infinite waiting time in the job queue.

pdf

Controlled Unclassified Information Framework
Jessica Ackerman (Purdue University)
Abstract

Many researchers at universities struggle to begin and maintain cybersecurity compliance for their data. They are often left in the dark, and don't know where to start. Even once they begin, they often take shortcuts leading to an even less secure environment for their data. The CUI framework at Purdue is attempting to remedy this problem. Using NIST 800-171 for a basis and mapping out other data types such as HIPAA, FERPA, and NIST 800-53 to the 171 controls, creating an environment that is more secure and easier for the researchers to meet compliance.

Design and Implementation of an Educational Web Based Application for Remote Sensing Concepts
Derek Morris Jr (Elizabeth City State University)
Abstract

Geographic Information Systems (GIS) and Remote Sensing is an industry with a high workforce demand and even higher environmental importance. There are many programs in which young adults are introduced to classes/careers that, otherwise may not come into contact with. It's these types of programs that expose young adults to the open possibilities and career paths that lay ahead of them, and in doing so these young adults are given a chance to discover their passions, and plan the rest of their academic careers accordingly. It was for this reason that we felt that the training modules within the GAIA Crossroads textbook needed to be updated to take a more modern approach.

Validating photogrammetric processes using Jetstream and reverse engineering
Matthew Mercer (Indiana University), Jenny Zhao (Massachusetts Institute of Technology), and Tassie Gniady and Eric Wernert (Indiana University)
Abstract

This project focuses on photogrammetry--the process of making 3D objects from 2D photographs--and the reverse engineering of this complex process. While a number of disciplines from cultural heritage to natural sciences use photogrammetry, the algorithms used by the most popular (and accurate) software packages are black box because of their proprietary nature. This project takes a synthetic 3D object which can be fully understood on the digital level and reverse engineers the photogrammetry process to determine exactly how accurate the modeling process is. This is of key importance to anyone trying to make reliable models available to remote researchers, especially when items such as fossils are 3D modeled. The researcher needs to have confidence in the final model to draw valid conclusions if they are unable to visit the collection in person. A virtual machine (VM) was created on Jetstream to import synthetic models, capture them, and export those captures for processing. Utilizing a parallel processing workflow on Jetstream, the speed up for creating 3D models allowed for refinement and comparison of workflows and models with much shorter turn-around times.

Visualizing (veterinary) medical data sets with Jetstream
Alan Nguyen (Purdue University), Yvan Pierre (Valencia College), and Scott Birch and Eric Wernert (Indiana University)
Abstract

Computed tomography (CT) is a diagnostic imaging test using x-rays to create multiple detailed diagnostic images of internal organs, bones, soft tissue and blood vessels. It produces a data set of thin, cross-sectional "slices" for viewing and is much more detailed than conventional radiography. Clinicians use CT examination to diagnose cancers, detect abnormal blood vessels, discover disorders of the abdomen, bones, and joints, and to plan surgical interventions such as heart defect or vascular repair. Dedicated visualization workstations allow radiologists to make high-resolution examinations of the data for diagnoses, but understanding the image stacks can be challenging for clinicians without specialized skills, training, and experience. To aid and enhance diagnostic evaluation, we explored a cloud-based workflow using Jetstream. CT data sets were segmented or translated into regions-of-interest (ROI) and/or volumetric 3D reconstructions which were then exported as polygonal 3D surface models. Using data sets obtained via CT from different species, this project focused on compiling a medical imaging/segmentation workstation instance with open source software on Jetstream, importing sample data sets into the imaging software, viewing 2D image sequences volumetrically, setting custom transfer functions based on tissue density, and segmenting the anatomy into multiple ROI for export as stereolithography files. Post-processing and polygon mesh editing techniques (smoothing, transient reduction, decimation) were demonstrated as the model was optimized for 3D printing or online distribution. Results were rendered into 2D graphical representations, the 3D models were deployed into interactive or virtual reality environments, or additively-manufactured into real-world objects for visual and tactile examination.

Parallel Simulation of AC Optimal Power Flow Utilizing the PETSc Library
Rylee Sundermann (South Dakota State University), Hong Zhang (Argonne National Laboratory), and Timothy Hansen and Jung-Han Kimn (South Dakota State University)
Abstract

We present a scalable AC Optimal Power Flow Simulation which demonstrates the capabilities of the new data management object DMNetwork in the PETSc Library (Portable Extensible Toolkit for Scientific Computation) developed at Argonne National Laboratory. Our main focus was DMNetworks data and topology management, capability to simulate multiphysical phenomena over large scale network systems, scalability, and results of large scale ACOPF simulations.

3D Simulations of Biofilms using the Modified Cahn-Hilliard Equation
Rylee Sundermann, Nicholas Stegmeier, Nathan McClanahan, Jung-Han Kimn, and Jeffery Doom (South Dakota State University)
Abstract

We present a 3D implementation of the modified Cahn-Hilliard (MCH) equation utilizing the Portable Extensible Toolkit for Scientific Computation (PETSc) to simulate the growth of a biofilm in an aqueous environment. Biofilm dynamics under shear flow are investigated for a variety of circumstances, and the parallel performance of the implementation is evaluated.

Tuesday 6:30pm-8:30pm
Crystal Foyer and Crystal B
Visualization Showcase · Reception, Vis Showcase, Interactive Visualization Demonstrations, Visualization Showcase Presentation Award

Visualization Showcase Reception

Visualization at the Texas Advanced Computing Center
Greg Foss, Greg Abram, Joao Barbosa, Anne Bowen, Ayat Mohammed, Paul Navrátil, Francesca Samsel, Dave Semeraro, Andrew Solis, Jo Wozniak, and Nitesh Kartha (Texas Advanced Computing Center, University of Texas at Austin)
Abstract

'Visualization at the Texas Advanced Computing Center (TACC)'

This video compilation of animation shorts and tests represents a sample of TACC's current visualization projects. Supporting scientists in fields as diverse as micro electronics, nuclear fusion, Arctic Ocean Atlantification, and cancer metastasis, the animations as well demonstrate visualization research in software raytracing, in situ machine learning, and interactive techniques.

Featured are contributions from TACC visualization specialists Greg Abram, Joao Barbosa, Anne Bowen, Greg Foss, Ayat Mohammad, Paul Navrátil, Francesca Samsel, David Semeraro, Andrew Solis and Jo Wozniak.

Waller Creek Water Quality: Visualization Through Augmented Reality
Greg Foss, Dave Semeraro, and Andrew Solis (Texas Advanced Computing Center, University of Texas at Austin); Mary Poteet (University of Texas at Austin); Fidelia Hernandez (Highline College); Jennie Youn (Emory University); Anne Bowen (Texas Advanced Computing Center); and Matias Fragoso (Texas Advanced Computing Center, University of Texas at Austin)
Abstract

Waller Creek is an urban waterway and Austin landmark that runs from just north of the University of Texas (UT) campus through the eastern edge of downtown into Lady Bird Lake. Its riparian ecosystem hosts tremendous biodiversity, but its proximity to the city center presents conservation challenges. By collecting water conductivity data, Dr. Mary Poteet from UT Geological Sciences is studying the creek's water quality with the goal identifying locations indicating influx that shouldn't be there eg. from city drinking water, sanitary sewers or other pollution sources. An important aspect of her work is communicating her research results to city officials and legislature to inform conservation efforts and also help locate leaks in the city's plumbing. To more effectively communicate the results of her water surveys, she requires many layers of map data to be visualized in context. This allows the pollutant sources to be identified, and also traced to outflows so that the effects on the entire watershed can be assessed at once. The work presented here utilizes augmented reality with Microsoft's HoloLens to present an overlay of maps showing Poteet's collection locations and results with watershed topography. Maps of the non-point source surveys highlight locations where anomalously high levels of conductivity -- which are likely locations with sewage leaks, were found.

HiperViz: Interactive visualization of CPU Temperatures inHigh Performance Computing Centers
Visualization Showcase Presentation Award
Ngan Nguyen and Tommy Dang (Texas Tech University)
Abstract

This showcase presents HiperViz, an interactive visual interface for real-time monitoring of the health status of high-performance computing systems. The objective of our visualization system is threefold: 1) to provide a visual interface for tracking CPU temperatures in large data centers, 2) to detect abnormal events and their temporal and spatial correlations, and 3) to assist system administrators in troubleshooting and maintenance.

pdf

Vapor 3: A modern 3D visualization suite for Geophysical Sciences
Scott Pearse (NCAR)
Abstract

Vapor v3.1 is an open-source visualization and analysis tool for physical sciences. This release is a key milestone in Vapor's development, and this video will demonstrate the improvements that have been made since Vapor 2.6 was released two years ago.

With this most recent release, Vapor has undergone a refactorization of its code base that spans a 15 year history. Many shortcomings have been addressed, and these improvements aim to help scientists produce science that is superior to what has been previously possible with 3D visualization.

Some of these areas include improvements to Vapor's data model, user interface, color application, and its general feature set.

Side-by-side comparisons of Vapor 2.6 will be done with Vapor 3.1 so old users can come up to speed with what's new. New users will get a high level view of how to get their data into Vapor 3.1, start rendering, and how to get help. Vapor's future roadmap will finally be discussed, along with a discussion on how users can request features, submit bugs, and even contribute their own code.

pdf

A Mixed Reality Exploration of Groundwater Conservation Districts
Andrew Solis, Jo Wozniak, and Suzanne Pierce (Texas Advanced Computing Center, University of Texas at Austin); Elizabeth Hood (Texas Alliance of Groundwater Districts); and Matias Fragoso (Texas Advanced Computing Center, University of Texas at Austin)
Abstract

The growing ecosystem of emerging technologies allows researchers to view their data from different perspectives, enabling them to gather insight. Mixed Reality is a powerful example, letting users view virtual objects overlaid in their existing world. This gives users a clear sense of how the virtual representation of the data they are viewing correlates to their natural environment. The Texas Advanced Computing Center (TACC) demonstrates an example of using a mixed reality headset to explore Geo-referenced information. Our application renders a virtual 3D map, using the Microsoft HoloLens, that is separated into individual groundwater conservation districts across Texas. Users use gestures to select districts, compared to the limited mouse and click selection required from a traditional GUI. Primary development used Unity3D, with data provided by the Texas Alliance of Groundwater Districts. This project resulted from work related to the Planet Texas 2050 initiative at the University of Texas at Austin, focusing on sustainability for future generation Texans. The first demonstration was at the Orange and Maroon legislative day February 5th, 2019 in Austin. The all day event is a platform to allow UT Austin and Texas A&M alumni to meet with their district representatives and discuss future goals and legislation for a better Texas. Future work involves a collaborative application where users can discuss the same virtual content, and not be limited by the view of a single person or screen.

Visualization of Hurricane Michael
David Bock (National Center for Supercomputing Applications), Matthew Bilskie (Louisiana State University), and Stephen Medeiros (University of Central Florida)
Abstract

Hurricane Michael made landfall near Mexico Beach, FL on October 10 as a Category 5 storm. Measurements of peak water levels revealed storm surge as high as to 4.37 m (NAVD88). Areas surrounding Mexico Beach, Port St. Joe, St. George Island, and Apalachicola experienced wide-spread flooding and extensive damage. To gain an understanding of storm surge and related overland flooding, we simulate the water level and wave response from Hurricane Michael using a tightly-coupled ADCIRC+SWAN model of the northern Gulf of Mexico (NGOM), NGOM3. The ADCIRC (Advanced Circulation) code computes water levels and depth-averaged currents via the shallow water equations while SWAN (Simulating Waves Nearshore) computes relative frequency and direction of wind-waves from the action balance equation. Both models utilize the same NGOM3 unstructured finite element mesh (5.5 million nodes) that spans the western north Atlantic Ocean, Caribbean Sea, and Gulf of Mexico and includes detailed representation of the Mississippi, Alabama, and Florida panhandle coastal floodplain. Model resolution ranges from 20 m – 200 m across the overland regions and describes the significant hydraulic features such as bay/inlet systems, sounds, estuaries, coastal river, and the Gulf Intracoastal Waterway that can convey or inhibit storm surge flows. The NGOM3 model is a result of over a decade of effort and improvements. A custom visualization system is used to visualize a variety of parameters including wave height, water surface elevation, wind, and current.

