Center for
High Performance
Computing

NEWS

THE UNIVERSITY OF UTAH

FALL 2010

Research

CHPC Helps Reveal the Secrets of the "Dreaded" Lake Effect

by Jim Steenburgh, Trevor Alcott and Kristen Yeager, Department of Atmospheric Sciences

Lake-effect snow, produced when cold winter winds move across warm lake water, brings joy to skiers, headaches to commuters, and ulcers to meteorologists. Utah meteorologists call it the "dreaded" lake-effect because it remains difficult to predict despite dramatic advances over the past few decades in numerical weather prediction. Funded by the National Science Foundation and supported by CHPC, Professor Jim Steenburgh, Atmospheric Sciences, heads a new research project that aims to make the lake-effect less dreaded and better anticipated. University of Utah graduate students Trevor Alcott and Kristen Yeager, along with faculty and students at Weber State University and Hobart and William Smith Colleges in upstate New York are also contributing.

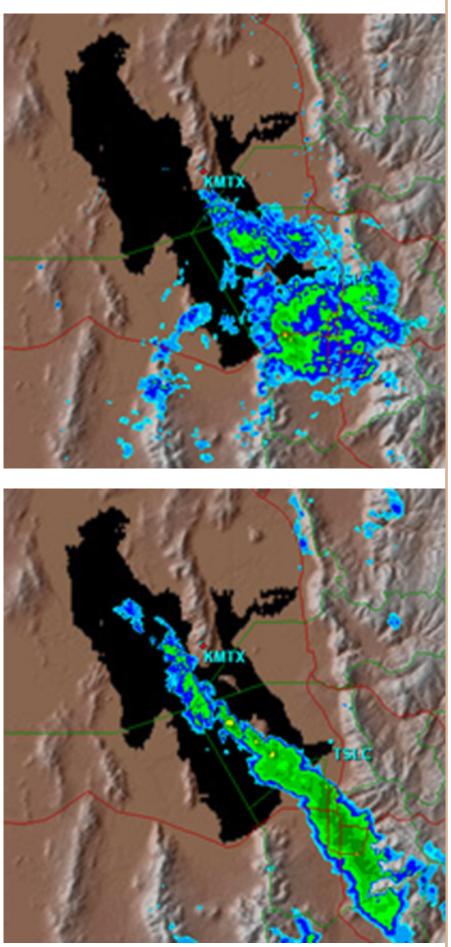

Fig. 1. Radar images of Great Salt Lake lake-effect storms, non-banded and banded modes

In addition to the Great Salt Lake's impact on snowfall, lake-effect (and ocean-effect) snowstorms have been documented around the world, including over and downstream of the Sea of Japan, North Channel, Great Lakes, and even Lake Tahoe, Pyramid Lake, and New York's Finger Lakes. Most research has concentrated on the large bodies of water such as the Great Lakes, which have produced single-event snowfalls of more than 10 feet on the Tug Hill Plateau downstream of Lake Ontario.

The Great Salt Lake, with its widely varying surface area ranging from 2800 to 4700 km^2, is approximately 1/25th the size of Lake Superior. It is quite shallow, with a mean depth of only three meters, and is surrounded by tall mountains and a highly heterogeneous land surface. Forecasting the Great Salt Lake effect is complicated by the lake's small size, which increases sensitivity to small changes in the upstream flow, and the difficulties in specifying the temperature of the lake and the surrounding land-surface conditions, which include the flat desert of the Great Basin and the mountainous terrain of the Wasatch, Oquirrh and Stansbury ranges. Satellite imagery can help specify lake-temperature and land-surface conditions, but only if cloud-free conditions exist prior to a potential storm. The Great Salt Lake lake-effect storms also come in a number of "flavors" or modes. Organized bands produce heavy

(Continued on Page 3)

CHPC Network I/O Performance

The process of data access and network performance is of great concern to users and CHPC staff, who consistently monitor performance in order to troubleshoot and make improvements. From June 22 to July 7, 2010 (14.8 days), CHPC had transferred in aggregate, summed for both incoming and outgoing transmissions, over 280.4 Terabytes of data. That works out to be a daily average of ~18TB/day. During this time no network errors, overruns, frame issues, or even dropped packets were measured on the public interfaces on CHPCFS, even though 326,427,342,845 packets were passed.

