Center for High-Performance Computing

NEWS

THE UNIVERSITY OF UTAH

FALL 2015

Research

Watching Nanomaterials Assemble at CHPC

By Prof. Michael Grünwald, Department of Chemistry

My son and I like to build remote control cars. The path that leads from a disordered pile of plastic parts and metal screws to a fine race car is straightforward and fun: step after step, we collect the pieces that need to be assembled and put them together according to the instructions. In fact, this assembly strategy is the blueprint for much human building activity and applies almost generally to the construction of houses, machines, furniture (in particular the Swedish kind), and many other objects of our daily lives.

Large objects, that is. Building small things, as it turns out, requires a strikingly different approach. Consider, for instance, the "objects" illustrated in Figure 1: A porous crystal structure made from intricately arranged metal ions and organic molecules (a "metal-organic framework"), and an ordered arrangement of nanoparticles (a "superstructure"), which themselves consist of many thousands of atoms. These structures are examples of "nanomaterials", objects that derive their unusual properties from their fascinating microscopic structure. Because of their large pores, metal-organic frameworks like the one in Figure 1a can be used to store hydrogen gas, filter CO_2, or separate molecules by shape. Depending on the kinds of nanoparticles used, superstructures such as the one in Figure 1b can be used to alter the direction of light, or act as new kinds of solar cells.

How should we go about making these materials? Can we not use our expertise gained from building cars and houses and simply arrange the building blocks (atoms, molecules, or nanoparticles) in the desired pattern? Not likely. In contrast to the bricks that make up a house (their number is approximately 5,000, according to Google), or the parts that make up a car (approximately 30,000), the building blocks of nanomaterials come with two crucial challenges: they are too many and they are too small. While we do have the technology to move single molecules or nanoparticles around, we certainly lack the time and money required to build entire solar cells consisting of 10^{20} particles.

Luckily, Nature has solved the problem of building with small blocks for us. One glimpse inside the human body reveals a breathtaking array of "bio-nanomaterials", intricate patterns and fascinating machines that perform complex functions. Of course, from the cell membrane, to the ribosome, to actin filament networks, these materials are not being built like cars. Rather, they arise through a process called self-assembly, in which bio-molecular building blocks spontaneously arrange to form a complex pattern amidst the thermal noise of perpetually jiggling molecules in solution. The "instructions" on how to build these patterns are encoded in the mutual interactions of building blocks.

(Continued on Page 2)

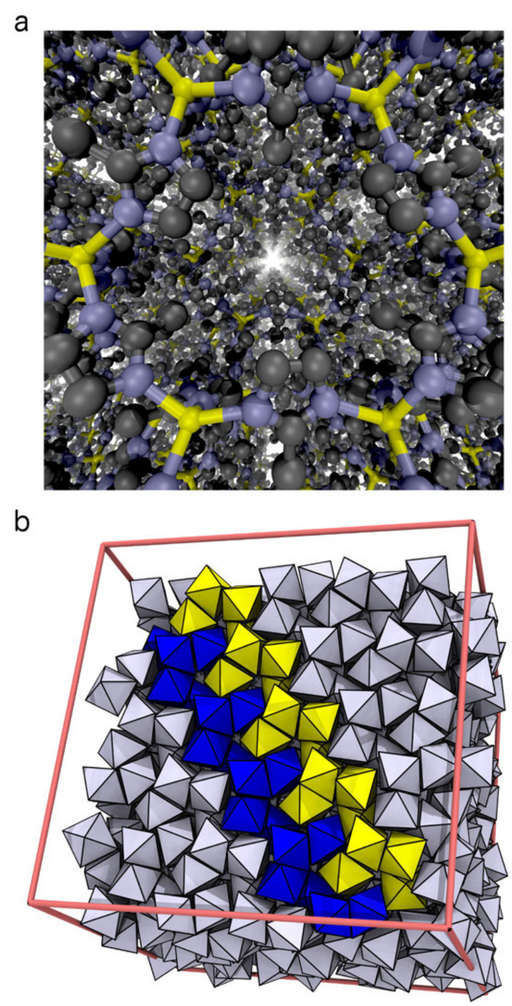

Figure 1: Snapshots from simulations of two types of nanomaterials. (a) A highly porous metal-organic framework (ZIF-8), consisting of Zn ions (yellow spheres) and methylimidazolate linkers (nitrogen atoms are colored blue, carbon atoms are colored gray, hydrogen atoms are not shown). (b) A superstructure formed from octahedral silver nanocrystals. The pink frame indicates the boundaries of the simulated region. A few nanocrystals are colored yellow and blue to highlight features of the complex structure they form.

