Center for High Performance Computing NEWS

THE UNIVERSITY OF UTAH FALL 2011

Research

Massively Parallel 3D Inversion of Gravity and Gravity Gradiometry Data

by Martin Čuma, Glenn Wilson, and Michael S. Zhdanov

When searching for hydrocarbon and mineral resources, among other things, below the earth's surface, geologists measure the gravity response of the subsurface, thereby giving clues to the hidden properties. Typical surveys with airborne gravity measurements cover hundreds to thousands square kilometers with measurement locations every few meters, resulting in large quantities of data that are finite, noisy and, therefore, in need of efficient processing. Interpretation of gravity surveys remains a challenge despite over a century of research and application. Using the ember cluster computer at CHPC and data from the gravity gradiometry survey of the Vredefort crater in South Africa, we demonstrated how such data can be inverted to 3D density models of unprecedented scale (Figure 1) within hours of data collection. This methodology will enable geologists to better assess the resources beneath Earth's surface.

Structural interpretations of gravity and gravity gradiometry data are often based on some form of Euler deconvolution, wavelet analysis, or analytic signal method. While such methods may provide information about the density of the sources, it is not immediately obvious how this information can be quantified in terms of the density distribution within a 3D earth model. For this reason, inversion of the gravity data to a 3D density distribution is an important step in quantitative interpretation. Generalized inversion methods discretize the 3D earth models into cells of constant density. The inversion process minimizes the difference between the observed gravity and gravity gradient data and the response of the predicted 3D model density. Regardless of the inversion methodology used, all geological constraints manifest themselves as regularization that can be quantified through a choice of data weights, model upper and lower bounds, model weights, an a priori model, and the type of stabilizing functional used.

For gravity, computational complexity increases linearly with the size of the problem. There are two major obstacles in large-scale 3D inversion. The first one being that storing the kernels of the forward modeling operators requires a large amount of computer memory. Even a small-sized 3D inversion of thousands of data points to 3D earth models with hundreds of thousands of cells can exceed memory available on a desktop computer. The second obstacle is the amount of CPU time required to apply the dense matrix of the forward modeling operator to the data and model vectors. Several strategies have been developed that reduce memory requirements and CPU time, including data compression and modeling based on FFT convolution product. However, these methods have significant limitations that render them less useful for routine data processing. Therefore, our approach is to use massively parallel 3D inversion. Our inversion methodology uses the re-weighted regularized conjugate gradient method for minimizing the objective functional. Additionally, we have incorporated a wide variety of regularization options in our 3D inversion algorithm. For a detailed description of our methodology see: http://www.publish.csiro.au/nid/228/paper/PVv2011n152p29.htm (Cont. page 2)

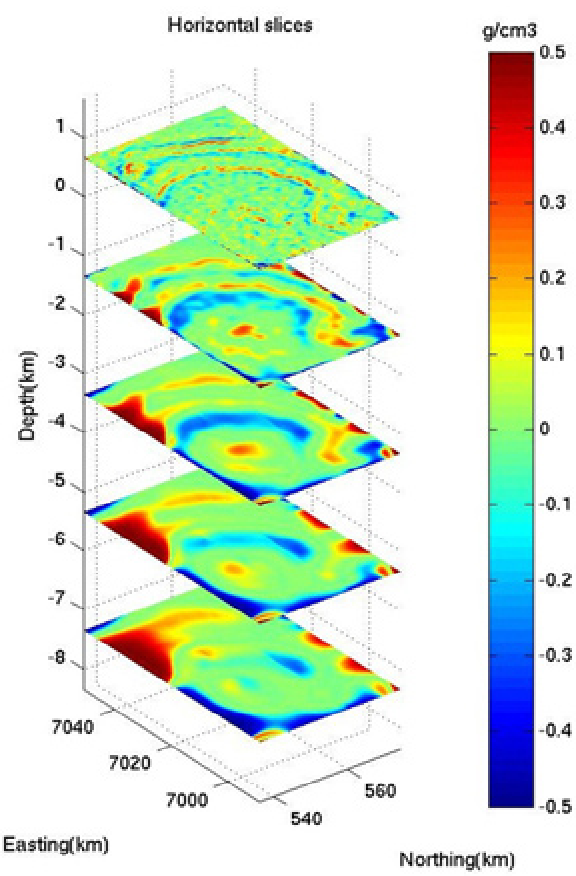

Figure 1. Comparison of horizontal cross sections through the 3D relative density model obtained from joint inversion of all FALCON® gravity gradiometry components from the survey of the Vredefort crater in South Africa. The relative density shown is on top of background density of 2.67 g/cm3

