Cornell University
Center for Advanced Computing

Cornell News Highlights

Fall 2018/Winter 2019

Contact: Paul Redfern, red@cac.cornell.edu

Cornell Named Training Partner for $60 Million Frontera Supercomputer Project

Cornell CAC staff members will develop online, asynchronous training modules for the $60 million NSF-funded *Frontera* supercomputer to be deployed at the Texas Advanced Computing Center (TACC) in the summer of 2019. As an operations partner in the Frontera project, Cornell will author content for online courses using its proven *Cornell Virtual Workshop (CVW)* platform, as well as adapt content created by the rest of the Frontera team.

Cornell Part of $25M NSF Effort to Untangle Future Physics Data

Laboratory of Elementary Particle Physics director Peter Wittich and his team, including CAC senior research associate Steve Lantz, will develop software to handle the tsunami of data expected from the CERN LHC upgrade. The Cornell effort is part of the *Institute for Research and Innovation in Software for High-Energy Physics (IRIS-HEP)* project.

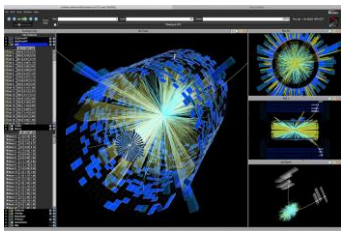

Lifka Serving as NSF CISE Advisory Committee Member

David Lifka, Cornell's vice president for information technologies and CIO, and director of the Center for Advanced Computing is serving on the NSF Directorates for Computer, Information Science & Engineering Advisory Committee. Members of the Committee provide advice on the impact of NSF support policies and programs on the CISE community, provide oversight on program management and performance, and provide advice to the CISE Assistant Director on special issues.

Cornell Digital Agriculture Team Wins $1 Million Tulane prize

The Cornell Computational Agriculture Initiative was funded by the USDA to solve agricultural problems using advanced computing technologies. CAC projects included collaborating with scientists to improve high-resolution weather data accessibility for farmers. Today, results of that initial collaboration and years of subsequent research have come to fruition: Cornell developers won the $1 million *Tulane Nitrogen Reduction Challenge* prize for *Adapt-N*, a tool used to improve nitrogen management for corn. Yara USA plans to distribute Adapt-N worldwide.

adapt-N

Cornell Investigating Open Cloud Marketplace

Cornell and its University at Buffalo and UC Santa Barbara partners received a $995K supplemental award from the NSF to investigate how to integrate public cloud into the *Aristotle Cloud Federation* portal in order to provide researchers with the information they need to compare and select public clouds for research. Learn more about the Aristotle project at https://federatedcloud.org/.

Cloud Usage Grows in Social Sciences and Economics

More than thirty faculty, staff, and student researchers affiliated with the *Cornell Institute for Social and Economic Research (CISER)* recently shifted some of their research workloads from clusters to *Red Cloud*.

Use of Cornell's on-premise cloud by CISER has grown to over 200 Red Cloud subscriptions. Thus far researchers find the on-demand resource effective and timely in meeting their computational and statistical analysis needs. Social sciences computing can be an ideal fit for the cloud considering statistical software is often used in a pleasingly parallel mode at moderate scale. Researchers especially like the availability of large memory instances on Red Cloud: 28 core instances have 192GB RAM. CAC staff Kim Burlingame and Ben Trumbore built 5 ready-to-use cloud images for the researchers: a Windows image with IBMSPSS, MATLAB, Mathematica. R, Rscript, RStudio, SAS, and Stata/MP and 4 custom Linux images that include Python, Miniconda, PostgreSQL, DataGrip, Kate editor, QIME2, and PICRUSt.

Cornell Tech Campus Built for the Digital Age

The Next Web reports that New York currently boasts 6,300 startups, with tech jobs growing at a rate of 30 percent. Billions have been invested in incubators, accelerators, and universities. One notable player is Cornell's campus on Roosevelt Island whose 2017 dedication was attended by New York governor Andrew Cuomo, New York City mayor Bill de Blasio, former mayor Mike Bloomberg (who donated $200 million to help start *Cornell Tech*), Cornell University president Martha Pollack, Technion president Peretz Lavie, and Cornell Tech dean Daniel Huttenlocher.

Cornell Tech is the first campus ever built for the digital age, bringing together academia and industry to create pioneering leaders and transformational new research, products, companies, and social ventures. Faculty and graduate students collaborate extensively with tech-oriented companies and pursue their own startups. When fully completed, the campus will include 2 million square feet of buildings, over 2 acres of open space, and will be home to more than 2,000 graduate students and hundreds of faculty and staff.

Cornell Engineering Specialties Ranked Top 10

Eight Cornell engineering specialties were ranked in the top 10 by the 2019 *US News & World Report* Best National University Rankings. They include biological/agricultural, civil, computer, electrical/electronic/, environmental/environmental health, industrial/manufacturing, materials, and mechanical. Engineering also ranked in the top 10 for undergraduate engineering programs in Ph.D.-granting institutions.

Fast Radio Bursts "Twist and Shout"

An international team of astronomers has found that the Cornell-discovered fast radio burst FRB 121102—a brief, gigantic pulse of radio waves from 3 billion light years away—passes through a veil of magnetized plasma. This causes the blasts to "twist and shout," which will help scientists determine the source. This finding was featured on the cover of *Nature.*

The "shouting" represents the bursts, and the "twisting" describes a physical phenomenon called Faraday rotation, which occurs as radio waves pass through a magnetized plasma, explained James Cordes, Professor of Astronomy. Measurement of the twisting provides further scientific detail on the origin of FRB 121102. The data were culled the Arecibo Observatory and confirmed by the Green Bank Observatory. The radio source and its environment are unique and not seen previously. In a flash lasting less than a millisecond, the burst radiates energy equal to our sun's output for an entire day – from the *Cornell Chronicle*.