(Continued from Page 1)

Inspired by Nature's example, self-assembly has become a main building strategy for synthetic nanomaterials. The idea is quite appealing: Throw a large number of nanoparticle building blocks into solution and recover the final self-assembled material, cheaply and quickly. Unfortunately, things rarely go that smoothly and the formation of erroneous patterns or structures different from the desired one are the rule rather than the exception. The problem lies in the deficient set of instructions that we pass on to our synthetic building blocks. Interactions between the building blocks that assemble in living systems have been optimized by evolution over the course of some 4 billion years; humans, on the other hand, have been in the business of self-assembly for merely 20 years and we still know very little about how to best engineer our nanoparticle building blocks to assemble into a desired nanomaterial.

In our research group, we address this problem with the help of computer simulations. What kinds of mistakes can happen during self-assembly? How can nanoparticle interactions be adjusted to avoid these defects, or induce assembly of a particular pattern? How do experimental conditions influence self-assembly outcomes? To answer these questions, we build mathematical models of nanoparticle building blocks and explore their self-assembly with molecular dynamics computer simulations. By sampling many trajectories of self-assembly, we collect the necessary statistics to learn about the microscopic mechanisms of nanomaterial formation. We analyze these trajectories with methods of statistical mechanics and enhanced sampling techniques.

The complexity of the models we use depends very much on the physical system we are interested in and is limited by the capabilities of today's computing hardware. While interactions between atoms and nanoparticles are in principle determined by the quantum nature of electrons, numerically evaluating the physical laws of quantum mechanics is only feasible for short, picosecond-long simulations involving a few atoms. Self-assembly, by contrast, involves many thousands of atoms or more, and time scales of microseconds or longer. We therefore use classical "force-fields" that describe interactions between atoms and nanoparticles in a computationally less demanding, coarse-grained fashion. Gauging the accuracy of the predictions we make using these models is an important part of our research. In some cases, however, the complex quantum nature of interactions simply cannot be forced into the narrow form of simple classical models. In these cases, we apply techniques from machine learning, so called neural networks, to predict complicated quantum interactions with much enhanced computational efficiency.

The computational resources available on CHPC clusters are invaluable for all our scientific endeavors. A single molecular dynamics simulation of metal-organic framework formation, for instance, typically runs simultaneously on tens of processors that communicate with each other and share information about the positions of particles and their mutual interactions. In this scheme, single processors are assigned a small spatial region of the simulated system, perhaps a cubic volume of 10 nanometers across. These kinds of simulations also benefit enormously from high-performance graphics processing units (GPUs) available on CHPC resources. The architecture of these GPUs, originally developed to power graphically demanding video games, consists of many sub-processors and lends itself naturally for the decomposition of large simulated systems into smaller regions. In fact, watching (simulated) nanomaterials assemble on high-powered CHPC resources is almost as fun as building remote control cars.

CHPC Presentation - Scientic Visualization for HPC

Aaron Knoll, research computer scientist at the University's Scientific Computing and Imaging Center (SCI), will give the final lecture, "How to Make Pretty Pictures for Science!," in CHPC's fall presentation series. Dr. Knoll will provide an introduction to scientic visualization, describing how visualization is used in practice, common algorithms (isosurfacing, volume rendering, stream lines), popular software tools (VTK, ParaView, VisIt, SCIRun, Visus), and remote and parallel visualization.

Everyone is invited to the presentation.