An Exploratory Sunburst Tool for XSEDE and HPC Resources
Andrew Solis, Latrell Gaither, and Briana Bradshaw (Texas Advanced Computing Center, University of Texas at Austin)
Abstract

There is a wide array of research being done on supercomputers today. Everything from the natural sciences, to the humanities and government utilize HPC resources in a multitude of ways. It is important to gain insight into what type of research topics and projects are using these HPC systems, and how many CPU resources each requires. The Texas Advanced Computing Center (TACC) demonstrates the early stages of an information visualization application that can be used to explore the types of research being done on HPC systems. Here we demo a sunburst chart, similar to a pie chart, where each slice represents a field of science that is utilizing our HPC resource. The size of a slice correlates to the amount of cores being used for a particular field of science. Users are able to select a field of science by clicking it's associated slice, which then creates a new sunburst. The resulting sunburst is separated by individual projects that are conducting research for the same field of science. The size of the slice again correlates to how many resources a particular project is using. Users can select a project's slice to gather more information about the particular project. Primary development was done in D3, a javscript framework for online web visualizations. Future intent is to create a simple framework where users who maintain similar systems can utilize our application to gain insight into what types of research is being done on their own systems.

pdf

Flight Paths: Mapping Our Changing Neighborhoods
Visualization Showcase Presentation Award
David Bock (National Center for Supercomputing Applications) and Elizabeth Wuerffel (Valparaiso University)
Abstract

Flight Paths: Mapping Our Changing Neighborhoods is a multi-media initiative analyzing factors that contribute to the fracturing of urban neighborhoods, communities, and regions in post-industrial America. It documents the changing racial and economic demographics of Gary and Northwest Indiana, including the rise of black political power and opportunity in the '60s and '70s, the flight of white residents and businesses to the suburbs, and the automation and underemployment of the steel mills. The website will animate migration of residents over time and host a curated portion of historical materials and professionally edited interviews (residents and scholars) within an interactive map of Northwest Indiana. The Flight Paths website will open with a visualization of population growth and density across the United States from 1890-2010, visually connecting the story of Gary to the story of the nation. As the population growth reaches its peak in Gary during the 1960s, the animation will zoom into Northwest Indiana, both to indicate that Gary is the website's main area of interest and to illustrate how the surrounding cities grew as Gary began to empty out. This initial visualization will function in two primary ways: 1) frame the national story of urban decline and suburban growth through the lens of Gary, Indiana, and 2) to entice users to enter the site through an aesthetically engaging, data-driven story prompt: 180,000 people lived in Gary, Indiana, by 1960. Fifty years later, only 80,000 people lived there, an experience shared by other rust belt cities in America. Why?

Ala Wai Watershed: 'A'ohe hua o ka mai'a i ka lā ho'okahi
Zachary Baptiste and Mark Speck (Chaminade University); Anne Bowen (TACC, University of Texas, Austin); and Matias Fragoso (Texas Advanced Computing Center, University of Texas at Austin)
Abstract

'A'ohe hua o ka mai'a i ka lā ho'okahi: is a native Hawaiian saying that means "no task is too big when done together by all". Pollution is a huge global problem whose solution beckons every individual's participation every moment of the day. In order to engage and empower an individual, you must first make the problem relative to them. Storytelling narratives have historically been an important way knowledge is passed on in traditional Hawaiian culture. This visualization showcase aims to integrate Hawaiian storytelling traditions and data-science to engage the local population around the Ala Wai Watershed and invite them to take action.

The 2 mile long Ala Wai Canal was designed and artificially created in the 1920's to drain the area that would become the extremely popular tourist destination, Waikiki beach on the island of O'ahu in Hawaii. The canal runs between the ecologically important Ko'olau Mountain Range down to the densely populated Waikiki beach. Dramatic elevation changes combined with high rainfall poses safety risks due to flash flooding. The Ala Wai canal is also extremely polluted and poses health risks due to disease and chemical exposure.

Efforts to monitor, understand and quantify pollutant loadings into the Ala Wai Canal from various sources are underway. This work will provide an immersive Virtual Reality perspective of the Ala Wai Canal network stations pollutant levels and water flow over an Arc-GIS framework integrated with a Hawaiian storytelling perspective to provide a culturally-informed water map of the Ala Wai area.

HiperVR: A Virtual Reality model for Visualizing Multidimensional HealthStatus of High Performance Computing Systems
Ngan Nguyen and Tommy Dang (Texas Tech University)
Abstract

This paper proposes a virtual reality model for situational aware-ness and health monitoring of high-performance computing systems. The virtual reality model takes into account the following design requirements: 1) Mimicking data center spatial layout, 2) Tracking temporal changes, and 3) Monitoring multidimensional health statuses such as temperature, CPU load, memory usage, fan speed, and power consumption. To show the effectiveness of our design, we demonstrate the developed virtual reality model on a medium-scale high-performance computing system of 10 racks and 467 hosts at a university. The work was developed using weekly feedback from both industrial and academic domain experts.

pdf

Wave breaking of a parametrically resonated normal mode in Fluid Dynamics
Jason Yalim (Arizona State University)
Abstract

Animation of different fields along the mid-plane sections of a short-depth rectangular cuboid filled with stably stratified fluid. The container is subject to harmonic forcing in the vertical direction, and the animation is over 5.4 forcing periods. The harmonic forcing parametrically resonates a normal mode of the system so strongly that it leads to fantastic wave-breaking.

This is apparent in all four fields that are plotted: the temperature (upper left), the vorticity (how much the fluid is spinning, bottom left), the thermal dissipation (the heat generated due to thermal effects, upper right) , and the momentum dissipation (the heat generated due to viscosity, bottom right).

Return to Top

Tuesday 9:00pm-11:59pm

Tuesday 9:00pm-11:59pm
Field
Special Event · Hackathon

SGCI Hack@PEARC19

SGCI Hackathon (sciencegateways.org)
Abstract

Join us for three days of science gateway hacking with some of the best programmers in the field!

In conjunction with PEARC19 in Chicago, SGCI is sponsoring a Hackathon to challenge up-and-coming programmers for tomorrow's science gateways. Undergraduate and graduate students will compete on gateway projects established by SGCI partners and community members.

Join us on the social media app SLACK to get the latest updates! Follow us on Twitter at @sciencegateways The competition will take place from Sunday, July 28, to Wednesday, July 31, in the Hyatt Regency Chicago where the PEARC19 conference is being held. Final awards will be announced on Thursday, August 1.

Return to Top

Wednesday 7:00am-5:00pm

Wednesday 7:00am-5:00pm
Regency Ballroom Registration Desk - Regency Foyer
Registration · Registration

Conference Registration and Information Desk

Abstract

Conference Registration and Information Desk Hours: Sunday – 12:00pm to 8:00pm Monday – 7:00am to 6:30pm Tuesday – 7:00am to 8:00pm Wednesday – 7:00am to 5:00pm Thursday – 7:00am to 12:00pm

Return to Top

Wednesday 8:00am-3:00pm

Wednesday 8:00am-3:00pm
Stetson F
Special Event · Private, Closed Meeting, Primary Room, Special Event

Advanced Computing for Social Change Institute

Kelly Gaither and Rosalia Gomez (TACC), Ruby Mendenhall (University of Illinois at Urbana-Champaign), Linda Akli (SURA), and Sue Fratkin (Fratkin Associates)
Biographies
Abstract

Advanced Computing for Social Change is a competition for undergraduate students from diverse disciplines and backgrounds who want to work collaboratively to:

Learn to apply data analysis and computational thinking to a social challenge Experience the latest tools and techniques for exploring data through visualization Expand skills in team-based problem solving Learn how to communicate ideas more effectively to the general public

To be eligible for this exciting program, you must:

Participants will be selected via an application/selection process. To be eligible, students must be full time undergraduate students at an accredited college/university; Not plan to graduate before May 2020; Have a minimum overall GPA of at least 2.5/4.0 (or equivalent); Be able to attend the full challenge, July 27-August 1; Complete the online application form before the deadline.

The costs of airfare, lodging, meals, and conference registration will be provided.

Return to Top

Wednesday 8:00am-5:00pm

Wednesday 8:00am-5:00pm
Crystal B
Posters/Visualization · Poster, Student Poster, Vis Showcase

Posters and Visualization

Abstract

All Exhibitors, Posters, and Visualization Showcase entries will be available at this time, for viewing at your convenience.

Return to Top

Wednesday 8:15am-10:30am

Wednesday 8:15am-10:30am
Regency AB
Plenary Session · Plenary, Platinum Exhibitor Plenary

Plenary Session B

Chair: Marisa Brazil (Campus Champions, Purdue University)

Redefining Today's HPC
Patricia Damkroger (Intel Corporation)
Biography
Abstract

In this talk we will look at the paradigm shift that is happening in HPC. It's no longer just simulation and modeling. It now includes other data-centric applications such as advanced analytics and AI. We will take a deep dive into what this means, the current and future trends we are seeing the market place and look at how industries today are already moving toward this new Data-Centric HPC.

Future of HPC in the Cloud
Ross Thomson (google.com)
Abstract

Today's research computing demands lightning-fast speed, vast data storage, and intensive processing power in order to advance discoveries across disciplines, from genomics to climate change. With high performance computing (HPC) in the cloud you can solve problems faster, reduce queue times for large batch workloads, and relieve compute resource limitations. Attend our talk as we explore the future of HPC in the cloud to help accelerate breakthroughs and unlock new scientific frontiers. We'll highlight how customers are breaking computing boundaries with Google Cloud Platform as well as discuss TPUs for accelerating data analysis and hybrid HPC.

Are There Closets in the Cloud?
Tim Carroll (microsoft.com)
Biography
Abstract

Commodity clusters had a profound change on computational science and more importantly, the science it supported. Cloud represents an equally compelling boost for research. Though HPC users have had the ability to run workloads on public clouds since 2007, the percentage of workloads running in production remains relatively low. This talk will examine the coming changes in both technology and funding mechanisms that will accelerate cloud adoption specifically for the academic community.

Return to Top

Wednesday 10:00am-5:00pm

Wednesday 10:00am-5:00pm
Toronto/Hong Kong/Acapulco
Exhibition · Exhibition

Exhibition

Joseph Paris
Abstract

The PEARC conference series works to integrate and meet the collective interests of our growing community by providing a forum for discussing challenges, opportunities, and solutions among the broad range of participants in the research computing community. The exhibitors program provides an opportunity to advance exciting new collaborations while shaping solutions through engaging with technology and service providers across both industry and academia.

Return to Top

Wednesday 10:30am-11:00am

Wednesday 10:30am-11:00am
Regency Foyer/Concourse Level Foyer/Crystal Foyer
Break · Break

Morning Break

Return to Top

Wednesday 10:45am-12:30pm

Wednesday 10:45am-12:30pm
Comiskey
Special Event · Special Event

Advanced Computing For Social Change Presentations

Advanced Computing for Social Change Institute
Kelly Gaither and Rosalia Gomez (TACC), Ruby Mendenhall (University of Illinois at Urbana-Champaign), Linda Akli (SURA), and Sue Fratkin (Fratkin Associates)
Biographies
Abstract

Advanced Computing for Social Change is a competition for undergraduate students from diverse disciplines and backgrounds who want to work collaboratively to:

Learn to apply data analysis and computational thinking to a social challenge Experience the latest tools and techniques for exploring data through visualization Expand skills in team-based problem solving Learn how to communicate ideas more effectively to the general public

To be eligible for this exciting program, you must:

Participants will be selected via an application/selection process. To be eligible, students must be full time undergraduate students at an accredited college/university; Not plan to graduate before May 2020; Have a minimum overall GPA of at least 2.5/4.0 (or equivalent); Be able to attend the full challenge, July 27-August 1; Complete the online application form before the deadline.

The costs of airfare, lodging, meals, and conference registration will be provided.

Return to Top

Wednesday 11:00am-12:30pm

Wednesday 11:00am-12:30pm
Crystal A
Tutorial · Tutorial Quarter-day

Introduction to Docker and Singularity

Ritu Arora, Charlie Dey, and John Cazes (Texas Advanced Computing Center)
Abstract

A container is a portable software unit that has one or more applications and all their dependencies bundled together in a single package. Containers make software and data distribution easy, and can thus save time in software installation and maintenance. They are also useful in ensuring the reliability and reproducibility of the applications by future-proofing them against the changes in hardware and system software. In this beginner-level, quarter-day tutorial (1.5 hours duration), the participants will get an overview of containers and their relevance in cloud computing and high performance computing. The tutorial presenters will demonstrate the use of Docker and Singularity – the two open-source container solutions that are popular in the advanced research community. A short hands-on session on Docker will also be included. The participants will be required to create an account on DockerHub to participate in the hands-on session and would need to bring their laptop.

pdf

Wednesday 11:00am-12:30pm
Regency AB
Paper · Facilitation, Technical Paper, Student Technical Paper, Cluster Management

Managing Clusters II

Chair: Christopher Phillips (Purdue University)

Harmony: A Harness Monitoring System for the Oak Ridge Leadership Computing Facility
Cameron Kuchta (Beloit College) and Reuben Budiardja and Verónica Vergara Larrea (Oak Ridge National Laboratory)
Abstract

Acceptance of a new system requires extensive testing and is often comprised of hundreds of tests. Summit, the latest flagship supercomputer at the Oak Ridge Leadership Computing Facility (OLCF), and the number one system in the Top500 list, completed its acceptance testing in 2018. To execute acceptance, the acceptance test (AT) team utilizes the OLCF test harness, a tool developed at the OLCF that automates the launch and verification of all acceptance tests. Acceptance requires analysis of test results and classification of all test failures. The sheer number of tests involved makes performing these tasks challenging. To complete these tasks more efficiently, in addition to lessen the personnel burden during acceptance testing, we developed a harness monitoring system for the OLCF test harness called Harmony.