New XSEDE Online Training Topics Available

New and updated *Cornell Virtual Workshop* training topics are available at the XSEDE user portal—*Introduction to Jetstream*, *Profiling and Optimization on KNL Clusters*, *Vectorization*, and *Introduction to R*. *Cornell Virtual Workshops* are freely available at all times to the entire scientific community of researchers, HPC practitioners, students, and educators at https://portal.xsede.org/online-training/. Additional topics are under development and will be available soon. These include *Scientific Computing with Python*, *Python for High Performance*, *Introduction to Wrangler*, and *Using the Jetstream API*. Cornell's online training is developed in collaboration with cyberinfrastructure experts at leading universities and centers, including the Texas Advanced Computing Center, the Indiana University Pervasive Technology Institute, and UC Berkeley.

Myers Joins CAC

Chris Myers, Ph.D. Physics, has joined the Center for Advanced Computing as a senior research associate. Besides providing Cornell faculty research support, Myers is co-PI on a $1.9M NIH grant examining the spread of diseases in bees. The decline in bee populations is a major concern as pollinators—especially wild and managed bees—are critical to native ecosystems and agricultural crops, providing the equivalent of billions of dollars in pollination annually. Myers is also an adjunct professor in the Department of Physics and is currently developing and teaching CAC's first *eCornell* course called *Introduction to Python for Data Science*.

XSEDE Software Toolkits Available for Campus Systems

Speaking at the *NERCOMP Conference*, CAC deputy director Rich Knepper said providing computational resources to researchers can be difficult, especially in terms of getting the necessary expertise to meet research needs. The NSF-funded XSEDE project provides access not only to resources but also to software toolkits that can be implemented on campus resources with minimal difficulty. XSEDE also offers consulting to help implement local resources that look like the supercomputing systems available at national centers. Knepper and his XSEDE team, including CAC staff, are building an *XSEDE Community Software Repository* that allows schools to implement a range of software, including cloud solutions, to support research computation needs and create a community of practice among implementing institutions. Visit the XSEDE CSR at https://software.xsede.org/.

Aristotle REU Students Make Meaningful Contributions to Science

Eleven NSF-funded Research Experience for Undergraduates (REU) students made meaningful contributions to *Aristotle Cloud Federation* science use case projects to date at Cornell University and the UC Santa Barbara.

In summer 2018 at Cornell, Plato Deliyannis implemented a flexible radio astronomy pipeline architecture and coded a friend-of-friends algorithm; Cindy Wu wrote a script to construct multiple species microbial community models and code to analyze the exchange reactions; and, Peter Cook analyzed the operating conditions for wind turbines and wrote scripts to analyze Weather Research and Forecasting models using an array of statistical methods and learned how to employ parallel computing for efficiency. At UC Santa Barbara, William Berman worked on DrAFTS, a cloud spot pricing analysis tool, and Gareth George, revamped the DrAFTS data management infrastructure and analyzed changes in AWS spot pricing, and is now engaged in developing a portable version of AWS Lambda so that science applications that use Lambda can be executed in remote locations. The Aristotle project is grateful to the NSF for funding these student experiences.

Dartmouth Joins Aristotle Cloud Federation to Explore Federated Cloud Computing Model

Dartmouth is joining the Aristotle Cloud Federation, a group of universities that are working together to build a cloud computing resources to be used for academic research. Dartmouth joins project leader Cornell University and co-PIs University at Buffalo and UC Santa Barbara, in exploring the computing model, which the three schools have been building under a NSF Office of Advanced Cyberinfrastructure award. “Were excited about this opportunity,” says Mitchel Davis, vice president and chief information officer at Dartmouth. “A federated cloud model has the potential to facilitate resource-sharing between campuses and is emerging as an important consideration in cyberinfrastructure planning. Davis says he hopes to provide Dartmouth researchers with access to cost-effective, on-demand systems, software, and data in a hybrid fashion that will let them choose from private, public, and federated cloud offerings.

Brazier Leading Multi-Messenger Astrophysics Cyberinfrastructure Planning

Multi-Messenger Astrophysics is an exciting new field of science that combines traditional astronomy with the brand new ability to measure phenomena such as gravitational waves and high-energy neutrino particles that originate from celestial objects. Under a recently announced NSF award, Cornell CAC computational scientist Adam Brazier will lead community planning for scalable cyberinfrastructure to support multi-messenger astrophysics. Brazier will be responsible for providing an in-depth analysis of cyberinfrastructure needs and collaborations, and for developing a strategic plan for a scalable cyberinfrastructure institute, including its mission and services.

The Cornell University Center for Advanced Computing (CAC) is supported by the Cornell Office of the Vice Provost for Research, with additional support from the National Science Foundation, National Institutes of Health and other partners and supporters, including Amazon, Boeing, Dell, eCornell, Globus, Johns Hopkins University, Indiana University, MathWorks, Red Hat, Texas Advanced Computing Center, University at Buffalo, and UCSB. Over 85 companies, universities and, public agencies have participated in our Partner Program since its inception. Visit https://www.cac.cornell.edu or email red@cac.corenll.edu to learn more about the Center.