On Page 2 are charts of usage patterns for a weekend, a month and a year. As summarized by Joe Breen, CHPC's assistant director, networking: "A tuned file system, tuned nodes and tuned network humming along supporting cool research."

For suggestions on improving your own performance, see the article by Brian Haymore and Sam Liston, "Input/Output in the HPC Environment" on Page 4.

CHPC Network I/O Performance during a weekend, a month and a year.

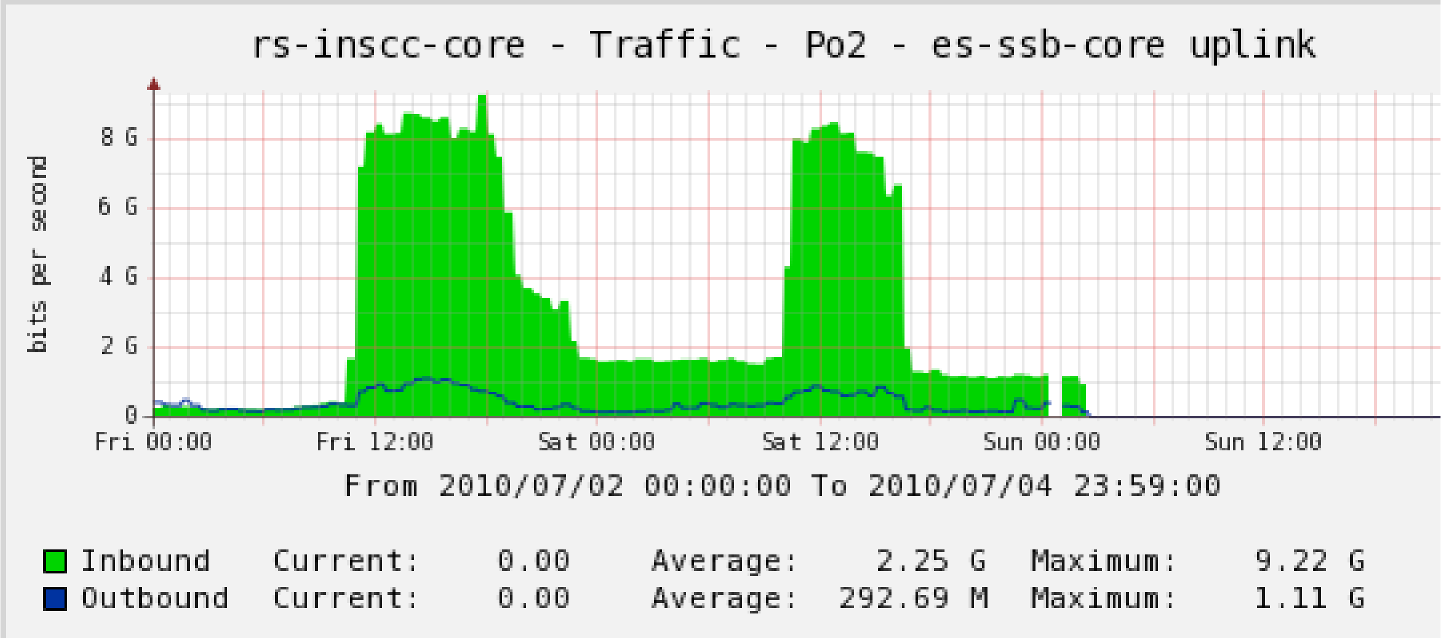

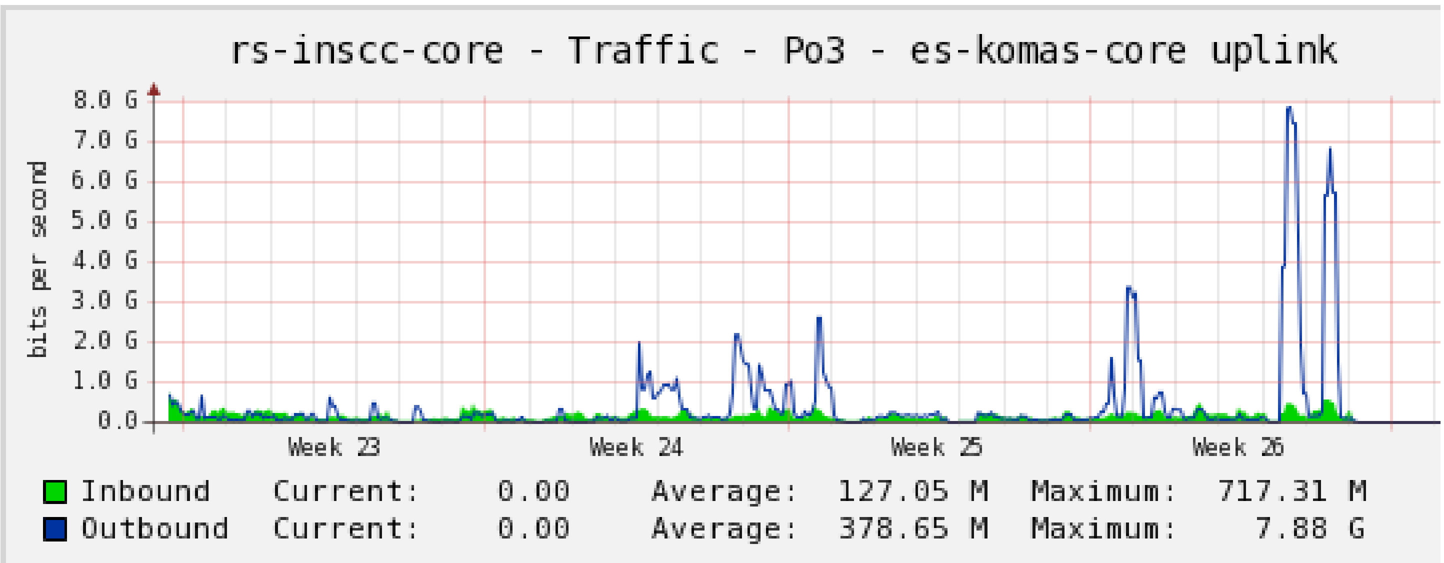

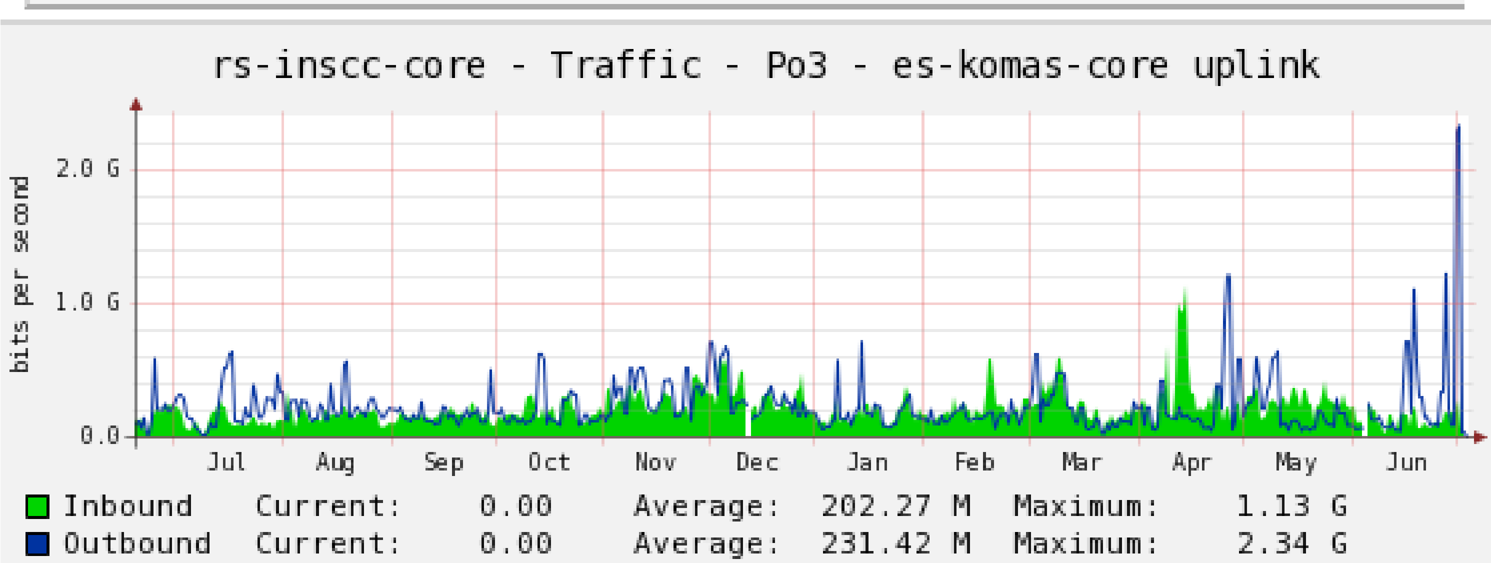

(Continued from Page 1)

snowfall over a narrow area, but non-banded events are widespread and less intense (Fig. 1). The processes responsible for these different lake-effect morphologies have yet to be identified.