Harmony consists of three distinct modules: monitoring, recording, and reporting modules.

Harmony's monitoring module ensures that tests launched by the harness are progressing in the job queue and restarted correctly after any failure. The recording system ingests results generated by the test harness into a database, and the reporting module provides a Django-based website to filter through tests, allowing staff to analyze, describe, and categorize any test failure.

Harmony is developed mainly in Python and uses MySQL, Django, and leverages the Slack and LSF APIs. The code is open source and publicly available. This paper presents Harmony's design and fully describes its modules it provides. Harmony's modular design allows it to be customized for other testing purposes and to be used at different HPC centers.

pdf

Arbiter: Dynamically Limiting Resource Consumption on Login Nodes
Dylan Gardner, Robben Migacz, and Brian Haymore (Center for High Performance Computing, University of Utah)
Abstract

In a high-performance computing (HPC) environment, shared resources are often the most capricious and unreliable. The performance of each process is affected by the behavior of all connected users. This applies particularly to the login nodes of HPC clusters, which are used by multiple people at a time. Often, policies govern their use, typically limiting users to small, lightweight tasks such as scripting, compiling, submitting batch jobs, and staging data. If a user is behaving poorly, such as using a significant portion of CPU time or consuming a large fraction of the total memory, processes using the same resources can be slowed or run out of memory. Policies on login nodes that prohibit such behavior, however, may not be enforced automatically because of technical limitations.

Arbiter is a daemon that overcomes such limitations by using cgroups, a feature of the Linux kernel, to monitor and limit usage. It can both enforce default limits—to ensure the server remains responsive—and penalize users who are consuming resources immoderately while still allowing for brief testing of workloads better suited to computational resources. Arbiter tracks the total memory and CPU usage of a user, reduces the quantity available for a period of time if usage is excessive, and sends emails to inform users of policies and potentially impactful behaviors. Throttling performed by Arbiter encourages users to use computational resources for intensive tasks (through the batch system), thereby supporting the reliability and responsiveness of login nodes.

pdf

A Resource Utilization Analytics Platform Using Grafana and Telegraf for the Savio Supercluster
Nicolas Chan (UC Berkeley)
Abstract

Understanding high performance computing cluster utilization patterns is key for decision making and efficient resource allocation. Cluster utilization statistics are useful for both cluster administrators and users, providing the ability to identify potential issues, plan jobs to minimize time spent in the queue, and identify constrained resources where funds might be focused in the future. Much of the existing data is accessible by the command line, but presenting the statistics visually allows for easier identification of trends and interactive exploration. Programs such as XDMoD exist to perform a similar function, but are not as well-suited towards use-cases of smaller clusters. Instead, we use a stack of open source software, allowing for a high degree of flexibility, including multiple database backends and more user-friendly querying and visualization. Based on open source software, similar software stacks can be deployed at other clusters to fit their existing infrastructure and needs. We present the workflow used by Berkeley Research Computing to consolidate existing data, collect additional utilization information, and display the relevant charts for different use-cases on the Savio supercluster. We have begun collecting, integrating, and displaying job information from Slurm; account association information; and CPU metrics collected with Telegraf. With this data, the Grafana visualization framework is able to present broad summary statistics, such as aggregated usage by campus department, all the way down to the CPU usage on a single node over the course of a job, and all with flexibility for a wide variety of possible needs and use-cases.

pdf

Wednesday 11:00am-12:30pm
Wrigley
Paper · Facilitation, Technical Paper, Cybersecurity, Student Technical Paper, Best Student Paper Facilitation

Cybersecurity

Chair: Lars Koesterke (TACC)

Trusted CI Experiences in Cybersecurity and Service to Open Science
Andrew Adams (Pittsburgh Supercomputing Center); Kay Avila and Jim Basney (NCSA); Dana Brunson (Internet2); Robert Cowles (BrightLite Information Security); Jeannette Dopheide and Terry Fleury (NCSA); Elisa Heymann (University of Wisconsin-Madison); Florence Hudson (Independent Consultant); Craig Jackson, Ryan Kiser, and Mark Krenz (Indiana University); Jim Marsteller (Pittsburgh Supercomputing Center); Barton Miller (University of Wisconsin-Madison); Sean Peisert (Berkeley Lab); Scott Russell, Susan Sons, and Von Welch (Indiana University); and John Zage (NCSA)
Abstract

This article describes experiences and lessons learned from the Trusted CI project, funded by the US National Science Foundation (NSF) to serve the community as the NSF Cybersecurity Center of Excellence (CCoE). Trusted CI is an effort to address cybersecurity for the open science community through a single organization that provides leadership, training, consulting, and knowledge to that community. The article describes the experiences and lessons learned of Trusted CI regarding both cybersecurity for open science and managing the process of providing centralized services to a broad and diverse community.

pdf

Managing Privilege and Access on Federated Edge Platforms
Christopher Weaver (University of Chicago), Joe Breen (University of Utah), Robert Gardner (University of Chicago), and Shawn McKee (University of Michigan)
Abstract

The SLATE (Services Layer at the Edge) platform supports collaborative, multi-institution scientific computing through federation of containerized edge services. This paper considers issues of trust between resource providers and developers of orchestrated services which span administrative domains. The context for discussion is the SLATE federation architecture and application deployment pattern. The major features are a custom central API server which implements the federation, partitioning of user and group applica- tions on edge clusters, and a curated catalog of applications which can be installed within the federation.

pdf

PULSAR:Deploying Network Monitoring and Intrusion Detection for the Science DMZ
Best Student Paper Facilitation
Shivam Trivedi, Lauren Featherstun, and Nathan DeMien (Purdue University); Callum Gundlach (William Henry Harrison High School); Sagar Narayan (Purdue University); Jacob Sharp (William Henry Harrison High School); and Brian Werts, Lipu Wu, Carolyn Ellis, Lev Gorenstein, Erik Gough, Alex Younts, and Xiao Zhu (Purdue University)
Abstract

The Purdue Live Security Analyzer (PULSAR) is a state-of-the-art, high speed network monitoring and intrusion detection system designed to enhance the security of Purdue University's research cyberinfrastructure. PULSAR project goals include empowering domain scientists to conduct research at Purdue with heightened cybersecurity requirements and engaging undergraduate students through the design, deployment and operation of advanced cyberinfrastructure. Deployment strategies and design decisions are discussed, ultimately providing a recipe book for other institutions to use as a guide for effective implementation of a large scale intrusion detection system for Science DMZs.

pdf

Wednesday 11:00am-12:30pm
Crystal C
Paper · Technical Paper, Student Technical Paper, Applications, Frameworks

Frameworks I

Chair: Shane Sanders (The Jackson Laboratory)

Web-based Visualization and Querying of Food and Beverage Endorsements by Celebrities
Srijith Rajamohan and Faiz Abidi (Virginia Tech)
Abstract

This paper explores the development of a web-based dashboard for the exploration of the nature of food and beverage endorsements by celebrities in the United States. While traditional methods of analysis require the translation of user requests to queries in one of the many languages that analysts use, the end user with no formal training in these languages is usually unable to get meaningful information without assistance. In this work, we propose the creation of a web-based dashboard that allows the user to use widgets to submit queries that can be executed on the dataset. To meet the design requirements of visualizing this data, two visual encodings are matched to the dataset under consideration, which illustrate the hierarchical relationships that exist between the various entities, within the confines of limited visual space. Finally, the dashboard concept is extended to a dataset larger than machine RAM using out-of-core processing that allows users to work on such data without resorting to distributed resources.

pdf

Extensible Framework for Analysis of Farm Practices and Programs
Sandeep Puthanveetil Satheesan (National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign); Rabin Bhattarai (Department of Agricultural and Biological Engineering, University of Illinois at Urbana-Champaign); Shannon Bradley (National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign); Jonathan Coppess (Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign); Lisa Gatzke (National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign); Rishabh Gupta and Hanseok Jeong (Department of Agricultural and Biological Engineering, University of Illinois at Urbana-Champaign); Jong Lee, Gowtham Naraharisetty, and Michal Ondrejcek (National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign); Gary Schnitkey (Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign); and Yan Zhao and Christopher Navarro (National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign)
Abstract

We present an open source extensible web framework for the analysis of different farm practices and programs and easy dissemination of their results to the users. Currently, this framework is being applied to two use cases - a web-based decision support system for cover crop management and a web-based farm program analysis tool to assist farmers, academics, and policymakers to understand programs and policies surrounding the Farm Bill. Through the first use case, we address the problem of bridging the gap between the scientific research that happens in labs and experimental plots and the day to day practices and decisions taken by the farmers in the fields. Specifically, this use case focuses on the practice of cover crops, their management, and the impact on reducing nutrient runoff into water bodies. Through the second use case, we address the problem of predicting the expected payout amounts and measured risk or probability of payment for different government insurance programs authorized by the 2018 Farm Bill, namely the Agriculture Risk Coverage (ARC) and Price Loss Coverage (PLC). This helps the farmers compare these programs based on forecasted crop yields and prices. In this paper, we describe the overall architecture of the framework and its major components, the use cases that are currently benefiting from using this framework and share screenshots of the web applications developed using this framework for those use cases. We also share our plans for future work and conclusions about applying this framework to the two use cases.

pdf

openDIEL: A Parallel Workflow Engine and Data Analytics Framework
Frank Betancourt and Kwai Wong (The University of Tennessee), Efosa Asemota and Quindell Marshall (Morehouse College), Daniel Nichols (The University of Tennesseee), and Stan Tomov (The University of Tennessee)
Abstract

openDIEL is a workflow engine that aims to give researchers and users of HPC an efficient way to coordinate, organize, and interconnect many disparate modules of computation in order to effectively utilize and allocate HPC resources. A GUI has been developed to aid in creating workflows, and allows for the specification of data science jobs, including specification neural network architectures, data processing, and hyperparameter tuning. Existing machine learning tools can be readily used in the openDIEL, allowing for easy experimentation with various models and approaches.

pdf

Wednesday 11:00am-12:30pm
Gold Coast
Paper · Technical Paper, Machine Learning/Artificial Intelligence, Adoption

Facilitate Machine Learning Adoption

Chair: Geoffrey Lentner (Purdue University)

Traps, Pitfalls and Misconceptions of Machine Learning applied to Scientific Disciplines
Davide Del Vento and Alessandro Fanfarillo (NCAR)
Abstract

In the last decade, Machine Learning has experienced a dramatic increase in performance on a wide variety of tasks, including computer vision, speech recognition, text parsing, and language translation, just to name a few. This has corresponded to an understandable hype especially for the remarkable results achieved in some cases. Therefore, practitioners of Scientific Disciplines have become interested in utilizing new Machine Learning techniques, and have sometimes started doing so with mixed success.

The purpose of this paper is to describe some of the common Traps, Pitfalls and Misconceptions of Machine Learning as relevant to the Scientific Discipline, and how to avoid them. In fact, in our experience, there is a number of Traps, Pitfalls and Misconceptions in which many people fall, but which could instead be easily avoided.

It is not the intention of the authors to provide any criticism to the work of other practitioners, particularly not to the ones working on the cutting edge of what is currently possible: in these cases expert researchers may well be doing exactly what we recommend here to avoid, and for a good reason. However we believe that the advice provided here will be useful, and perhaps even a reference, for the newcomers of the field.

pdf

Publishing and Serving Machine Learning Models with DLHub
Ryan Chard (Argonne National Lab); Logan Ward, Zhuozhao lI, Yadu Babuji, Anna Woodard, Steven Tuecke, Kyle Chard, and Ben Blaiszik (University of Chicago); and Ian Foster (Argonne National Laboratory)
Abstract

In this paper, we introduce the Data and Learning Hub for Science (DLHub). DLHub serves as a nexus for publishing, sharing, discovering, and reusing machine learning models. It provides a flexible publication platform that enables researchers to describe and deposit models by associating publication and model-specific metadata and assigning a persistent identifier for subsequent citation. DLHub also supports scalable model inference, allowing researchers to execute inference tasks using a distributed execution engine, containerized models, and Kubernetes. Here we describe DLHub and present four scientific use cases that illustrate how DLHub can be used to reliably, efficiently, and scalably integrate ML into scientific processes

pdf

Open Compass: Accelerating the Adoption of AI in Open Research
Paola Buitrago and Nicholas Nystrom (Carnegie Mellon University, Pittsburgh Supercomputing Center)
Abstract

Artificial intelligence (AI) has immense potential spanning research and industry. AI applications abound and are expanding rapidly, yet the methods, performance, and understanding of AI are in their infancy. Researchers face vexing issues such as how to improve performance, transferability, reliability, comprehensibility, and how better to train AI models with only limited data. Future progress depends on advances in hardware accelerators, software frameworks, system and architectures, and creating cross-cutting expertise between scientific and AI domains. Open Compass is an exploratory research project to conduct academic pilot studies on an advanced engineering testbed for artificial intelligence, the Compass Lab, culminating in the development and publication of best practices for the benefit of the broad scientific community. Open Compass includes the development of an ontology to describe the complex range of existing and emerging AI hardware technologies and the identification of benchmark problems that represent different challenges in training deep learning models. These benchmarks are then used to execute experiments in alternative advanced hardware solution architectures. Here we present the methodology of Open Compass and some preliminary results on analyzing the effects of different GPU types, memory, and topologies for popular deep learning models applicable to image processing.

pdf

Wednesday 11:00am-12:30pm
New Orleans
Panel · Panel

Data Science Support Services for Researchers

Christina Maimone and Christina Maimone (Northwestern University), Alex Storer (Stanford Graduate School of Business), Ista Zahn (Harvard T.H. Chan School of Public Health), and Simo Goshev (Boston College)
Abstract

Across fields, innovative research increasingly requires computational and data skills that domain researchers lack. The demand for support from researchers extends well beyond the computational resources and support services research computing groups have traditionally provided. Researchers also need assistance with computationally-driven data collection, wrangling, analysis, and visualization tasks, as well programming and software development support. The need for support ranges from qualitative analysis to machine learning.