Tuesday, December 1st
1 - 3 PM
INSCC Auditorium

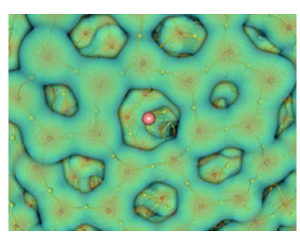

Images: Aaron Knoll and Atilla Gyulassy (SCI)
Data: KC Lau (Argonne National Lab)

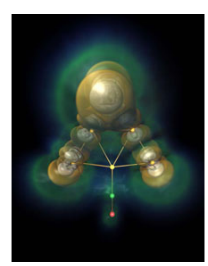

Data: Aslihan Sumer (ARI), Julius Jellinek (ANL)

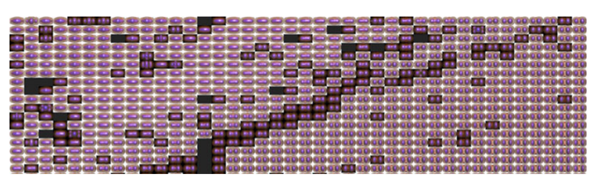

DFT computation - nuclear fission Plutonium (Data: Nicolas Schunck, LLNL)

CHPC Users

Highlight Research at SC16

Each year CHPC has a booth at The International Conference for High Performance Computing, Networking, Storage and Analysis (SC) that highlights some of the research done with CHPC computing resources. This year the conference is in Austin, Texas, but in 2016 it will be held once again in Salt Lake City. The dates are November 14th to the 18th, 2016. Next year's meeting will give many University of Utah researchers a close-to-home opportunity to highlight their work through posters and video presentations at the CHPC booth. Now is the time to start planning.

Contact Sam Liston if you want your research highlighted.

SC is sponsored by ACM (Association for Computing Machinery) and IEEE Computer Society. The conference offers a complete technical education program and exhibition to showcase the many ways high performance computing, networking, storage and analysis lead to advances in scientific discovery, research, education and commerce. This premier international conference includes a globally attended technical program, workshops, tutorials, a world class exhibit area, demonstrations and opportunities for hands-on learning.

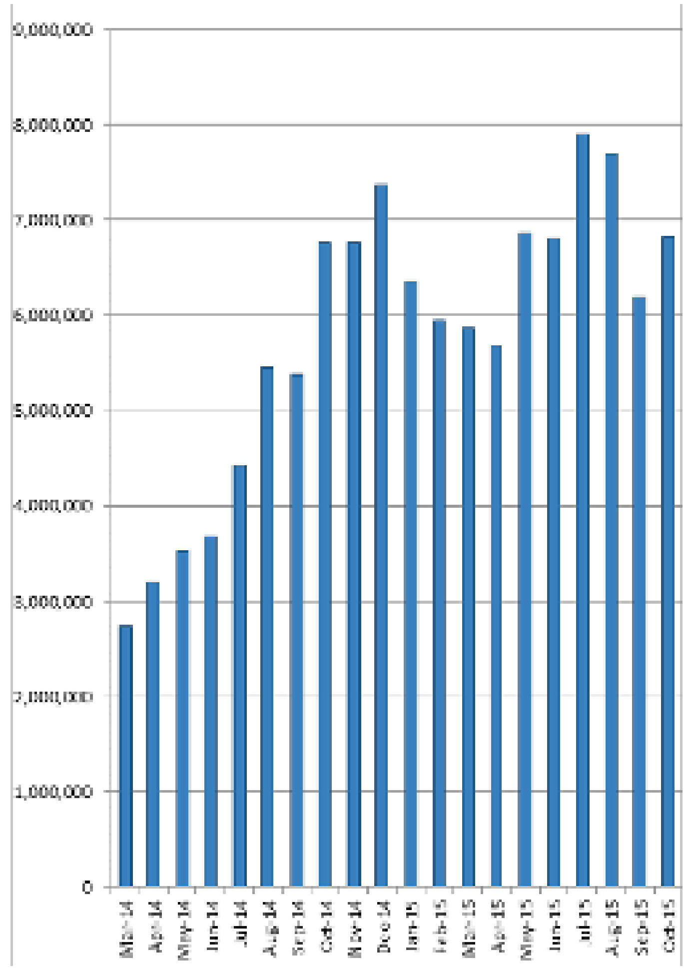

Monthly core hours used on all CHPC clusters showing growth in usage

CHPC's booth at SC 2012 in Salt Lake City

CHPC's Local and National Collaborations

by Joe Breen, Network Initiatiives

Beyond operations, CHPC has been very active in collaborations locally and nationally. In order to help University of Utah scientists and users be able to advance their own fields of study, CHPC is actively exploring new infrastructure technologies and research that will enable the delivery of services anywhere, anytime, with the correct network and security characteristics.