The U's research project involves three major components: (1) a radar-based climatology, (2) an observational field program, and (3) numerical modeling. The radar-based climatology has identified 177 lake-effect events since 1997 with data from several sources, including the KMTX radar station located at Promontory Point and North American Regional Reanalysis, which provides gridded atmospheric

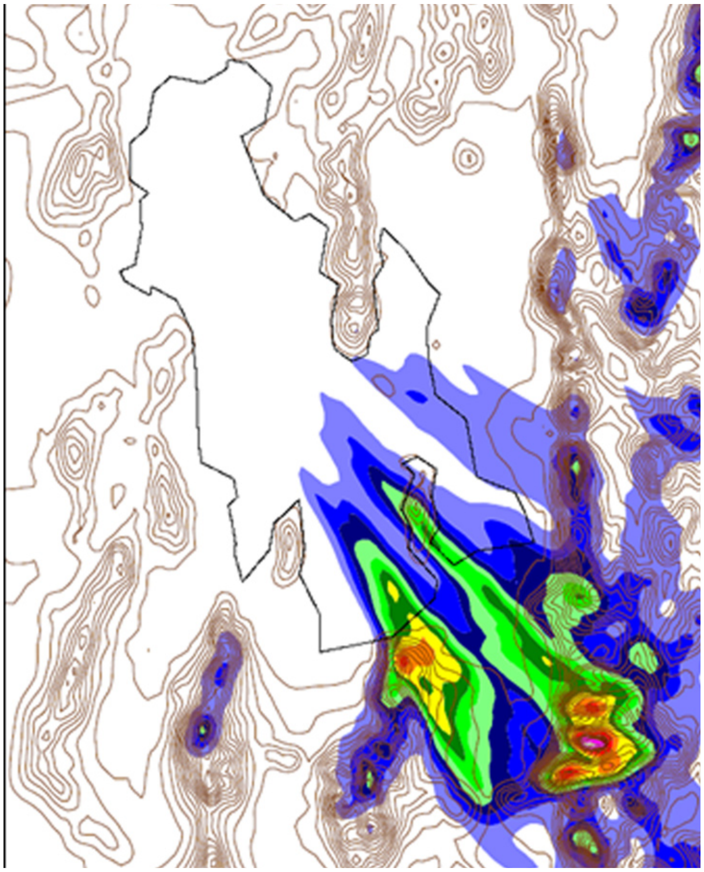

Fig. 2. WRF-Model simulation of a Lake Effect Storm

and land-surface analyses for the western United States. The data are being mined statistically to better understand the environmental conditions that control lake-effect initiation, mode, and intensity. The project is also examining how much the lake effect contributes to snowfall and the hydrologic cycle of the Great Salt Lake basin.

The observational field program is called the Sounding Observations of Lake-Effect Precipitation Experiment (SOLPEX). The objectives of SOLPEX are: (1) to identify the upstream and near-lake environmental conditions that control the initiation, mode, and intensity of lake-effect precipitation, (2) to determine the strength and depth of east-side land breezes and their role in lake-effect events, and (3) to obtain the upper-air data needed to better initialize and validate numerical simulations of lake-effect storms.