As universities recognize this need, research computing groups, libraries, centers, and departments are being asked to provide these data science support services in a variety of ways. This panel provides an overview of what such services entail, how they are being provided at different institutions, and the challenges that come with providing data science support for researchers. Session participants and panelists will discuss factors to consider when developing data science support services for researchers, how they manage such services, and what careers for specialists providing support services look like.

pdf

Wednesday 11:00am-12:30pm
Water Tower
Panel · Panel

Stop Chasing Unicorns in the Global Gig Economy

Elizabeth Leake (University of Iowa, STEM-Trek Nonprofit); Ruth Marinshaw (Stanford University); Hakizumwami Birali Runesha (University of Chicago); Michael Papka (Argonne National Laboratory); Preston Smith (Purdue University); and Ben Rogers (University of Iowa)
Abstract

Stop Chasing Unicorns in the Global Gig Economy A PEARC19 Panel Submission by Ben Rogers (University of Iowa) Venue: Chicago Hyatt Regency Hotel (151 E. Wacker), July 28-August 1, 2019.

Most research computing center directors realize there's a global shortage of skilled personnel; they devote more time searching for qualified staff than they'd like to. Energy spent chasing unicorns could otherwise be used to support life-saving and transformative research. But since most directors ultimately compete for the same finite pool of financial and human resources, some might be reluctant to talk about their problems; they may fear it would expose a perceived weakness. But if these problems aren't addressed, they will have a detrimental impact on global research outcomes, national economic competitiveness and security, if they haven't already. This 90-minute panel discussion will explore recruiting and retention challenges faced by center directors who are willing to show their hands.

Format:

Panelists will present 10-minute talks on novel strategies they've used to attract, train and retain a larger pool of prospects. These lightening talks will be followed by a moderated Q&A for the balance of time.

The crux of the matter:

HPC systems administrators, cybersecurity specialists, scientific software analysts, network engineers, and systems architects with three or more years of experience are in high demand. Specialists who work for universities or government agencies are aggressively recruited from within and by commercial enterprises that can afford to pay twice as much and offer more attractive benefit packages.

At the same time, the workforce lacks diversity. Decades of broader engagement efforts have only been marginally successful. Most believe we should enlist a broader pool of prospects, and engage them earlier; make the workplace more welcoming for women, minorities, people with disabilities and veterans so more will feel comfortable entering and staying in the academic and government cyberinfrastructure workforce. Statistically, many drop out for various reasons. This panel will foster continued discussion about certain themes identified during the 2016 "Advancing Research Computing on Campuses (ARCC) conference panel titled, "A New Career Path: The Cyberinfrastructure Professional." ARCC16 was held on the University of Illinois Urbana-Champaign campus at the National Center for Supercomputing Applications. Attendees agreed that it had become increasingly difficult for campus-based centers to recruit and retain skilled staff. Universal demand for advanced skills peaked as university enrollments declined, federal government investments waned and state belts tightened. Corporate salaries increased, while university wages stagnated and new employee benefits were trimmed. To make it even worse, as soon as specialists are trained well enough to manage a trusted research computing environment—even student workers with a year or two of experience noted on their LinkedIn profiles—they are increasingly targeted by corporate recruiters and become a flight risk.

Since the ARCC panel in 2016, others have noted that the situation hasn't improved; in fact, it's worse in regions where employment upon graduation is uncertain, and more continue to live with their parents well into their 30's. Some students are even tempted to drop out of school and enter the workforce if they're approached by recruiters. They may have a difficult time visualizing the longitudinal benefits of a college degree. This situation presents a moral dilemma for university supervisors, advisers and parents who hope their students will graduate on time, but they also realize that attractive salary and benefit packages are difficult to refuse when many early-career professionals are under- or unemployed and defaulting on student loans.

But from a global workforce preparedness standpoint, once a specialist departs from academia, they leave a huge void; the ripple effect is felt around the world. With a graduating class of 200 computer science majors, you might only find two or three good candidates. University HPC sysadmin-trainers have an opportunity to reach students—they can identify and engage the needle-in-the-haystack prospects and guide them toward the career track. ________

For history and context, please read this February 13, 2019 HPCwire article by Elizabeth Leake (STEM-Trek Nonprofit; University of Iowa) who sat on the ARCC16 panel that was moderated by Gwen Jacobs (U-Hawaii); with panelists Ruth Marinshaw (Stanford); John Towns (NCSA/XSEDE); and Paul Wilson (UW-Madison).

https://www.hpcwire.com/2019/02/13/iowa-grows-its-own-to-fill-the-hpc-workforce-pipeline/

Leake has participated on many broader engagement programs through SC, ISC, SIAM-CSE, TeraGrid/XSEDE/PEARC and CHPC (South Africa) since 2008. Through her experience with these organizations and work for STEM-Trek, a global, grassroots nonprofit organization that supports HPC workforce development, she is confident that the shortage of skilled research computing personnel is a global problem, and the prospect pool is increasingly mobile. Pending her availability, Leake is invited to correspond about the panel discussion, and will share outcomes with HPCwire.

pdf

Wednesday 11:00am-12:00pm
Atlanta
Special Event · Champions, Special Event

Campus Champions Fellows

Campus Champions Fellows Presentations
Robert Sinkovits (San Diego Supercomputer Center) and Marisa Brazil (Purdue University)
Abstract

The XSEDE Campus Champions Fellows program partners Campus Champions with ECSS staff and research teams to work side by side for a year on real-world science and engineering projects. Fellows will develop expertise within varied areas of cyberinfrastructure, and they are already well positioned to share their advanced knowledge through their roles as the established conduits to students, administrators, professional staff, and faculty on their campuses. In addition to the technical knowledge gleaned from their experiences, the individual Fellows will benefit from their personal interactions with the ECSS staff and will acquire the skills necessary to manage similar user or research group project requests on their own campuses. The Campus Champions Fellows program is a unique opportunity for a select group of individuals to learn first-hand about the application of high-end cyberinfrastructure to challenging science and engineering problems.

The 2018-19 Fellows who will be sharing their experiences over the past year are: Peter Hawrylak, University of Tulsa; Chet Langin, Southern Illinois University at Carbondale; Xinlian Liu, Hood College; Gil Speyer, Arizona State University; and Mohammed Tanash, Kansas State University.

To learn more about the Campus Champions Fellows Program, visit:https://www.xsede.org/web/site/community-engagement/campus-champions/fellows.

Return to Top

Wednesday 12:30pm-1:30pm

Wednesday 12:30pm-1:30pm
Gallery 5,6,7
Student Program · Student Program, Private, Closed Meeting

Student-Mentor Lunch

Ester Soriano (NCSA)
Biography
Abstract

The Student-Mentor Lunch is a very well received event at the end of the conference. We have had it in PEARC17 and PEARC18 the last day of the conference. It gives an opportunity to all mentors and mentee to exchange impressions, anecdotes, photos, cards :) and do a lot of networking that may be of a lot of value for our students.

This year, we would like try "high tables" so everybody is standing up while eating and able to move around and chat with others

Wednesday 12:30pm-1:30pm
Regency CD
Lunch · Lunch

Lunch

Wednesday 12:30pm-1:30pm
Crystal C
Co-located Meeting · Co-located Meeting

XDMoD User Group Meeting

Tom Furlani (Roswell Park Comprehensive Cancer Center)

Return to Top

Wednesday 1:30pm-3:00pm

Wednesday 1:30pm-3:00pm
Comiskey
Student Program · Student Program

Student Modeling Presentations

Wednesday 1:30pm-3:00pm
Regency AB
Paper · Facilitation, Technical Paper, Student Technical Paper, Science Gateways

Science Gateways

Chair: Carlos Redondo (Texas Advanced Computing Center)

Serverless Science Gateway Development for Ca2+ binding site prediction on Amazon Web Services: Case Study
Melchizedek Mashiku and Neranjan Edirisinghe (Georgia State University)
Abstract

In this paper, we discuss the development of a science gateway; Identifying Ca2+ -binding sites in proteins using java application, developed by Dr. Jenny Yang at the Chemistry department, Georgia State University. Starting with a Protein Data Bank (PDB) x-ray or NMR structure file, MUGC application predicts calcium binding sites using a graph theory based algorithm. The project creates a science gateway to provide access to the MUGC algorithm using tools provided by Amazon Web Services. The full-stack solution uses S3 storage, AWS Lambda functions, and API gateway to relay the PDB files to the back-end computing in EC2. Architecture for a full stack serverless processing pipeline is implemented which allows users to access the application. The design is optimized for scalability, reliability, security, performance, and cost.

pdf

Petrel: A Programmatically Accessible Research Data Service
William Allcock and Benjamin Allen (ANL); Rachana Ananthakrishnan, Ben Blaiszik, and Kyle Chard (UChicago); Ryan Chard and Ian Foster (ANL); Lukasz Lacinski (UChicago); Michael Papka (ALN); and Rick Wagner (UChicago)
Abstract

We report on our experiences deploying and operating Petrel, a data service designed to support science projects that must organize and distribute large quantities of data. Building on a high-performance 3.2 PB parallel file system and embedded in Argonne National Laboratory's 100+ Gbps network fabric, Petrel leverages Science DMZ concepts and Globus APIs to provide application scientists with a high-speed, highly connected, and programmatically controllable data store. We describe Petrel's design, implementation, and usage and give representative examples to illustrate the many different ways in which scientists have employed the system.

pdf

ORCID Integration Into Science Gateways
Carrie Arnold and Maytal Dahan (Texas Advanced Computing Center, University of Texas at Austin); Ester Soriano (National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign); and David Hart (University Corporation for Atmospheric Research)
Abstract

With science gateways becoming more popular every day, both researchers and gateway developers are facing challenges in collecting, displaying, and verifying researcher career information, especially across multiple projects and organizations users are involved in. ORCID has attempted to simplify this for both researchers and organizations. ORCID is a non-profit organization that provides researchers with a persistent digital identifier, known as the ORCID iD, that links them with their work and allows them to share their activities with participating websites, including science gateways.

pdf

Wednesday 1:30pm-3:00pm
Crystal C
Paper · Technical Paper, Student Technical Paper, Applications, Frameworks

Frameworks II

Chair: Yinzhi Wang (Texas Advanced Computing Center)

Building A Scalable Forward Flux Sampling Framework using Big Data and HPC
Ryan DeFever, Walter Hanger, and Sapna Sarupria (Clemson University); Jon Kilganon (West Chester University of Pennsylvania); Amy Apon (Clemson University); and Linh Ngo (West Chester University of Pennsylvania)
Abstract

Forward flux sampling (FFS) is an established scientific method for sampling rare events in molecular simulations. However, as the difficulty of the scientific problem increases, the amount of data and the number of tasks required for FFS is challenging to manage with traditional scripting tools and languages for high performance computing. The SAFFIRE software framework has been developed to address these challenges. SAFFIRE utilizes Hadoop to manage a large number of tasks and data for large scale FFS simulations. The framework is shown to be highly scalable and able to the support large scale FFS simulations. This enables studies of rare events in complex molecular systems on commodity cluster computing systems.

pdf

Connecting Diffusion MRI to Skeletal Muscle Microstructure: Leveraging Meta-Models and GPU-acceleration
Noel Naughton (University of Illinois at Urbana Champaign) and John Georgiadis (Illinois Institute of Technology, University of Illinois at Urbana Champaign)
Abstract