Our 3D inversion algorithm has been implemented as a multilevel parallel application. The 3D inversion domain is divided in a distributed fashion over Message Passing Interface (MPI). On a fine-grained level, loops over the data points and a few other auxiliary loops within each MPI process are further parallelized with a shared memory OpenMP standard. This two-level approach has multiple advantages. It reduces the number of MPI communicating processes, minimizing communication stress on the network. It also saves memory, since data structures need to be replicated by each process and most of the data is shared by the OpenMP threads. Finally, it allows for better locality of the processes/threads on the node's boards and sockets, which improves data transfers to/from the main memory. The data locality is critical on modern non-uniform memory architecture (NUMA) computers as the number of with CPU cores grows.

In a typical cluster configuration, we can run several MPI processes per cluster node. Each of these processes launches a number of OpenMP threads - one thread per processor core. The current generation of clusters have two hexa-core CPUs (i.e., 12 cores) per node. We have found that it is optimal to run one MPI process per socket (i.e., two per node), with six OpenMP threads per MPI process. The advantage of this is the ability to pin the process to the CPU socket, so that it does not move from one socket to another, which improves the memory performance. We have found that without pinning, the performance can degrade by up to 20%.

Our 3D inversion is relatively light in MPI communication, thanks largely to the linearity of the forward modeling operators. Most MPI communication consists of accumulation of the sensitivities and the regularization as reduction operations. As a result, the program exhibits excellent parallel scaling. Parallel scaling is usually evaluated with two different metrics. The first one is called strong scaling. It measures the performance of a fixed problem size with an increasing number of processors. The second metric is weak scaling. It relates the time to complete one unit of work on one processing element to the time to perform N units of work on N processing elements. In both cases, ideal (linear) scaling is 100%. Any scaling below 100% is sublinear, and any scaling above 100% is superlinear. As a side note, it is possible to achieve superlinear scaling due to hardware architectural features that multiprocessor programs can exploit.

Our results were run on CHPC's ember cluster that has 260 nodes, each equipped with two hexa-core (i.e., 12) Intel Xeon CPUs running at 2.8 GHz with 24 GB of RAM and QDR InfiniBand interconnect. Figure 2 shows the parallel scaling efficiency of the Vredefort case study. In the case of strong scaling, as depicted by the blue line in Figure 2, we chose a 3D model with about 11 million cells and 600,000 data. The scaling efficiency is excellent from 18 to 288 cores. We see a drop at 576 cores. This is due to running 12 rather than 6 cores per process, i.e. one rather than two processes per node. The memory load is much more uneven for the single MPI process sharing threads on both CPU sockets in the node, which decreases the efficiency by 15%. The weak scaling, depicted by the red line in Figure 2, varied the number of inversion cells from about 11 million cells on 18 cores to about 350 million cells on 576 cores. Again, the scaling is nearly linear with a 1% to 2 % difference, which can be attributed to system noise. We draw two conclusions from our scaling analysis. First, our 3D inversion software shows linear scaling and is expected to scale well to thousands of cores. Second, we have identified that process and thread locality is critical in achieving optimal performance, and that one MPI processes should be bound to each socket.

The effectiveness of our approach has been demonstrated with a case study for 3D inversion of data from a FALCON® airborne gravity gradiometry survey of the Vredefort crater in South Africa, which included the joint inversion of over 600,000 gravity gradient data to a 3D earth model with over 350 million cells. The computational time for the inversion totaled about 24 hours using a cluster with 576 CPUs.

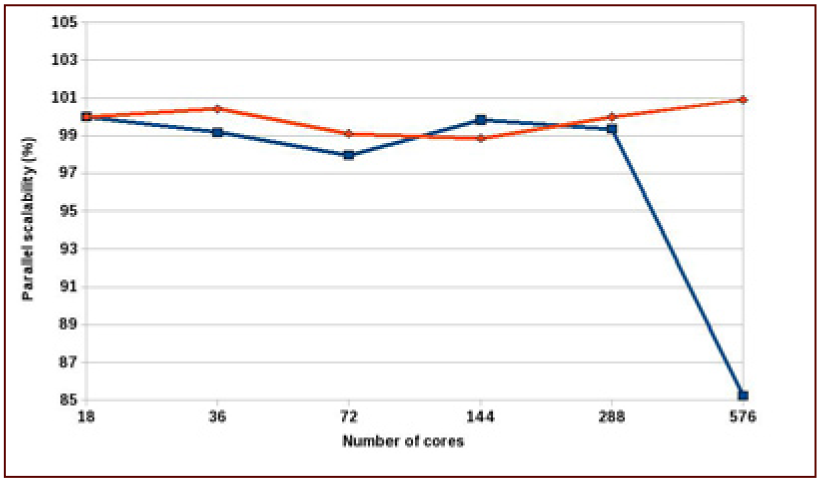

Figure 2. Parallel scaling efficiency for 3D inversion of the Vredefort FALCON® data. Strong scaling is shown in blue, and weak scaling is shown in red.