Enabling large data transfers is a fundamental requirement for CHPC's support of its constituents. A NSF CC-NIE collaboration with Univ of Utah domain scientists has allowed CHPC and Univ of Utah School of Computing Flux researchers to explore delivering a Science DMZ performance environment as a "slice" of compute and networking dynamically. This work has enabled CHPC to collaborate locally and nationally with different fields of genomics. CHPC has been working with multiple collaborators in Clemson and the National Institute of Health to characterize aspects of the network and to setup national scale network and compute environments for large plant genomics transfers. This work has fed into additional Clemson and CHPC collaboration exploring the use of Clemson developed software defined network techniques to optimize large transfers.

Large transfers and the ability to provision network resources dynamically requires good troubleshooting tools. CHPC staff have been collaborating with the Internet2 SDN working group and University of Houston faculty in exploring new techniques to trace networks across a dynamic circuit. The emerging SDNTrace software and protocol is a result of this work. Active measurement and characterization of dynamic circuits end-to-end requires new techniques. A collaboration of CHPC, University of Indiana Crest researchers, and Dept of Energy ESnet Advanced Network group has been exploring active measurement in virtual networks across a national footprint. This ongoing work has received acceptance as a formal demo in the upcoming SC15 conference. CHPC has been collaborating with Internet2, Indiana University, Mid-Atlantic Crossroads (Univ of Maryland), Florida International University, the Open Network Lab, and others in exploring alternative network routing techniques across the US.

In parallel to these activities, CHPC and the School of Computing Flux group have collaborated on a NSF Major Research Instrumentation grant named Apt that has explored the ability to dynamically provision a HPC cluster with specific characteristics. Flux used some of the software techniques from Apt to develop the national Cloudlab infrastructure. CHPC is exploring using this new Cloudlab infrastructure to expand the dynamic HPC ideas at a national scale. In order for these national scale HPC ideas to be successful, they require federated credentials and authorization. CHPC and UIT staff have been working with Clemson, Raytheon BBN, University of Illinois and other collaborators on the Fedushare project to accomplish this vision.

New dynamic infrastructure allows emerging areas of domain science, new research across disciplines, and new research across geographic boundaries. These emerging areas of research require additional user support, often across institution. The ACI-Refs collaboration, of which CHPC is a part, leverages user support skills across institutions and builds relationships that allow for better support for emerging research. This support is allowing CHPC exploration in different growth areas such as High Throughput Computing through the Open Science Grid. CHPC staff are investigating how to implement Open Science Grid nodes in tandem with existing HPC nodes and how to engage users in best use of each paradigm.

Continual collaboration, exploration and work at a national and local scale is enabling CHPC to advance quickly in order to support the emerging needs of the University of Utah faculty and their collaborators.

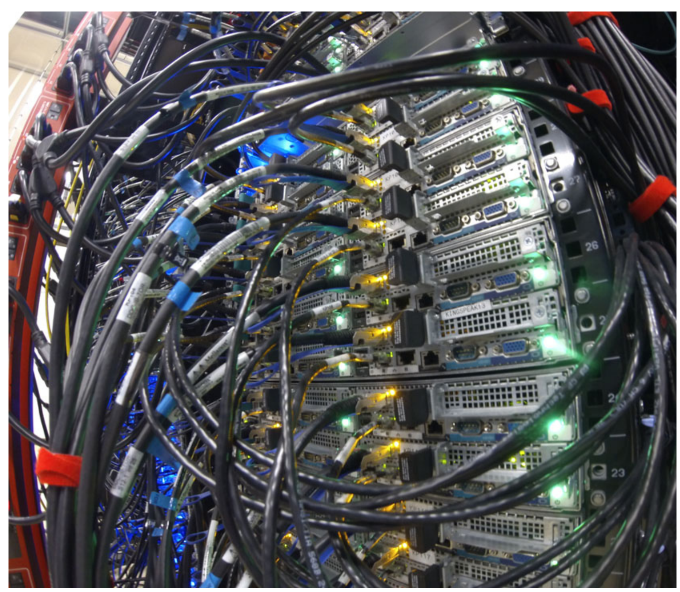

Well connected.
(Kingspeak's cabling) Photo by Sam Liston

News

What's New at CHPC?