Due to its non-invasive nature, diffusion-weighted MRI (dMRI) has shown promise as a method to quantify skeletal muscle's microstructure; however, connecting the dMRI signal of muscle to the underlying microstructure is difficult. Numerical models of dMRI can parameterize this relationship, but the associated computational expense has prohibited extensive use. Here, efficient numerical methods are presented to address this problem. In particular, a meta-model representation of a lattice Boltzmann model of dMRI is formulated and shown to be both accurate and several orders of magnitude faster to evaluate. It is also demonstrated how such a meta-model can help inform dMRI pulse profile selection in encoding microstructural information into the dMRI signal. Additionally, histologically-informed simulations are performed, allowing comparison of the numerical model's simplified parameterization with the more complex topology of skeletal muscle. Finally, an efficient inversion method is proposed to infer microstructural parameters of muscle from dMRI signal using a GPU-accelerated numerical model. The inversion method is able to infer microstructural parameters from dMRI signal when the underlying geometry matches the numerical model's, however, the simplified numerical model does not agree with simulations of more complex muscle tissue.

pdf

In-situ Analysis and Visualization of Earthquake Simulation
Dawei Mu (National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign); Jon Moran (GigaIO, Inc.); Hui Zhou (The Second Monitoring and Application Center, CEA); Yifeng Cui, Ronald Hawkins, and Mahidhar Tatineni (San Diego Supercomputer Center); and Steve Campbell (GigaIO, Inc)
Abstract

Nowadays, numerical simulation has played a vital role in analyzing and assessing earthquakes and their affects. Storage I/O performance and network bandwidth have not kept pace with the growth of computing power; as a result, post-processing has become a bottleneck to end-to-end simulation performance. One approach to solving this performance imbalance is to reduce the amount of output data by implementing in-situ visualization, which constructs the visualization concurrent with the simulation. In this paper, we propose a new software utility named "awp-odc-insitu" that is based on the well-known open-source seismic simulation software "awp-odc-os" and capable of performing in-situ visualization functionality by employing the open-source data analysis and visualization application "ParaView" and its in-situ library "ParaView Catalyst". Moreover, the paper discusses the implementation of in-situ functionality and analyzes the performance and efficiency of the "awp-odc-insitu" code to demonstrate the code is of potential use in practical seismic research.

pdf

Wednesday 1:30pm-3:00pm
Wrigley
Paper · Technical Paper, Workforce, Cyberinfrastructure Skill, Best Paper Workforce

CyberInfrastructure Skill Development I

Chair: Adam Erck (Doane University)

Cultivating the Cyberinfrastructure Workforce via an Intermediate/Advanced Virtual Residency Workshop
Henry Neeman (University of Oklahoma), Hussein Al-Azzawi (University of New Mexico), Dana Brunson (Internet2), William Burke (George Washington University), Dirk Colbry (Michigan State University), Jeff Falgout (United States Geological Survey), James Ferguson (University of Oklahoma), Sandra Gesing (University of Notre Dame), Joshua Gyllinsky (University of Rhode Island), Christopher Simmons (University of Texas at Dallas), Jason Simms (Lafayette College), Mohammed Tanash (Kansas State University), Daniel Voss (Beaumont Research Institute), Jason Wells (Bentley University), and Scott Yockel (Harvard University)
Abstract

Cyberinfrastructure (CI) Facilitation is the process of helping researchers to use research computing systems and services to advance their computing-intensive/data-intensive research goals. The growing need for CI Facilitation isn't being met by traditional academic degree and certificate programs, so informal education is required. The Virtual Residency (VR) is a program that teaches key CI Facilitation skills to CI Facilitators. Using a combination of (a) workshops, (b) biweekly conference calls, (c) a Grant Proposal Writing Apprenticeship and (d) a new Paper Writing Apprenticeship, the VR has been teaching CI Facilitation since 2015. During the summers of 2015-17, the annual VR workshop was at an introductory level, driving demand for a higher level workshop. In 2018, the VR workshop was presented at a level described as intermediate, but in practice it included a great deal of advanced content, concentrated on institutional CI leadership. The 2018 focus areas were: (1) in-depth CI expertise in areas of rapidly changing technology; (2) CI leadership; (3) funding acquisition skills; (4) outreach strategies, techniques, and skills; (5) communication skills. The 2018 VR workshop served 216 participants from 147 institutions in 42 US states and 2 US territories plus 2 other countries.

pdf

The CyberAmbassador Training Program
Astri Briliyanti, TJ Nguyen, Julie Wilson Rojewski, Kathleen Colbry, and Dirk Colbry (Michigan State University)
Abstract

The CyberAmbassador project aims to provide professional skills training for CyberInfrastructure (CI) professionals, with the goal of developing "CyberAmbassadors" who are prepared to facilitate and lead multidisciplinary, computationally-intensive research. This NSF funded program (Award Number 1730137) has the following objectives: (1) Develop curriculum that focuses on professional skills (communications, teamwork, leadership) within the context of large scale, multidisciplinary computational research; (2) Pilot, evaluate and revise the curriculum; (3) Develop "Train the Trainers" workshops to broaden the impact of the curriculum and connect with external partners to ensure the longevity of the program beyond the timeframe of the grant. This paper introduces the core curriculum, describes different modes of delivering content that we have piloted, shares preliminary evaluation results from two particularly relevant cases, and offers initial lessons learned.

pdf

The Professionalization of Cyberinfrastructure Personnel?
Best Paper Workforce
Nicholas Berente (University of Notre Dame); Stanley Ahalt (RENCI; University of North Carolina, Chapel Hill); James Bottum (Clemson University); Dana Brunson (Internet2); Joel Cutcher-Gershenfeld (Brandeis University); James Howison (University of Texas at Austin); John King (University of Michigan); Henry Neeman (University of Oklahoma); John Towns (University of Illinois, Urbana-Champagn; NCSA); Nancy Wilkins-Diehr (SDSC); and Susan Winter (University of Maryland)
Abstract

Cyberinfrastructure (CI) personnel are critical to next-generation scientific research. Although funding agencies and universities increasingly realize this, enabling resources do not grow as fast as the need for the CI skills and capabilities. There are two main issues with respect to CI personnel: CI personnel are scarce and cyberinfrastructure is different from traditional information technology ("IT"). A potential answer to these issues is to professionalize the CI field. In this article, we draw on organization science to argue that a strong form of professionalization does not make sense for the CI workforce. However, we do encourage the development of standardized roles for CI personnel, and we recommend starting with three such roles: system facing, software facing, and researcher facing CI roles. A credentialing effort around these three roles could strengthen the field as well. We conclude with a discussion of directions forward.

pdf

Wednesday 1:30pm-3:00pm
Gold Coast
Paper · Student Technical Paper, Machine Learning/Artificial Intelligence, Performance Evaluation, Phil Andrews Award, Best Student Paper Machine Learning

Performance Evaluation and Improvement

Chair: Weijia Xu (Texas Advanced Computing, University of Texas)

MagmaDNN: Accelerated Deep Learning Using MAGMA
Daniel Nichols, Kwai Wong, and Stanimire Tomov (University of Tennessee, Knoxville); Lucien Ng (Chinese University of Hong Kong); Sihan Chen (The Chinese University of Hong Kong); and Alex Gessinger (Slippery Rock University)
Abstract

MagmaDNN is a deep learning framework driven using the highly optimized MAGMA dense linear algebra package. The library offers comparable performance to other popular frameworks, such as TensorFlow, PyTorch, and Theano. C++ is used to implement the framework providing fast memory operations, direct cuda access, and compile time errors. Common neural network layers such as Fully Connected, Convolutional, Pooling, Flatten, and Dropout are included. Hyperparameter tuning is performed with a parallel grid search engine.

MagmaDNN uses several techniques to accelerate network training. For instance, convolutions are performed using the Winograd algorithm and FFTs. Another approach is MagmaDNNs custom memory manager, which is used to reduce expensive memory transfers. Finally, training is accelerated by distributing batches across GPU nodes.

This paper provides an overview of the MagmaDNN framework and how it leverages the MAGMA library to attain speed increases. This paper also addresses how deep networks are accelerated by training in parallel and further challenges with parallelization.

pdf

A Novel Pruning Method for Convolutional Neural Networks Based off Identifying Critical Filters
Phil Andrews Award, Best Student Paper Machine Learning
Mihaela Dimovska and Travis Johnston (Oak Ridge National Laboratory)
Abstract

Convolutional Neural Networks (CNNs) are one of the most extensively used tools in machine learning, but they are still not well understood and in many cases they are over-parameterized, leading to slow inference and impeding their deployment on low-power devices. In the last few years, many methods for decreasing the number of parameters in a network by pruning its output channels have been suggested, but a very recent work has argued that random pruning of channels performs on-par with state-of-the-art pruning methods. While random and other pruning methods might be effectively used for lowering the number of parameters in a CNN, none of these methods can be used to gain any further understanding of the model that the CNN has built. In this work, we propose a novel method for pruning a network, that at the same time can lead to a better understanding of what the individual filters of the network learn about the data. The method proposed aims to keep only the filters that are ``important'' for a class. We define a filter as important for a class if its removal has the highest negative impact on the accuracy for that class. We demonstrate that our method is better than random pruning on two networks used on the EMNIST and CIFAR10 datasets. By analyzing the important filters, we find that the important filters in the pruned networks learn features which are more general across classes. We demonstrate the importance and applicability of that observation in two transfer-learning tasks.

pdf

Characterizing the Execution of Deep Neural Networks on Collaborative Robots and Edge Devices
Matthew Merck, Bingyao Wang, Lixing Liu, Chunjun Jia, Arthur Siqueira, Qiusen Huang, Abhijeet Saraha, Dongsuk Lim, Jiashen Cao, Ramyad Hadidi, and Hyesoon Kim (Georgia Institute of Technology)
Abstract

In this paper, we examined and analyzed the effects of executing deep neural networks collaboratively on edge devices. Our method uses model parallelism to distribute tasks from DNN layers to Raspberry Pis, which are mounted on and powered by iRobots. Our measurements showed that while adding additional devices to a network of edge devices increases static energy, it decreases the dynamic energy used by each device. It also increases the average number of inferences performed by each device. We found that execution of DNNs on a distributed robot system can lead to unpredictable power consumption, which may in turn worsen Raspberry Pi performance. Another unreliable aspect is the network latency, which increases and decreases often based on a variety of often uncontrollable factors. Finally, through our measurement of the Raspberry Pi temperature during DNN execution, we found that DNN computation can increase temperature by as much as 16 degrees Celsius.

pdf

Wednesday 1:30pm-3:00pm
Water Tower
Panel · Panel

Transforming Scientific Computing with Cloud Technologies

Martial Michel (Data Machines Corp.), John Lowe (Indiana University), Timothy Randles (Los Alamos National Lab), Rion Dooley (Data Machines Corp.), Robert Killen (University of Michigan), Evan Bollig (Minnesota Supercomputing Institute), and Jonathan Mills (NASA)
Abstract

Cloud computing technologies bring expanded capabilities to enhance or even replace the traditional batch HPC model. The maturity of cloud technologies such as OpenStack or Kubernetes have enabled the coexistence and interoperability of HPC and the cloud. However, integrating HPC with public and private clouds poses architectural complexities and requires operator expertise.

In this panel, we will continue the conversation we started in 2018 and will discuss and engage the audience in topics of virtualized computing, hybrid cloud and hpc, utilizing cloud for infrastructure management, and non-traditional scientific workflows. These topics will include discussion of virtual machines, containers, bare metal infrastructure, and science gateways and their role in the scientific community.

pdf

Wednesday 1:30pm-3:00pm
Field
Co-located Meeting · Co-located Meeting

Eastern Regional Network: Simplifying Multi-Campus Research Collaborations

Barr von Oehsen (Rutgers Univerity)
Biography
Abstract

The Eastern Regional Network (ERN) is a regional organization dedicated to simplifying multi-campus collaborations and partnerships that advance the frontiers of research, pedagogy, and innovation. The ERN mission is to provide layered and transparent access to shared data and computing facilities for research projects located at partner sites throughout the region. The project team consists of leaders among university research computing providers, regional research and education networks, Internet2, and Commercial Cloud Providers who have been working together to address the growing need for a regional research platform to support a diverse set of science and education needs. Through a Kubernetes and SLURM hybrid solution, including edge components, Data Transfer Nodes (DTNs), and storage, coupled with regional NSF CC* and DIBBS projects, and the Open Science Grid (OSG), we connect a diverse set of local and cloud research platforms and Science DMZs into a single cohesive federated regional DataDMZ designed to facilitate the sharing of data and access to advanced computing resources that include traditional HPC, High Throughput Computing (HTC), HADOOP/SPARK, machine learning, and containerization. The resulting layered approach offers the research community across the region access to a broad range of services and resources that's not available on any one campus alone. This gathering is intended for ERN members attending PEARC to discuss action items outlined during a recent ERN workshop as well as give those interested in learning more about the ERN a chance to ask questions.

Return to Top

Wednesday 1:30pm-3:30pm

Wednesday 1:30pm-3:30pm
Stetson BC
Special Event · Hackathon

SGCI Hack@PEARC19

SGCI Hackathon (sciencegateways.org)
Abstract

Join us for three days of science gateway hacking with some of the best programmers in the field!

In conjunction with PEARC19 in Chicago, SGCI is sponsoring a Hackathon to challenge up-and-coming programmers for tomorrow's science gateways. Undergraduate and graduate students will compete on gateway projects established by SGCI partners and community members.