The Vredefort crater is known as the largest and oldest impact structure on Earth, with a diameter of 250 to 300 km and age about 2 billion years. The impact structure has since been deformed via erosion and tectonic processes, though the center remains largely unaltered. We can clearly see multiple rings with density highs and lows, corresponding to the denser basement that was uplifted by the asteroid impact and lighter sedimental sequences which have been deposited subsequently. The impact site shows a pronounced higher density anomaly located below the surface, the source of which is a subject of ongoing discussion. These results agree well with the known geology and independent analyses of the same data.

NEWS

CHPC Moves to New Storage Architecture

by Guy Adams, Assistant Director, CHPC

With the recent deployment of ember, our latest cluster, we implemented a parallel file system architecture from HP called X9000 for the cluster scratch space. The X9000 line is a scale out network attached storage (NAS) product built upon the technology HP purchased from IBRIX. After considering the feedback from the performance and features of this new storage, we decided to partner with HP to use this technology for our users' home and group directory storage.

HP's X9320 IBRIX network attached storage

Our storage requirements have escalated from tens of Terabytes to hundreds of Terabytes. Our first purchase of this storage will support eleven departments, deploying around 330 usable Terabytes. Within the X9000 line, there are several disk options based upon acceptable performance requirements. Since our home and group storage needs are based on capacity first and performance second, we are moving forward with X9320 couplets using two Terabyte disks.

By migrating to this technology CHPC can offer a more feature-rich storage solution. For example, researchers were previously limited to a single-name space of 16 Terabytes. When their needs moved beyond 16 Terabytes, they were given another file system and they were responsible for moving files and keeping track of their usage. With the X9320s we have the ability to grow a single-name space into the hundreds of Petabytes. Another feature is the ability to provide snapshots. This feature is configurable per research group, and in some cases even finer grain. The snapshot feature will enable researchers to retrieve an earlier version of a file from a snapshot directory. Based on their needs, this can be configured to be snapshot hourly, daily, a few days or months.

The X9320 can be populated with faster disk media than the 2 Terabyte disks if required by the research group. This allows us to tier the files between different types of disk speeds without impacting the research group. For example, if a research group needs the ability to write files as fast as possible, but after a couple of weeks the group doesn't need high speed, the system will tier it to lesser performing disks. A roadmap consideration of this product is to tier to 'tape.' CHPC has requested this feature in future releases as this would allow our storage to be a true hierarchical storage management system.

After we move to the new downtown datacenter next year, this system will support both 10 gigabit Ethernet and the quad data rate 40 gigabit Infiniband. This will be the first home and group storage deployment that will use the two fabrics shared by our supercomputer clusters and department desktops. During that phase of the deployment we hope to offer better direct access to our users' home and group storage from the supercomputer clusters.

A Sample of Research Using CHPC Resources

DeMille, T., Cheatham, T.E., Molinero, V. (2011). "A Coarse-Grained Model of DNA with Explicit Solvation by Water and Ions." J. Phys. Chem. B. 115(1): 132 - 142.

Facelli, J. C. (2011). "Chemical shift tensors: Theory and application to molecular structural problems." Progress in Nuclear Magnetic Resonance Spectroscopy 58: 176 - 201.

Huang, M., F. Cavallo, Feng Liu and M. G. Lagally (2011). "Nanomechanical Architecture of Semiconductor Nanomembranes." Nanoscale 3(96).

Wang, Z., Y. Zhang, and Feng Liu (2011). "Formation of hydrogenated graphene nanoripples by strain engineering and directed surface self-assembly." Phys Rev. B 83: 041403 (R).

Alcott, T. I., Steenburgh, W. J. (2010). "Snow-to-liquid ratio variability and prediction at a high elevation site in Utah's Wasatch Mountains." Wea. Forecasting 25(1): 323 - 337.