SOLPEX will observe several lake-effect storms this winter. Two teams of scientists comprised of University of Utah graduate and undergraduate students will be deployed during potential lake-effect snowstorms, one to the immediate north of the Great Salt Lake near the ghost town of Kelton and the other to the east shore of the Great Salt Lake near the Antelope Island causeway. These two teams will brave snow, wind and cold to launch weather balloons that will provide profiles of wind and temperature during lake-effect storms. The group will study the data to determine lake-effect sensitivity to upstream flow conditions and to improve knowledge of how land breezes and downslope flows from east of the Great Salt Lake affect the initiation and mode of lake-effect storms.

CHPC's computational resources and support play a central role in the field program. Vital to SOLPEX are numerical forecasts produced on CHPC-supported infrastructure using the Advanced Research Weather Research and Forecasting model (WRF-ARW). Current operational weather prediction models used for weather prediction in the United States are run with a 12-km grid spacing, which is insufficient to adequately resolve lake-effect storms. WRF-ARW forecasts will be produced for SOLPEX at 1.3 km grid spacing, which permits the explicit prediction of lake-effect snowbands and improves representation of topographic and lake-surface processes. These forecasts will be used to plan observing efforts, with SOLPEX data ultimately used to validate their accuracy and reliability.

Numerical simulations by the WRF-ARW will also be used to better understand the physics and predictability of lake-effect events (Fig. 2). Of particular interest is the role of topography in event initiation and intensity. Preliminary work shows that flow splitting around the Raft River range of northwest Utah influences lake-effect band development and intensity over the Great Salt Lake. Such effects are not presently incorporated in numerical operational weather prediction.

SC10 New Orleans, LA

Once again CHPC will have a booth at SuperComputing 2010 in November. Designed by Sam Liston, our booth will showcase CHPC supported research. SC12 will be in Salt Lake City at the Salt Palace. This will be a great opportunity to show off our area's computing and networking strengths. If you would like your research highlighted, contact Sam.

ARTICLE

Input/Output in the HPC Environment

by Brian Haymore and Sam Liston

As increasingly fast computational power has allowed researchers to work on larger and more complex problems the size and quantity of the output from those problems has also increased. During the past decade CHPC has seen a great increase in researchers' need for both high performance and high capacity storage. CHPC offers several types of storage to meet these needs.

Users have home directories (/uufs/chpc.utah.edu/common/home/uNID), which are intended for the storage of critical and volatile data. The expectation is that this space is "snappy."

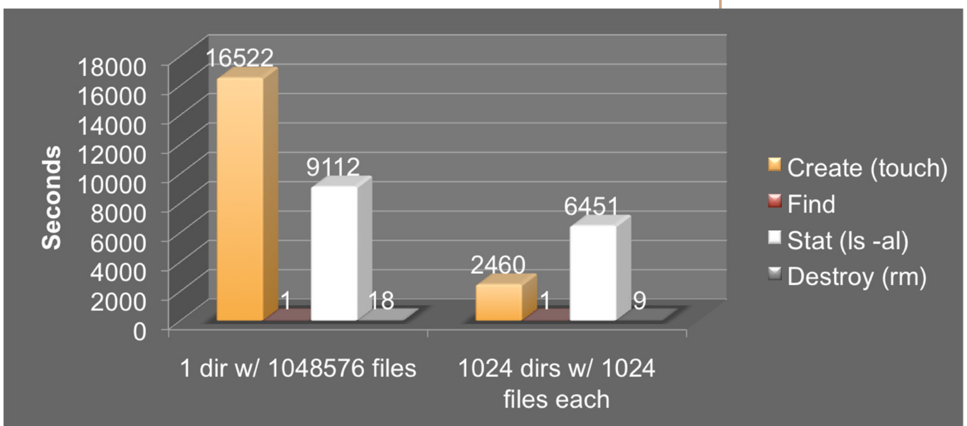

Fig. 1 Hierarchical file system improves performance

Depending on CHPC's arrangements with each department, home directory space may be backed up on a regular basis. Groups may also purchase dedicated space (uufs/chpc.utah.edu/common/home/pi_grp) that can be used for the storage of data from active group projects. Usage and backup policy of this space are determined by the owner group. CHPC also provides network mounted scratch file systems (/scratch/serial) to be used by all users as temporary storage space for their computational output of jobs on the clusters. There is no expectation of data retention on these file systems, therefore there is no back up. The final storage option is that of the local disks (scratch/local) unique to each machine.