by Anita Orendt

- **CHPC Presentations are now broadcast via Lync**

We are now using Lync (Skype for Business) to provide remote access to our training sessions. The information page for each presentation and the email reminders now include the information on how to join the presentation. We still recommend attendance in person if possible, especiallyfor the hands-on presentations. The presentation schedule with the links to each one can be found at https://www.chpc.utah.edu/presentations/

- **Hadoop cluster**

CHPC has a small Hadoop cluster called Elephant. This cluster has two login nodes, two name nodes, and 24 data nodes (with 3.7 TB storage per node), and has a replication factor of 3. It is running the HortonWorks Data Platform (HDP) 2.2.4. Additional information about the cluster can be obtained at http://bit.ly/1JuVtzn. Access to the cluster is by request; please submit a request to issues@ chpc.utah.edu with 'Access to the Hadoop Cluster' in the subject line. If you would like additional information about Hadoop, you can also email us at issues@chpc.utah.edu.

- **Application Database**

CHPC now has a searchable database of all applications that we have installed on CHPC clusters. The link to the database can be found on the CHPC Software Documentation page, http://bit.ly/1POv9Bx. The database includes information on the package including the version, the clusters it will run on, a link to the software web site, the installation path and the best person at CHPC to contact about the usage.

- **Xalt and Modules**

In conjunction with the September downtime, CHPC added a change to the login scripts such that all users are now set up to use modules, even if they still are running the older chpc provided login files. There are two basic modules that are always loaded: chpc and xalt. Xalt is a package that is used to gather data on the applications being run. CHPC will use this data to help us trouble shoot user issues as well as to gather information on the use of the various packages available at CHPC. We request that users check their scripts and remove any 'module purge' lines as this line will unload these modules and keep this data from being collected.

- **Data Center Tours**

This Fall CHPC began offering tours of the Downtown Data Center that houses the CHPC compute and storage resources. The first tour was in early October and due to the high interest level we have scheduled additional tours on November 5 and December 3. Tours last about 90 minutes and are limited to 20 participants. All tour participants must be registered and must bring ID (driver's licenses are preferred). Please see CHPC News at http://bit.ly/1ObGzLf for a link to the registration site as well as for directions to the Data Center. After these initial tours, we will move to a quarterly schedule. We can also provide tours for groups upon request.

- **Video Tutorials**

CHPC is producing a series of short "How To" training videos. Albert Lund, writer and narrator, and Nathaniel Ellingson, video editor, have created tutorials on the following subjects:

- **Introduction to Linux**
- **Connecting to CHPC Resources**
- **Introduction to FastX**
- **Intro to Managing your Environment with Modules**
- **Basics of Slurm**
- **Submitting Batch Jobs with Slurm**
- **Using Slurm for Interactive Work**

We have additional videos planned on topics such as matlab, R, duo, git & svn, the CHPC allocation process, security best practices. We are interested in your feedback on these videos and also on suggestions of additional topics that we should cover. The links to these videos are at https://www.chpc.utah.edu/documentation/videos/

- **Info for Python Users**

Are there python modules that you need to be available on CHPC? Are they modules that are very specialized to your research? In this case, you could install the packages in your own directory. Check out the new software documentation page on Python Virtual Environments & Installation of Packages that can be found at http://bit.ly/1S4RlJf. On this page there are step by step instructions to accomplish this task. If you have any questions on the process or on Python in general, please contact Dr. Wim Cardoen at CHPC.

- **Upcoming Changes to scratch file systems**

In the next few months there will be a number of changes coming to the scratch file system offerings. The first change will be the addition of a new Lustre scratch file system that will be about 750TB in size. Once this new file system is deployed, CHPC will schedule a draining and downtime of the existing /scratch/kingspeak/serial followed by the retirement of the /scratch/ibrix/chpc_gen file system. Watch for announcements of these changes.