Join us on the social media app SLACK to get the latest updates! Follow us on Twitter at @sciencegateways The competition will take place from Sunday, July 28, to Wednesday, July 31, in the Hyatt Regency Chicago where the PEARC19 conference is being held. Final awards will be announced on Thursday, August 1.

Return to Top

Wednesday 1:30pm-5:00pm

Wednesday 1:30pm-5:00pm
New Orleans
Tutorial · Tutorial Half-day

Rogues Gallery: Addressing Post-Moore Computing

Jeffrey Young, Jason Riedy, and William Powell (Georgia Institute of Technology)
Abstract

The Rogues Gallery is a new collaborative, experimental testbed hosted at Georgia Tech that is focused on tackling "rogue" architectures for the post-Moore era of computing. Some of these devices have roots in the embedded and high-performance computing spaces, but many of the expected post-Moore technologies are limited to custom prototypes as with quantum, neuromorphic, and reversible computing devices.

This tutorial will present a brief overview of the Rogues Gallery including related tools and resources like benchmark suites for novel architectures and will focus on providing hands-on experience with two hardware rogues. Attendees will have an opportunity to learn about and program for the Emu Chick system, a near-memory computing architecture for sparse applications, and the Field Programmable Analog Array (FPAA), a mixed analog/digital platform for implementing machine learning and neuromorphic designs. We will provide and work through a set of demonstration codes based on Sparse Matrix-Vector Multiply for the Emu, and we will explore the open-source toolset and virtual machine image used to program the FPAA. Attendees will have an opportunity to continue their investigation into using post-Moore technologies by requesting a free account to access the Rogues Gallery at the end of the tutorial.

pdf

Wednesday 1:30pm-5:00pm
Crystal A
Tutorial · Tutorial Half-day

Managing HPC Software Complexity with Spack

Levi Baber (Iowa State University); Adam Stewart (University of Illinois, Urbana-Champaign); and Gregory Becker and Todd Gamblin (LLNL)
Abstract

The modern scientific software stack includes thousands of packages, from C, C++, and Fortran libraries, to packages written in interpreted languages like Python and R. HPC applications may depend on hundreds of packages spanning all of these ecosystems. To achieve high performance, they must also leverage low-level and difficult-to-build libraries such as MPI, BLAS, and LAPACK. Integrating this stack is extremely challenging. The complexity can be an obstacle to deployment at HPC sites and deters developers from building on each others' work.

Spack is an open source tool for HPC package management that simplifies building, installing, customizing, and sharing HPC software stacks. In the past few years, its adoption has grown rapidly: by end-users, by HPC developers, and by the world's largest HPC centers. Spack provides a powerful and flexible dependency model, a simple Python syntax for writing package build recipes, and a repository of over 3,000 community-maintained packages. This tutorial provides a thorough introduction to Spack's capabilities: installing and authoring packages, integrating Spack with development workflows, and using Spack for deployment at HPC facilities. Attendees will leave with foundational skills for using Spack to automate day-to-day tasks, along with deeper knowledge for applying Spack to advanced use cases.

pdf

Return to Top

Wednesday 3:00pm-3:30pm

Wednesday 3:00pm-3:30pm
Regency Foyer/Concourse Level Foyer/Crystal Foyer
Break · Break

Afternoon Break

Return to Top

Wednesday 3:30pm-5:00pm

Wednesday 3:30pm-5:00pm
Atlanta
Special Event · Hackathon

SGCI Hack@PEARC19 Final Presentations

SGCI Hack@PEARC19
SGCI Hackathon (sciencegateways.org)
Abstract

Join us for three days of science gateway hacking with some of the best programmers in the field!

In conjunction with PEARC19 in Chicago, SGCI is sponsoring a Hackathon to challenge up-and-coming programmers for tomorrow's science gateways. Undergraduate and graduate students will compete on gateway projects established by SGCI partners and community members.

Join us on the social media app SLACK to get the latest updates! Follow us on Twitter at @sciencegateways The competition will take place from Sunday, July 28, to Wednesday, July 31, in the Hyatt Regency Chicago where the PEARC19 conference is being held. Final awards will be announced on Thursday, August 1.

Wednesday 3:30pm-5:00pm
Regency AB
Paper · Facilitation, Technical Paper, Research Computing Analysis, Gold Exhibitor Technical Talk

HPC/Research Computing Analysis

Chair: Verónica G. Vergara Larrea (Oak Ridge National Laboratory)

Assessment of financial returns on investments in cyberinfrastructure facilities: A survey of current methods
Craig Stewart, David Hancock, and Julie Wernert (Indiana University Pervasive Technology Institute); Thomas Furlani (Roswell Park Comprehensive Cancer Center); David Lifka (Cornell University); Alan Sill (Texas Tech University); Nicholas Berente (University of Notre Dame); D.F. "Rick" McMuillen (Texas A&M University); Thomas Cheatham (University of Utah); Amy Apon (Clemson University); Ron Payne (University of Illinois, Urbana-Champaign); and Shawn Slavin (Indiana University Pervasive Technology Institute)
Abstract

In recent years considerable attention has been given to assessing the value of investments in cyberinfrastructure (CI). This paper includes a survey of current methods for the assessment of financial returns on investment (ROI) in CI. Applying the financial concept of ROI proves challenging with regard to a service that, in most academic environments, does not generate a "sold amount" such as one would find in the buying and selling of stocks. The paper concludes with a discussion of future research directions and challenges in the assessment of financial ROI in CI. This work is intended less as a definitive guide than as a starting point for further exploration in the assessment of CI's value for scientific research.

pdf

Research Computing Desktops: Demystifying research computing for non-Linux users
Abhinav Thota, Le Mai Weakley, and Ben Fulton (Indiana University); Alexander Younts and Stephen Harrell (Purdue University); H. E. Cicada Dennis, Laura Huber, Scott Michael, and Winona Snapp-Childs (Indiana University); and Daniel Dietz, Christopher Phillips, and Xiao Zhu (Purdue University)
Abstract

Many members of the current generation of students and researchers are accustomed to intuitive computing devices and never had to learn how to use command-line based systems, which comprise the majority of high-performance computing environments in use. In the 2013-14 time frame, both Indiana University and Purdue university separately launched virtual desktop front-ends for their high performance computing clusters with the aim of offering an easier on-ramp to new users. In the last five years we iterated on and refined these approaches, and we now have over two thousand annual active users combined. Over 75% of those users say that the desktop services are either moderately or extremely important for their ability to use HPC resources. In this paper, we share our experience bootstrapping this new service framework, bringing in the end-users, dealing with runaway success, and making this service a sustainable offering. This paper offers a comprehensive picture of the driving motivations for desktops at each institution, reasons users like desktops, and ways of getting started.

pdf

Practical Insights from the HPC and AI Innovation Lab
Alex Filby (Dell EMC, HPC and AI Innovation Lab)
Abstract

One of the primary roles of the Dell EMC HPC and AI Innovation Lab is to evaluate new technologies and different options for HPC and AI system design, then simplify these design choices by analyzing performance across a range of applications. This information is then used to influence the design of our Ready Solutions and shared with the HPC research community. This talk will cover what the Lab is and how customers can take advantage of its resources. We will focus on recent insights gathered from the lab in two main areas: a new BIOS tunable named Adaptive Double DRAM Device Correction (ADDDC) and how a seemly small change can have a large impact on performance. Then we'll discuss profiling tools commonly used in the lab and their application to specific real-world examples.

Wednesday 3:30pm-5:00pm
Gold Coast
Paper · Facilitation, Technical Paper, Student Technical Paper, Project Management, ML/AI Applications

Project Management/Soft Skills

Chair: Carolyn Ellis (Purdue University)

The biggest failure of my career: what happened, what we did about it, and how we won't do it again.
Jonathon Anderson (University of Colorado Boulder)
Abstract

In 2017 the Research Computing team at University of Colorado Boulder began the process of refreshing the PetaLibrary, a campus-wide research and academic data storage service. Updates were motivated by aging hardware, the impending end of upstream support, and a hope to improve the service with lessons learned.

We selected a supplier with a seemingly well-architected proposal; but problems experienced during deployment and testing increased our workload beyond expectations. We were required to operate the legacy infrastructure longer than expected, leading to stress within the team. Finally, inadequate planning and risk mitigation led to an over-reliance on project success and a reticence to enforce requirements and deadlines.

The project began to see success through more open communication and by making adequate allowances for partial project failure. We found it necessary to take a more active role in the project's success, including the removal of components that were unable to meet acceptability requirements. This, coupled with strong contracting, allowed us to amend the solution and avoid total project failure. The amended solution also matches our institutional priorities of openness and flexibility more closely than the original design.

From this experience we have learned to perform large and complex migrations in discrete stages; to decouple legacy decommissioning timelines from deployment timelines; to insist on and provide adequate time for hands-on evaluation of new components; to maintain architectural authority and responsibility; and at all times to assess risk frankly and communicate openly.

pdf

Good Enough Project Management Practices for Researcher Support Projects
Christina Maimone (Northwestern University)
Abstract

Successful provision of data and computational services in support of researcher projects requires implementing project management practices to counteract tendencies towards scope creep, miscommunication, and inaccurate resource forecasting. Good project management practices also support service metrics and reporting, which are essential for communicating the impact and value of researcher support services. These practices do not need to be complicated or onerous to be successful, but they do need to be applied intentionally and consistently. Key practices include: creating project summary documents, determining appropriate client management strategies, tracking project resources and status, and regularly reviewing projects' progress.

pdf

Improving HPC System Performance by Predicting Job Resources via Supervised Machine Learning
Mohammed Tanash, Brandon Dunn, Dan Andresen, Yang Huichen, Adedolapo Okanlawon, and William Hsu (Kansas State University)
Abstract

The majority of our HPC users come from other disciplines than Computer Science. HPC users including computer scientists have difficulties and do not feel proficient enough to decide the required amount of resources for their submitted jobs on the cluster. Consequently, users are encouraged to over-estimate resources for their submitted jobs, so their jobs will not be killing due insufficient resources. This process will waste and devours HPC resources; hence, will lead to inefficient cluster utilization. We created a supervised machine learning model and integrated it into the Slurm resource manager simulator to predict the amount of required memory resources (Memory) and the required amount of time to run the computation. Our model involved using different machine learning algorithms. Our goal is to integrate and test the proposed supervised machine learning model on the Slurm. We used over 1000 tasks selected from our HPC log files to evaluate the performance and the accuracy of our integrated model. The purpose of our work is to increase the performance of the Slurm by predicting the amount of the required jobs memory resources and the time required for each particular job in order to improve the utilization of the HPC system using our integrated supervised machine learning model.

Our results indicate that for large jobs our model helps dramatically reduce computational turnaround time (from five days to ten hours for large jobs), substantially increased utilization of the HPC system, and decreased the average waiting time for the submitted jobs

pdf

Wednesday 3:30pm-5:00pm
Crystal C
Paper · Technical Paper, Student Technical Paper, Applications, Frameworks

Frameworks III

Chair: Joey Jones (Mississippi State University)

Combining OpenCL and MPI to Support Heterogeneous Computing on a Cluster
Ming Li, Peter Hawrylak, and John Hale (The University of Tulsa)
Abstract

This paper presents an implementation of a heterogeneous programming model which combines Open Computing Language (OpenCL) and Message Passing Interface (MPI). The model is applied to solving a Markov decision process (MDP) with value iteration method. The performance test is conducted on a high performance computing cluster. At peak performance, the model is able to achieve a 57X speedup over a serial implementation. For an extremely large input MDP, which has 1,000,000 states, the obtained speedup is still over 12X, showing that this heterogeneous programming model can solve MDPs more efficiently than the serial solver does.

pdf

Scalable Parallel Programming in Python with Parsl
Yadu Babuji, Anna Woodard, and Zhuozhao Li (University of Chicago); Daniel Katz (University of Illinois at Urbana-Champaign); Ben Clifford (University of Chicago); Ian Foster (Argonne National Laboratory, University of Chicago); Michael Wilde (ParallelWorks); and Kyle Chard (University of Chicago)
Abstract

Python is increasingly the \emph{lingua franca} of scientific computing. It is used as a higher level language to wrap lower-level libraries and to compose scripts from various independent components. However, scaling and moving Python programs from laptops to supercomputers remains a challenge. Here we present Parsl, a parallel scripting library for Python. Parsl makes it straightforward for developers to implement parallelism in Python by annotating functions that can be executed asynchronously and in parallel, and to scale analyses from a laptop to thousands of nodes on a supercomputer or distributed system. We examine how Parsl is implemented, focusing on syntax and usage. We describe two scientific use cases in which Parsl's intuitive and scalable parallelism is used.

pdf

The PreSync Project: Synchronization Automation in the Cactus Framework
Samuel Cupp and Steven Brandt (Louisiana State University) and Roland Haas (University of Illinois at Urbana-Champaign)
Abstract