The home directories, group data space, and the network mounted scratch are all shared and finite resources. Users should be good citizens and consider the impact of their usage on others. One way of doing this is by determining the most appropriate storage space for their files. The first step is to classify the data by considering these questions: How important is the data? Can it be easily recreated? Is the dataset currently in use? Will the data be used by others? Does the data need to be backed up? Important data should never be stored on the network mounted scratch file system.

After classifying their data and the appropriate storage site, users should consider the tools to use for data migration. SSH/SFTP is simple and portable, rsync is restartable and allows for file verification, tar via SSH is the most efficient to use when transferring many small files.

Consideration should be taken not only when determining where data is stored, but how is it stored. Directory structure and file count should be considered. Too many files in the same directory will result in poor performance. This can be remedied by organizing directories of file in a tree structure. Fig. 1 illustrates the performance improvement when organizing 1,048,576 files into a tree structure of 1024 directories, each containing 1024 files, verses a single directory containing all the files. Changing code to write files in a hierarchical directory structure produced a multiple times speed up.

Processing many small files over the network can produce slower performance as well. Fig. 2 illustrates that copying input files from the network file space to the local disk, processing the files, then copying the output back to the network space can be faster than processing the file directly on the network file space.

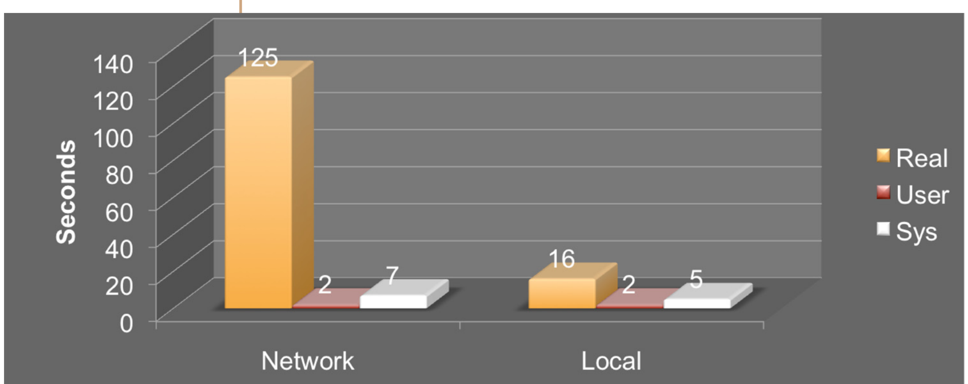

Fig. 2 Processing on local disk improves performance

Users who persistently experience poor performance from their network file systems should send a description of the problem to CHPC by emailing us at issues@chpc.utah.edu.

News

Coming Soon -- EMBER, a New HPC Cluster

by Julia Harrison, Associate Director

We are in the process of finalizing the configuration of our new cluster to be named “ember.chpc.utah.edu”. This cluster will have 636 cores for general campus usage, and approximately 62 Tb of /scratch space. The general campus portion of the new system is expected to be roughly three times the compute capacity of the retiring systems, and will be interconnected with a 40 Gigabit infiniband network. In addition, several research groups have purchased nodes, bringing the total core count for the new cluster up to 3144. Phil Smith’s ICSE group purchased the largest share and, therefore, will have 2244 cores available for their research. The total core count works out to be about 35 Tflops.

To make room for the new cluster (space, heating and cooling constraints) CHPC’s oldest clusters, delicatearch, tunnelarch and landscapearch, were retired on October 4th, 2010. Also note that the old /scratch files systems /scratch/ta, /scratch/da, /scratch/serial-old and /scratch/serial-pio were also retired. These old systems totaled 2.2 Tflops, where the general campus portion of the new system is estimated to be roughly 7.1 Tflops.