Selection of Recent Research Using CHPC Resources

Alshykhly, O. R., Fleming, A.M., Burrows, C. J. (2015). "Guanine Oxidation Product 5-Carboxamido-5-formamido-2-iminohydantoin Induces Mutations When Bypassed by DNA Polymerases and Is a Substrate for Base Excision Repair." Chem. Res. Toxicol. 28: 1861 - 1871.

Gorski, G., Strong, C., Good, S.P., Bares, R., Ehleringer, J.R., Bowen, G.J. (2015). "Vapor hydrogen and oxygen isotopes reflect water of combustion in the urban atmosphere." Proceedings of the National Academy of Sciences 112(11): 3247 – 3252.

Johnson, D. P., et al. (2015). "HDAC1,2 inhibition impairs EZH2- and BBAP-mediated DNA repair to overcome chemoresistance in EZH2 gain-of-function mutant diffuse large B-cell lymphoma." Oncotarget 6(7): 4863 - 4887.

Mougel, V., Santiago, C. B., Zhizhko, P.A., Bess, E.N., Varga, J., Frater, G., Sigman, M.S., Coperet, C. (2015). "Quantitatively Analyzing Metathesis Catalyst Activity and Structural Features in Silica-Supported Tungsten Imido–Alkylidene Complexes." Journal of the American Chemical Society 37(20): 6699–6704.

Nguyen, A. H., Molinero, V. (2015). "Identification of clathrate hydrates, hexagonal ice, cubic ice, and liquid water in simulations: the CHILL+ algorithm." J. Phys. Chem. B 119(8): 4104-4117.

Nguyen, B. T., Samimi, A., Simpson, J.J. (2015). "Recent advances in FDTD modeling of electromagnetic wave propagation in the ionosphere." Applied Computational Electromagnetics Society Journal 29(12): 1003 - 1012.

Shu, Y., Bolton, A.S., et al. (2015). "The Sloan Lens ACS Survey. XII. Extending Strong Lensing to Lower Masses." The Astrophysical Journal 8013: 71 - 88.

Starr, A.M., Moore, J.R., Thorne, M.S. (2015). "Ambient resonance of Mesa Arch, Canyonlands National Park, Utah." Geophy.Res. Lett 42: 6696 – 6702.

Wheeler O.W,. Carl, D.R., Hofstetter T.E., Armentrout, P.B. (2015). "Hydration Enthalpies of $Ba(2+)(H_2O)x$, x = 1-8: A Threshold Collision-Induced Dissociation and Computational Investigation." J Phys Chem A 119(16): 3800 - 3815.

CHPC Welcomes Utah State University

CHPC entered into an agreement with Utah State University to provide high performance computing support to USU research faculty and their students.

Recently four USU research groups purchased compute nodes on CHPC's kingspeak cluster and/or storage space and backup services. In addition we have other USU faculty and students who now have the same accessibility provided freely to University of Utah researchers.

Marc Mansfield, senior environmental modeler at the USU Bingham Research Center in Vernal, Utah, heads the effort to produce reliable computer models of winter ozone formation in the Uintah Basin as part of the effort to ameliorate air quality problems in the Intermountain West.

Robert Spall, Chair of the Mechanical & Aerospace Engineering Department, has been active in developing computational fluid dynamics algorithms for use in engineering education. The nodes his department purchased will be available to all the department's faculty and students.

Zachariah Gompert, Evolutionary Genetics and Genomics, has his lab working on the development of models and software "to measure evolutionary processes from genome-scale space-time allele frequency data, test for evidence of polygenic adaptation and study global and local ancestry in admixed lineages." The animals supplying the data are manakins (a small tropical bird), butterflies and stick insects.

Karen Kapheim, Biology, states her labs's "primary research objectives are to gain insights about the mechanistic underpinnings of variation in behavior, and to use these insights to investigate the evolutionary processes governing the origins, maintenance, and elaboration of social organization." She and her student are using bees in their studies.

What is CHPC?

The Center for High Performance Computing (CHPC) serves as an expert team to broadly support the increasingly diverse research computing needs on campus. These needs include support for big data, big data movement, data analytics, security, virtual machines, Windows science application servers and advanced networking.