Crosman, E. T. (2011). Idealized large-eddy simulation sensitivity studies of sea and lake breezes. Atmospheric Sciences. Salt Lake City, UT, University of Utah. Ph.D.

Thomas, K. C., Ethirajan, M., Shahrokh, K., Sun, H., Lee, J., Cheatham, T.E., Yost, G.S., Reilly, C.A. (2011). "Structure activity relationship of capsaicin analogues and TRPV1-mediated human lung epithelial cell toxicity." J. Pharm. Exp. Ther 337: 400-410.

For CHPC Users:

CHPC Web Site Profile Application

by Walter Scott

We have recently added many new features to the "profile" application on our web site (http://www.chpc.utah.edu/profile). The new features are aimed at providing users with more control over and better access to their accounts.

Past improvements have given users access to update their email address on our systems and to change the default linux shell. We also recently added functionality for users to view their disk usage on our filesystems and to see if they have exceeded any of their filesystem quotas.

Principle Investigators (PIs) can now do the following:

- Request an account for one of their project members.
- View pending account requests for their users and update the user's email address or send them to another PI.
- View past account requests and a list of current users in their project.
- View quota information for users in their proect.

Below is screen shot of the new feature as seen by the PIs. Users not identified as PIs will also be able to see the quota information on their profile page.