After a competative bid process, we chose Hewlett-Packard as our provider and will complete the purchase soon. We are hoping to have the system delivered and deployed for testing sometime in late November or early December 2010, and available to users by the new year. For users preparing their allocation proposals for Winter 2011, please include ember on your request (Due date: November 24, 2010)

CHPC welcomes to our staff four student employees: Brandon Day, Ricky Merida, Josh Spolsdoff, and Kyle McGuire (not pictured).

For CHPC Users:

On-Line Account Application Process

The CHPC account application process is now available online. Although we still accept the old paper form, we encourage everyone to use the online system as it is efficient and easy to use, thanks to the great work of our web developer, Walter Scott.

The process may be initiated either by the applicant or by the faculty advisor (PI) with whom the applicant will be working. Once the online form has been completed and submitted to CHPC, the system will automatically request verification of the applicant’s email address by sending an email to the applicant and waiting for a response. An applicant’s request will not be processed further without a response to this email verification request. If the applicant initiates the process, a request for approval of the application will be sent via e-mail to the faculty advisor. This approval must also be completed before the account will be created.

CHPC policy requires that the faculty advisors of non-faculty CHPC account holders, including post-docs, also have CHPC accounts. If the applicant identifies a faculty member who does not yet have an account, the system will create an account after CHPC has verified the faculty member’s status and email. We encourage faculty members who wish to create accounts for themselves and their research group members to speak with CHPC’s associate director, Julia Harrison, so she can identify the resource needs. She can be reached via email at julia.harrison@utah.edu or by phone 801-652-0019 The online application is on the CHPC website: https://www.chpc.utah.edu/apps/profile/account_request.php

CHPC Fall Presentations

10/7 Intro to Programming with Open MP -- Martin Cuma
10/21 Chemistry Packages at CHPC -- Anita Orendt
10/28 Mathematical Libraries at CHPC -- Martin Cuma
11/4 Statistical Resources at CHPC - Byron Davis
11/11 Using Python for Scientific Computing - Wim Cardoen
11/18 Using Gaussian09 and Gaussview - Anita Orendt
12/2 High Performance Networks and Long-Distance Data Transfers – Tom Ammon
12/9 Debugging with Totalview – Martin Cuma
12/16 Hybrid MPI – Open MP Programming – Martin Cuma

The presentations are in the INSCC Auditorium on Thursdays at 1:00 PM. The INSCC building is just north of the Park Administration Building on Presidents' Circle. For more information about the presentations please go to http://www.chpc.utah.edu/docs/presentations/. Everyone is welcome.

FYI

Published Research Using CHPC Resources

Examples of recently published research that used CHPC resources:

West, G. L., Steenburgh., W.J. (2010). "Life cycle and mesoscale frontal structure of an Intermountain cyclone." Monthly Weather Review 138 (7): 2528-2545.

Pendley, S. S., Yu, Y.B., Cheatham, III, T. E. (2009). "Molecular Dynamics Guided Study of Salt Bridge Length Dependence in Both Fluorinated and Non-Fluorinated Parallel Dimeric Coiled-Coil Proteins." Proteins 74: 612 - 629.

Kochanski, A., M.A. Jenkins, and S. Krueger, 2010: Wind Forecasting in the Fire Environment. Wildland Fire Workshop held in conjunction with 24th Annual Advance Combustion Engineering Research Center Conference, Brigham Young University, Utah.

Shingleton, Nick. (2010) "Coupling a Land-Surface Model to Large-Eddy Simulation to Study the Nocturnal Boundary Layer." MS Thesis. Department of Mechanical Engineering, University of Utah.

Lupton, E. M., Chen, L., Liu, F. (2010). "Uniaxial Strain in Molecular Nanowires: A Case Study of B-phase Polyfluorenes." Chemical Physics Letters 1: 1326.

Halling, M. D., Orendt, A. M., Strohmeier, M., Solum, M.S., Tsefrikas, V. M., Hirao, T., Scott, L.T., Pugmire, R. J., Grant, D. M. (2010). "Solid-state 13C NMR investigations of 4,7-dihydro-1H-tricyclopenta-[def,jkl, pqr]triphenylene (sumanene) and indeno [1,2,3-cd] fluoranthene: Buckminsterfullerene moieties." Phys. Chem. Chem. Phys. 12:7934 - 7941.