CHPC also operates a protected environment (PE) for researchers who work with data that is sensitive in nature. These resources have been reviewed and vetted by the Information Security Office and the Compliance Office as being an appropriate place to work with Protected Health Information (PHI).

These computing resources are available to all faculty at the University of Utah, their students and research staff.

CHPC Staff Directory

Administrative Staff	Title	Phone*	Email	Location
Tom Cheatham	Director	585-6318	tec3@utah.edu	414 INSCC
Julia D. Harrison	Associate Director	585-1869	julia.harrison@utah.edu	430 INSCC
Guy Adams	Assistant Director, Systems & Network	554-0125	guy.adams@utah.edu	424 INSCC
Anita Orendt	Assistant Director, Research Consulting & Faculty Engagement	231-2762	anita.orendt@utah.edu	422 INSCC
Janet Ellingson	Administrative Manager & Newsletter Editor	585-3791	janet.ellingson@utah.edu	405 INSCC

Scientific Staff	Expertise	Phone*	Email	Location
Wim Cardoen	Scientific Applications	971-4184	wim.cardoen@utah.edu	420 INSCC
Martin Cuma	Scientific Applications	652-3836	martin.cuma@utah.edu	418 INSCC
Albert Lund	Graduate Assistant	N/A	albert.lund@utah.edu	405-14 INSCC

Project Leads		Phone*	Email	Location
Wayne Bradford	Security	243-8655	wayne.bradford@utah.edu	426 INSCC
Joe Breen	Advanced Network Initiatives	550-9172	joe.breen@utah.edu	416 INSCC
Steve Harper	Virtualization	541-3514	s.harper@utah.edu	405-30 INSCC
Brian Haymore	HPC & Storage	558-1150	brian.haymore@utah.edu	428 INSCC
Sam Liston	Special Events/SC	232-6932	sam.liston@utah.edu	405-41 INSCC

Technical Support Staff		Phone*	Email	Location
Amanda Allen	Help Desk Assistant	N/A	N/A	405-3 INSCC
Irvin Allen	Systems Admin	231-3194	irvin.allen@utah.edu	405-40 INSCC
Nathaniel Ellingson	Tech Assistant	N/A	N/A	405-28 INSCC
Jake Evans	Network Engineer	718-1526	jake.evans@utah.edu	405-32 INSCC
Paul Fischer	Tech Assistant			
David Heidorn	Systems Admin	303-987-3072	david.heidorn@utah.edu	405-18 INSCC
Alan Navarro	Tech Assistant			
Aaron Pabst	Tech Assistant			
David Richardson	Systems Admin	550-9518	david.richardson@utah.edu	405-29 INSCC
Walter Scott	User Services	309-0763	walter.scott@utah.edu	405-39 INSCC
Alan Wisniewski	Network Support	580-5835	alan.wisniewski@utah.edu	405-38 INSCC

*All phone numbers are preceded by area code 801 unless otherwise noted.

THE UNIVERSITY OF UTAH

Center for High Performance Computing
155 South 1452 East, RM #405
SALT LAKE CITY, UT 84112-0190

Welcome to CHPC News!

If you would like to be added to our mailing list, please fill out this form and return it to:

Janet Ellingson
THE UNIVERSITY OF UTAH
Center For High Performance Computing
155 S 1452 E ROOM 405
SALT LAKE CITY, UT 84112-0190
FAX: (801)585-5366

Name:
Phone:

Department or Affiliation:
Email:

Address:
(UofU campus or U.S. Mail)

Thank you for using our Systems!

Please help us to continue to provide you with access to cutting edge equipment.

ACKNOWLEDGEMENTS

If you use CHPC computer time or staff resources, we request that you acknowledge this in technical reports, publications, and dissertations. Here is an example of what we ask you to include in your acknowledgements:

"*A grant of computer time from the Center for High Performance Computing is gratefully acknowledged.*"

Please submit copies or citations of dissertations, reports, pre-prints, and reprints in which the CHPC is acknowledged to: Center for High Performance Computing, 155 South 1452 East, Rm #405, University of Utah, Salt Lake City, Utah 84112-0190