The Einstein Toolkit (ETK) is an open-source software platform for computational simulations in relativistic astrophysics and gravitational physics. The Cactus Framework (Cactus) comprises the core component of the ETK and handles creation of distributed data structures, parallelism, I/O, checkpointing, etc. The current Cactus scheduling system relies on the manual synchronization of ghost zones for grid functions (i.e. distributed matrices). Unfortunately, deciding which subroutines should synchronize which grid functions is a non-trivial problem. Incorrect synchronization can result in over-synchronization (a performance problem) or under-synchronization (an error). This issue also creates a barrier for new users and collaborators. We have developed a new scheduling and synchronization method for Cactus to improve the accessibility and efficiency of the ETK. Synchronization and boundary condition application are now handled automatically by Cactus, and performance has been greatly improved. The new approach requires each subroutine to have "read" and "write" declarations for individual grid functions. Before scheduled subroutines run, Cactus checks these declarations and performs any synchronization or boundary condition updates as needed. This method removes the difficulty of deciding where synchronization should take place and removes unneeded synchronizations. The only knowledge required to provide these declarations is what grid functions a given subroutine uses and on which parts of the grid they are read or written. In addition, these changes have reduced the number of synchronizations by approximately five percent in unigrid simulations. This new system is a first step in improving Cactus' scheduling methods to use more advanced and scalable techniques.

pdf

Wednesday 3:30pm-5:00pm
Wrigley
Paper · Technical Paper, Workforce, Cyberinfrastructure Skill

CyberInfrastructure Skill Development II

Chair: Adam Erck (Doane University)

An Expansion of the User Support Services for the Research Computing Group at the University of Colorado Boulder
Shelley Knuth (University of Colorado, Research Computing); Joel Frahm, Andrew Monaghan, Daniel Trahan, Zebula Sampedro, Aaron Holt, Timothy Dunn, and Thomas Hauser (University of Colorado-Boulder); and Patrick Burns (Colorado State University)
Abstract

In any research computing (RC) environment, the role of the user support group is a combination of education, consulting, and outreach. As this role is the most public-facing of a research computing group's team, it is important to ensure an excellent level of support is provided to users. The Research Computing group at the University of Colorado (CURC) Boulder maintains a large-scale computing cluster with several hundred active users, among other services, and has done so since 2011. The user support group at CURC provides a variety of services intended to support these users. This paper describes those services the CURC group provides, as well as explores the various ways that these services have been improved in 2018-19. The impact on users is also assessed.

pdf

Cyber Infrastructure: Training and Mentoring - A way to engage students in their technology future
Diana Dugas and D. Brian Ormand (New Mexico State University)
Abstract

The Cyber Infrastructure: Training and Mentoring (CI:TraM) pro-gram is designed to engage with local-area high school students and undergraduates in Las Cruces, New Mexico to teach them technology skills as well as provide career mentoring. We have seen a variety of results, mostly aligned with the amount of effort and involvement the student demonstrates. As part of the program, students leave behind the "do for me" mindset of high school and early college and enter the "how do I" mindset of successful adults. CI:TraM is an individualized program designed to give students the freedom to explore career options in-depth, as well as expose students to the technology needed in their desired career choice.Students engage with professionals to determine their career paths and determine how to reach those goals both academically and financially. Students learn core technology skills, such as Excel, and discipline-specific ones, such as Matlab and Auto-CAD.The CI:TraM program has helped engaged students to solidify their career paths and feel more confident both professionally and as persons on college campuses. Although longitudinal data is forthcoming, students have identified their believed benefits of the program and its impact on their lives.

pdf

The Blue Waters Petascale Institute: Longitudinal Impact and Assessment-Driven Development of an Intensive, Hands-on Curriculum for Teaching Applications in HPC
Aaron Weeden (The Shodor Education Foundation, Inc.); Colleen Heinemann (University of Illinois at Urbana-Champaign); Skylar Thompson (University of Washington); Cameron Foss (University of Massachusetts Amherst); Mobeen Ludin (Yukon College); and Jennifer Houchins (The Shodor Education Foundation, Inc.)
Abstract

The Blue Waters Petascale Institute was a two-week intensive training program that ran annually from 2014–2018. The institute implemented a hands-on, inquiry-based approach comprised of exploratory activities, engaging discussions, and informative lectures. These allowed students to learn parallel programming, high-performance computing, and computational science concepts with relevant examples from a wide range of real-world, domain-specific applications and scenarios. The goal was to provide undergraduate and beginning graduate students with the background and skills necessary to be successful not only during the institute but also in year-long research projects and into their future academic and professional careers. This paper examines the longitudinal impact of the institute on its students and recounts how the institute curriculum evolved over time based on formative and summative assessments. Results show that, overall, students considered the institute a positive experience and felt they gained the confidence and knowledge necessary to be successful in their internships and in their future careers. The institute has made a lasting impact in providing students with the ability to produce top-notch research outcomes of publishable quality, and students have taken what they learned into graduate programs, industry positions, and other career directions.

pdf

Wednesday 3:30pm-5:00pm
Water Tower
Panel · Panel

Jupyter at Scale: HPC and Scientific Facilities

Michael Milligan (University of Minnesota Supercomputing Institute); Aaron Culich (University of California, Berkeley); Rollin Thomas (NERSC, LBNL); and Rick Wagner (University of Chicago, Argonne National Laboratory; Globus)
Abstract

Project Jupyter has developed a mature and robust ecosystem of open-source tools and components supporting interactive computing. Users of high-performance computing (HPC) systems, especially those new to HPC and representing the next generation of scientists, have begun to expect Jupyter to be as ubiquitously available as bash or other popular shells. In order to meet and manage these expectations, facilities must simultaneously respond to demand while tackling relevant issues of policy, security, infrastructure support, and scale.

HPC users, researchers, and staff have begun exploring how Jupyter can expose HPC resources (storage, batch queues, workflow management, data transfer) in new and more accessible ways. The "Jupyter at Scale: HPC and Scientific Facilities" panel will bring together a group of practitioners to discuss their experiences deploying, maintaining, supporting, and extending Jupyter and related technologies at their respective institutions.

pdf

Return to Top

Wednesday 4:00pm-5:00pm

Wednesday 4:00pm-5:00pm
McCormick
Co-located Meeting · Co-located Meeting

Campus Research Computing Consortium (CaRCC) Futures & Sustainability Town Hall

Dana Brunson (Internet2); Thomas Cheatham (University of Utah); Patrick Schmitz (University of California, Berkeley); and Lauren Michael (University of Wisconsin - Madison)
Biography
Abstract

All are welcome! The Campus Research Computing Consortium (CaRCC) invites all to discuss and provide feedback on the CaRCC roadmap, and in particular to provide the community's initial conceptualization for possible models of sustainability and collaboration for the future of CaRCC and related or alternative communities as we move to support, advocate for, and enhance research computing and data, research IT, or research cyberinfrastructure on our campuses, in our regions and in the nation. This initial discussion will help frame a potential workshop on sustainability in the Fall of 2019.

CaRCC - the Campus Research Computing Consortium - is an organization of dedicated professionals developing, advocating for, and advancing campus research computing and data and associated professions.

Through various working groups and based on community input via the 2017 stakeholder survey, members of the CaRCC community have developed a robust structure to connect professionals in the field, referred to as the People Network with tracks for researcher-facing, system-facing, and data-facing professionals; a research computing and data job families guide; a document articulating the value propositions for campus research computing; brought together representatives from many communities of CI professionals in an ecosystem awareness workshop; and in collaboration with Internet2 and EDUCAUSE are developing a campus research IT maturity model.

CaRCC's origins are in the NSF award OAC-1620695 "RCN: Advancing Research and Education through a national network of campus research computing infrastructures - The CaRC Consortium."

For this meeting we welcome all who are interested in helping formulate future directions of CaRCC-facilitated (or potentially other community's-facilitated) activities to serve the broader campus research computing community and ways to sustain those activities beyond the current NSF grant funding.

Following this discussion at PEARC19, interested participants will be encouraged to join in further activities to refine and implement the ideas generated. If you are interested in participating in this but cannot make the meeting, please write to info@carcc.org

For more information about CaRCC please visit https://carcc.org/

Return to Top

Wednesday 5:00pm-10:00pm

Wednesday 5:00pm-10:00pm
Hyatt Main Lobby
Special Event · Networking

Chicago White Sox vs. New York Mets

Please meet in the Hyatt main lobby, East tower, near the restaurant, to gather prior to boarding the buses to the White Sox ballpark.

Abstract

Please meet in the Hyatt Main Lobby, East tower, near the restaurant, to gather prior to boarding the buses to the White Sox ballpark

Transportation to and from game provided, as well as food and beverage at the game. Bus departure times: 5:00 PM, 5:15 PM, 5:30 PM

Return to Top

Wednesday 6:30pm-10:00pm

Wednesday 6:30pm-10:00pm
Hyatt Main Lobby
Special Event · Networking

Chicago's First Lady Architectural Boat Tour

Please meet in the Hyatt main lobby, East tower, near the restaurant, to gather prior to walking to meet the boat.

Abstract

Please meet in the Hyatt Main Lobby, East tower, near the restaurant, to gather prior to walking to meet the boat.

Architectural Boat Tour with food and beverage included.

Return to Top

Thursday 7:00am-12:00pm

Thursday 7:00am-12:00pm
Regency Ballroom Registration Desk - Regency Foyer
Registration · Registration

Conference Registration and Information Desk

Abstract

Conference Registration and Information Desk Hours: Sunday – 12:00pm to 8:00pm Monday – 7:00am to 6:30pm Tuesday – 7:00am to 8:00pm Wednesday – 7:00am to 5:00pm Thursday – 7:00am to 12:00pm

Return to Top

Thursday 8:15am-10:00am

Thursday 8:15am-10:00am
Regency AB
Plenary Session · Plenary

Plenary Session Session C

Chair: Andrew Sherman (Yale University); John Towns (NCSA, XSEDE)

NSF/CISE and advanced cyberinfrastructure: an update and a look at, and over, the horizon
Jim Kurose (National Science Foundation)
▸ Biography
▾ Abstract

> In this talk we provide an overview, an update, and a look into the future for the National Science Foundation Directorate of Computer and Information Science and Engineering, with a particular emphasis on the Office of Advanced Cyberinfrastructure, and NSF's Big Ideas.

Town Hall
John Towns (National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign); Manish Parashar (National Science Foundation); Nick Nystrom (Pittsburgh Supercomputing Center); Shawn Strande (San Diego Supercomputer Center, University of California at San Diego); and Robert J. Harrison (Stony Brook University)
▾ Abstract

> In this session a very brief description of the NSF Advanced Computing Systems & Services: Adapting to the Rapid Evolution of Science and Engineering Research solicitation will be presented followed by brief presentations from the announced awardees of this solicitation. The resources will be deployed and in production in 2020 and will be available for allocations requests via XSEDE's resource allocations process.

Return to Top

Thursday 10:00am-10:30am

Thursday 10:00am-10:30am
Regency Foyer/Concourse Level Foyer/Crystal Foyer
Break · Break

Morning Break

Return to Top

Thursday 10:00am-12:00pm

Thursday 10:00am-12:00pm
Toronto/Hong Kong/Acapulco
Exhibition · Exhibition

Exhibition

Joseph Paris
▾ Abstract

> The PEARC conference series works to integrate and meet the collective interests of our growing community by providing a forum for discussing challenges, opportunities, and solutions among the broad range of participants in the research computing community. The exhibitors program provides an opportunity to advance exciting new collaborations while shaping solutions through engaging with technology and service providers across both industry and academia.