What is CHPC?

The Center for High Performance Computing provides large-scale computing resources to University faculty and research staff to facilitate their research. CHPC is located in the INSCC building (just north of the Park administration building) and is responsible for the operation, maintenance and upgrade of their computing resources housed in INSCC, SSB and Komas.

The projects currently supported by CHPC come from a wide array of University disciplines that require large capacity computing resources, both for calculating the solutions of large-scale, two and three dimensional problems and for graphic visualization of the results.

If CHPC resources would be of use in your research, please go to our website www.chpc.utah.edu for more information.

The Updraft cluster. Photo by Sam Liston

CHPC Staff Directory

Administrative Staff	Title	Phone	Email	Location
Julio Facelli	Director	585-3791	julio.facelli@utah.edu	410 INSCC
Julia D. Harrison	Associate Director	585-3791	julia.harrison@utah.edu	430 INSCC
Guy Adams	Assistant Director, Systems	554-0125	guy.adams@utah.edu	424 INSCC
Joe Breen	Assistant Director, Networking	550-9172	joe.breen@utah.edu	426 INSCC
DeeAnn Raynor	Administrative Officer	581-5253	dee.raynor@utah.edu	412 INSCC
Janet Ellingson	Admin. Program Coordinator & 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	587-7770	martin.cuma@utah.edu	418 INSCC
Byron L. Davis	Statistics	585-5604	byron.davis@utah.edu	416 INSCC
Julio Facelli	Molecular Sciences	585-3791	julio.facelli@utah.edu	410 INSCC
Sean Igo	Natural Language Processing	N/A	sean.igo@utah.edu	405-16 INCSS
Anita Orendt	Molecular Sciences	231-2762	anita.orendt@utah.edu	422 INSCC
Ron Price	Software Engineer & Grid Architect	560-2305	ron.price@utah.edu	405-4 INSCC
Derek Huth	Scientific Support staff	N/A	N/A	405-8 INSCC

Technical Support Staff	Group	Phone	Email	Location
Irvin Allen	Systems	231-3194	irvin.allen@utah.edu	405-40 INSCC
Tom Ammon	Network	674-9273	tom.ammon@utah.edu	405-22 INSCC
Robert Bolton	Systems	528-8233	robert.bolton@utah.edu	405 -24 INSCC
Wayne Bradford	Systems	243-8655	wayne.bradford@utah.edu	405-41 INSCC
Erik Brown	Systems	824-4996	erik.brown@utah.edu	405-29 INSCC
Brandon Day	Systems	NA	NA	405 - 20 INSCC
Steve Harper	Systems	541-3514	s.harper@utah.edu	405-31 INSCC
Brian Haymore	Systems	558-1150	brian.haymore@utah.edu	428 INSCC
Eric Hughes	Systems	879-8449	eric.hughes@utah.edu	405-18 INSCC
Samuel T. Liston	Systems, Multimedia	232-6932	sam.liston@utah.edu	405-39 INSCC
Kyle McGuire	User Services	NA	NA	405 - 14 INSCC
Randy Madsen	Systems	N/A	randy.madsen @utah.edu	5700 HSEB
Ricky Merida	User Services	NA	NA	405 - 17 INSCC
Jimmy Miklavcic	Multimedia	585-9335	jimmy.miklavcic@utah.edu	296 INSCC
Beth Miklavcic	Multimedia	585-1067	beth.miklavcic@utah.edu	111 INSCC
David Richardson	Network	550-3788	david.richardson@utah.edu	405-38 INSCC
Walter Scott	User Services	309-0763	walter.scott@utah.edu	405-13 INSCC
Steve Smith	Systems	581-7552	steve.smith@utah.edu	405-25 INSCC
Josh Spolsdoff	Systems	NA	NA	405 - 19 INSCC
Neal Todd	Systems	201-1761	neal.todd@utah.edu	405-30 INSCC
Alan Wisniewski	Network	580-5835	alan.wisniewski@utah.edu	405-21 INSCC

The University of Utah seeks to provide equal access to its programs, services, and activities to people with disabilities. Reasonable prior notice is needed to arrange accommodations.

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

(room 405 of the INSCC Building)

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, preprints, 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