Return to Top

Thursday 10:30am-12:00pm

Thursday 10:30am-12:00pm
Water Tower
Paper · Data, Technical Paper, Workforce

Data Science

Chair: Kurt Showmaker (The University of Mississippi Medical Center)

Undergraduate Data Science and Diversity at Purdue University
Elizabett Hillery, Mark Ward, and Alex Younts (Purdue Univesity); Eric Adams (Purdue University); Jenna Rickus (Purdue Univesity); and Preston Smith (Purdue University)
▾ Abstract

> The vision of Purdue University's Integrative Data Science Initiative is to be at the forefront of advancing data science-enabled research and education. Tightly coupling theory, discovery, and applications, while providing students with an integrated data science-fluent campus ecosystem, this initiative is designed for college graduates at Purdue University to have Big Data experiences in academics. Faculty and staff of the university assist students through curriculum, research, residential life and professional development.
>
> In this paper, we present the framework for this initiative, which includes outlining course offerings, describing the residential living communities and workforce development on campus. We will finish by describing our initial successes.

pdf

Mining Online Training Log Data
Susan Mehringer and Christopher Myers (Cornell University, Center for Advanced Computing); Jennifer Houchins (The Shodor Education Foundation, Inc.); and Lorna Rivera (Georgia Institute of Technology; Center for Education Integrating Science, Mathematics, and Computing)
▾ Abstract

> Online training has been growing in popularity, and offers many advantages for both trainers and learners. Assessing the usage and impact of online material can be difficult, especially if content is made available to anyone and is not part of a course requiring formal enrollment. The Cornell Virtual Workshop (CVW) first offered online training on topics in high-performance computing and computational science in 1994, and ten years ago we began logging usage. We are now performing our first in-depth analysis of those log data to identify patterns in usage, so that we can better understand how users access the material, which types of topics and materials result in the greatest impact, how topic usage changes over time, and what types of presentation format might be preferred. While the CVW is built around a cohesive, sequential narrative for each training topic, we find that many users access our content in a more targeted fashion, suggesting that we rethink how we package our material. We anticipate that ongoing analysis using data science and machine learning methods will enable us to produce more useful training materials, and provide the educational community with valuable information about patterns in online material usage.

pdf

DeapSECURE: Empowering Students for Data- and Compute-Intensive Research in Cybersecurity through Training
Wirawan Purwanto, Hongyi Wu, Masha Sosonkina, and Karina Arcaute (Old Dominion University)
▾ Abstract

> As the volume and sophistication of cyber-attacks grow, cybersecurity researchers, engineers and practitioners rely on advanced cyberinfrastructure (CI) techniques like big data and machine learning, as well as advanced CI platforms, e.g., cloud and high-performance computing (HPC) to assess cyber risks, identify and mitigate threats, and achieve defense in depth. There is a training gap where current cybersecurity curricula at many universities do not introduce advanced CI techniques to future cybersecurity workforce. At Old Dominion University (ODU), we are bridging this gap through an innovative training program named DeapSECURE (Data-Enabled Advanced Training Program for Cyber Security Research and Education). We developed six non-degree training modules to expose cybersecurity students to advanced CI platforms and techniques rooted in big data, machine learning, neural networks, and high-performance programming. Each workshop includes a lecture providing the motivation and context for a CI technique, which is then examined during a hands-on session. The modules are delivered through (1) monthly workshops for ODU students, and (2) summer institutes for students from other universities and Research Experiences for Undergraduates participants. Future plan for the training program includes an online continuous learning community as an extension to the workshops, and all learning materials available as open educational resources, which will facilitate widespread adoption, adaptations, and contributions. The project leverages existing partnerships to ensure broad participation and adoption of advanced CI techniques in the cybersecurity community. We employ a rigorous evaluation plan rooted in diverse metrics of success to improve the curriculum and demonstrate its effectiveness.

pdf

Thursday 10:30am-12:00pm
Crystal A
Paper · Facilitation, Technical Paper, Cluster Management, Applications

Managing Clusters III

Chair: John Cazes (University of Texas at Austin)

Wrangling Rogues: A Case Study on Managing Experimental Post-Moore Architectures
Will Powell, Jason Riedy, Jeffrey Young, and Thomas Conte (Georgia Institute of Technology)
▾ Abstract

> The Rogues Gallery is a new experimental testbed that is focused on tackling "rogue'' architectures for the Post-Moore era of computing. While some of these devices have roots in the embedded and high-performance computing spaces, managing current and emerging technologies provides a challenge for system administration that are not always foreseen in traditional data center environments.
>
> We present an overview of the motivations and design of the initial Rogues Gallery testbed and cover some of the unique challenges that we have seen and foresee with upcoming hardware prototypes for future post-Moore research. Specifically, we cover the networking, identity management, scheduling of resources, and tools and sensor access aspects of the Rogues Gallery and techniques we have developed to manage these new platforms. We argue that current tools like the Slurm resource manager can support new rogues without major infrastructure changes.

pdf

Deploying a Top-100 Supercomputer for Large Parallel Workloads: the Niagara Supercomputer
Marcelo Ponce (University of Toronto, SciNet HPC Consortium); Ramses Van Zon, Scott Northrup, and Daniel Gruner (SciNet HPC Consortium, University of Toronto); Joseph Chen (Scinet HPC Consortium, University of Toronto); Fatih Ertinaz (SciNet HPC Consortium); Alexey Fedoseev and Leslie Groer (SciNet HPC Consortium, University of Toronto); Fei Mao (SciNet HPC Consortium); and Bruno Mundim, Michael Nolta, Jaime Pinto, Marco Saldarriaga, Vladimir Slavnic, Erik Spence, Ching-Hsing Yu, and Richard Peltier (SciNet HPC Consortium, University of Toronto)
▾ Abstract

> Niagara is currently the fastest supercomputer accessible to academics in Canada. It was deployed at the beginning of 2018 and has been serving the research community ever since. This homogeneous 60,000-core cluster, owned by the University of Toronto and operated by SciNet, was intended to enable large parallel jobs and has a measured performance of 3.02 petaflops, which put it at #53 in the June 2018 TOP500 list. It was designed to optimize throughput of a range of scientific codes running at scale, energy efficiency, and network and storage performance and capacity. It replaces two system that SciNet operated for over 8 years, the Tightly Coupled System (TCS) and the General Purpose Cluster (GPC). In this paper we will describe the transition process from these two systems, the procurement and deployment processes, as well as the unique features that make Niagara a one of a kind machine in Canada. We believe that there are important lessons and knowledge that can be transferred and applied in other supercomputer centers as ours.

pdf

Automatic Modulefile Generation in Environment Modules and Lmod Systems with mkmod
Kent Milfeld (TACC)
▾ Abstract

> After installation of an HPC application there is the chore of creating a modulefile for the modules system (MS), so that a user can easily load an environment for executing the application. The mkmod utility was designed to make modulefile generation simple. It features an auto-detection approach and a simple, intuitive input interface for creating the content of the modulefile, without knowing any details about the underlying language (Lua or Tcl). By also automatically integrating user modulefiles into the local MS, mkmod reduces the primary barriers that keep users from employing modules for their installed applications.

pdf

Thursday 10:30am-12:00pm
Gold Coast
Paper · Facilitation, Technical Paper, Student Technical Paper, Software, Best Paper Facilitation

Software

Chair: Scott Michael (Indiana University)

Providing a Unified Software Environment for Canada's National Advanced Computing Centers
Maxime Boissonneault (Université Laval, Compute Canada - Calcul Québec); Bart Oldeman (McGill University, Compute Canada - Calcul Québec); and Ryan Taylor (University of Victoria, Compute Canada - WestGrid)
▾ Abstract

> Exploiting an advanced computing platform consisting of several clusters distributed across the second-largest country in the world is challenging. Each cluster may run a different operating system, use a different generation of CPU, GPU, or network fabric, or be managed by a different team of system administrators. Presenting a unified software environment can tremendously facilitate the task of supporting researchers, but is challenging to implement. This is nevertheless what Compute Canada set out to do in 2016, in the midst of deploying a new generation of large clusters.
>
> We had to find software solutions to solve the challenges involved to achieve this goal. Distribution, portability and performance were three important technical criteria for us. We also had to consider the practicality of each approach for our users, and reproducibility of software installations performed by staff located at various sites across Canada.
>
> In this paper, we present the solution that we created, which has allowed Compute Canada to serve the needs of over 10,000 researchers across the country. This solution is used on over 20 different clusters with heterogeneous configurations, on processor architectures ranging from AMD's 2010 Magny-Cours to Intel's 2017 Skylake SP, with or without GPUs, with InfiniBand, Ethernet or OmniPath as the network fabric, and with Slurm or Torque/Moab as the scheduler. This stack provides a unified software environment to users, providing over 600 different scientific applications that are available in over 4,000 different combinations of version, compiler and CPU architecture.

pdf

A Continuous Integration-Based Framework for Software Management
Best Paper Facilitation
Samuel Khuvis, Zhi-Qiang You, Heechang Na, Scott Brozell, Eric Franz, Trey Dockendorf, Judith Gardiner, and Karen Tomko (Ohio Supercomputer Center)
▾ Abstract

> Automated software-installation tools such as EasyBuild and Spack have enabled HPC support staff to more quickly deploy software environments for HPC. However, to provide reliable user-facing software environments, any change made to the environments should be recorded and tested. In this paper, we will describe a continuous integration-based framework for software management in use at the Ohio Supercomputer Center (OSC). In this framework, we rely on homegrown tools for automated software build and installation and the ReFrame framework to setup the testing system which can be triggered on commit to a local Gitlab repository. We will document the workflow and explain how we can deploy the software environment quickly and confidently using these tools.

pdf

Mii: An Automated Search Engine for Environment Modules
Justin Stanley (Iowa State University)
▾ Abstract

> The use of environment modules in HPC systems is rapidly gaining traction due to the increasing complexity of managing research software. While environment modules improve organization for administrators, they can be unfriendly to end-users, especially those less familiar in UNIX-like shell environments. I created a search engine for modules that allows users to query available modules by the names of the commands they provide. Furthermore, I wrote integrations for the popular shells bash and zsh that completely automate the searching and the loading process.

pdf

Thursday 10:30am-12:00pm
Crystal C
Paper · Technical Paper, Applications, Collaboration

Collaborative Applications

Chair: Shane Sanders (The Jackson Laboratory)

GenApp, Containers and Abaco
Emre Brookes (University of Texas Health Science Center at San Antonio) and Joe Stubbs (Texas Advanced Computing Center)
▾ Abstract

> GenApp is an NSF-funded framework for rapid generation of applications including feature rich science gateways. GenApp is being successfully used to produce science gateways wrapping scientific programs. Its organization is designed to simplify the process of adding new features and capabilities to generated applications. A limited set of definition files define application generation. To bring a new executable into GenApp, one creates a single "module" definition file. The executable must run on some compute resource accessible by the generated application. Installations of the executable on target resources may be complex. To simplify portability of execution, we introduce automatic containerization of defined modules and integration of container execution. Abaco is an NSF-funded web service and distributed computing platform providing functions-as-a-service (FaaS) to the research computing community. Abaco implements functions using the Actor Model of concurrent computation. We introduce GenApp integration of execution with Abaco as a resource.

pdf

Enabling Real-time User Interaction For Decision Support: Experiences Extending A Local Agave Platform Metadata Service.
Sean Cleveland (University of Hawaii); Jared McLean (University of Hawaii - Hilo); Kolja Rotzoll (University of Hawai'i- WaterResources Research Cen); Scot Izuka (U.S Geological Survey, Pacific IslandsWater Science Cent); and Gwen Jacobs (University of Hawaii)
▾ Abstract

> The University of Hawai`i Information Technology Services Cyberinfrastructure team in partnership with the United States Geological Survey developed the Hawai`i groundwater recharge tool, a decision support tool, as part of the `Ike Wai Gateway to support water sustainability research for the state of Hawai`i. To enable the development of the tool within the existing University of Hawai`i Agave platform that the `Ike Wai Gateway is built on, the Metadata application programmatic interface (API) and Agave infrastructure needed to be extended to store and retrieve spatial data and handle a larger number of requests. The development team successfully added the functionality to the local University of Hawai`i Agave platform instance to support the storage of deserialized vector data and metadata that can be spatially queried across hundreds of thousands of records to enable decision support applications within a science gateway.

pdf

HUBzero®: Novel Concepts Applied to Established Computing Infrastructures to Address Communities' Needs
Sandra Gesing (University of Notre Dame) and Claire Stirm, Michael Zentner, Ben Haley, and Steve Clark (Purdue University)
▾ Abstract

> The science gateway framework HUBzero® has been enhanced and further developed since its initial vision in 1996 – always driven by requirements of the diverse communities applying HUBzero® for their research. HUBzero® is part of a computational landscape that has never evolved as fast as in the last decade. Novel frameworks and concepts on the user interface side such as Javascript libraries and Jupyter notebooks support communities in their working environment with easy-to-use user interfaces while novel technologies and concepts in the backend allow for effective and efficient modeling, simulations and processing research tools and data. HUBzero®'s enhancements include extensions for BOINC and XSEDE infrastructures on the backend while offering interactive computations and analytical tools via Jupyter Notebooks, RStudio, and other web applications as publishing environments. The paper goes into detail for novel developments for the three use cases PURR, nanoHUB and MyGeohub.

pdf

Return to Top

Thursday 12:00pm-1:30pm

Thursday 12:00pm-1:30pm
Regency AB
Lunch · Plenary, Lunch, Awards

Awards Lunch

Return to Top

Thursday 1:30pm-3:00pm

Thursday 1:30pm-3:00pm
Water Tower
Co-located Meeting · Private, Closed Meeting, Co-located Meeting

XSEDE All-Staff Meeting

John Towns (XSEDE)
▾ Abstract

> We would like for this event to occur immediately following the close of PEARC19. We expect 40-50 in attendance. The event should be 90 minutes.

Return to Top

Thursday 3:00pm-4:00pm

Thursday 3:00pm-4:00pm
Water Tower
Co-located Meeting · Private, Closed Meeting, Co-located Meeting

XSEDE ECSS All-Hands Meeting

Robert Sinkovits (SDSC) and Philip Blood (PSC)
▾ Abstract

> We will take this opportunity to interact face-to-face as an XSEDE ECSS team and address topics that will be chosen by ECSS leadership and staff prior to the meeting.

Return to Top