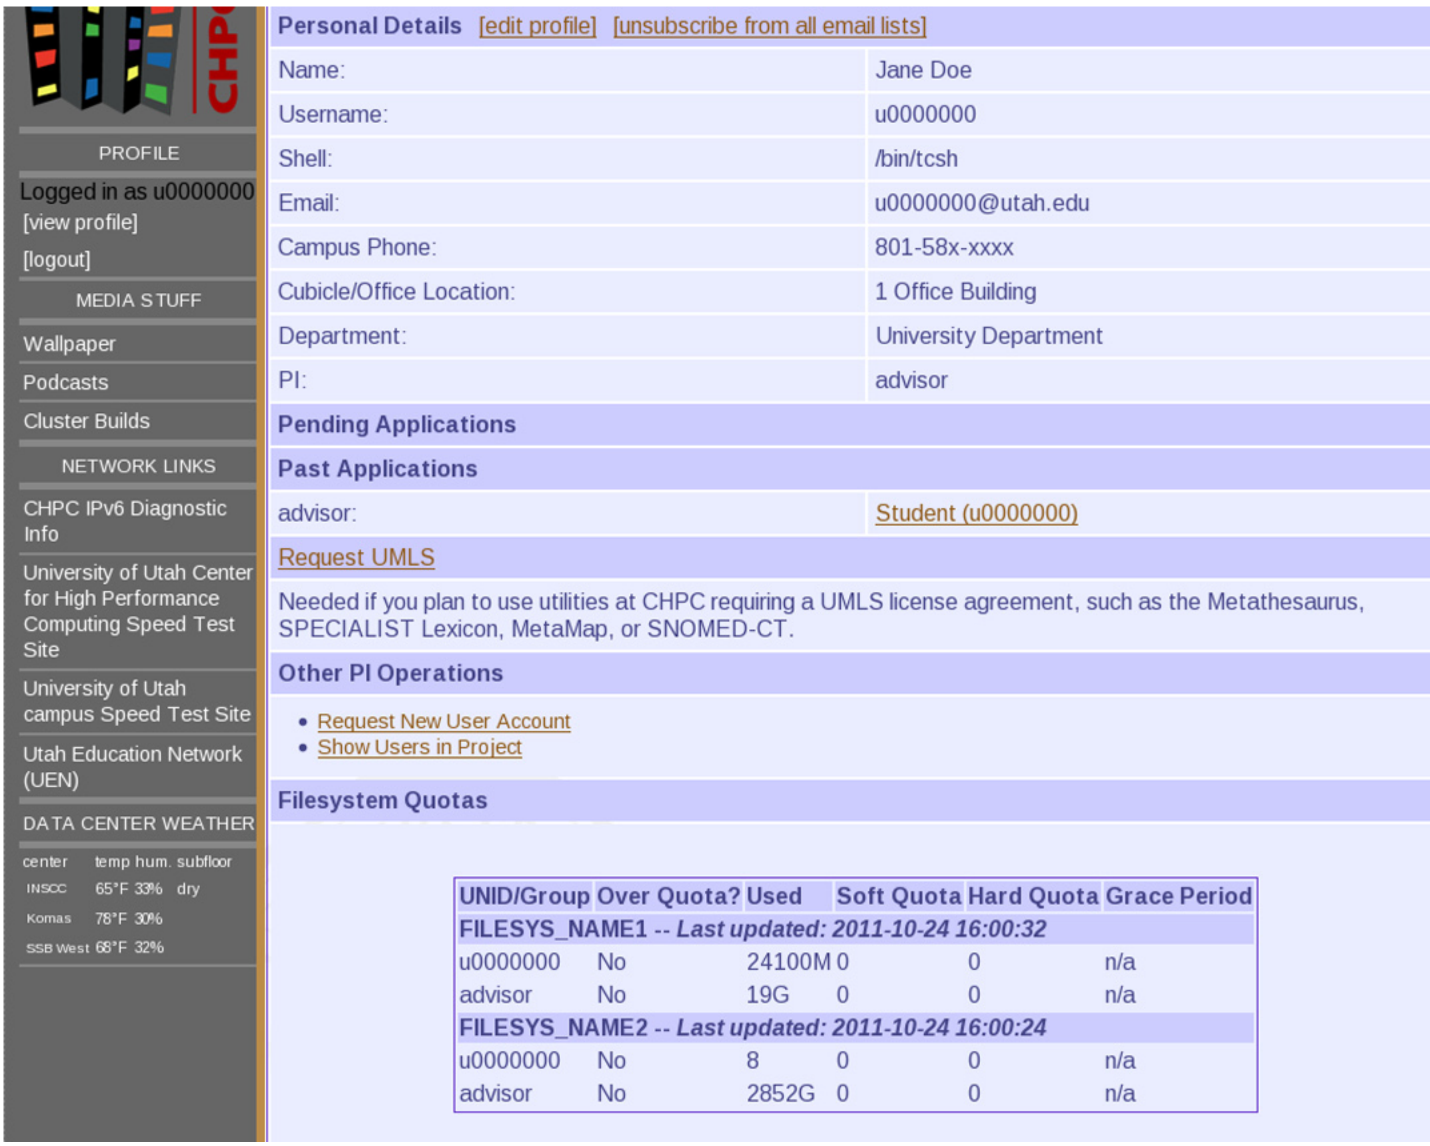

UNID/Group	Over Quota?	Used	Soft Quota	Hard Quota	Grace Period
FILESYS_NAME1 -- *Last updated: 2011-10-24 16:00:32*					
u0000000	No	24100M	0	0	n/a
advisor	No	19G	0	0	n/a
FILESYS_NAME2 -- *Last updated: 2011-10-24 16:00:24*					
u0000000	No	8	0	0	n/a
advisor	No	2852G	0	0	n/a

Research

CHPC Helps Geographer Identify Social Interaction Potential

by Janet Ellingson

For many of us the subject of geography pertains to the maps we drew and the state capitols we memorized in elementary school. Of course, the field is far more than borders and place names. And, as with most academic fields of study, during the 1970s computers expanded the ability of geographers to gather, analyze, and share geographical data, including data on urban land use and the movement of humans through a defined space. University of Utah geographer, Prof. Steven Farber, is now using the computational resources at CHPC to study the movement of people through a metropolitan area and to analyze the potential for social interaction.

The amount of interaction we have with others in our geographic region and the quality of that interaction contributes to community cohesiveness, economic output, and quality of life. Therefore, Prof. Farber has focused his research on identifying the characteristics of a region that contribute to the frequency and quality of social interactions. The practical application of his research will be in the planning and construction of urban environments that maximize opportunities for interaction and thereby foster greater social benefits. In addition, his research may identify existing barriers and constraints that discourage community cohesiveness. Such impediments can then be ameliorated.

People will meet with each other if there is a will to do so. Prof. Farber acknowledges this when he writes, "We recognize that if two people wish to have a joint activity, they can and will make it happen by arranging their daily, weekly, or monthly schedules accordingly." However, some communities seem more conducive to social interactions than others. Prof. Farber hopes that his research "can be used to evaluate the degree to which urban form, transportation networks, and commuter flows assist or inhibit joint activity participation in a metropolitan region provided there exists intent amongst people in a social network to do so."

The primary step in the research was to identify useful metrics. With computer modeling, Farber and his colleagues created 80 synthetic cities each populated with one million workers living and working within a 7x7 tessellation of 12 km (about 7.5 miles) square spaces. The research group gave each modeled city unique land-use (residential and work sites) and commute-flow patterns representing various levels of concentration, sprawl and willingness to travel. They then identified all the possible scenarios of movement from work to home through the space during a specific time and determined the potential for a person to meet with another person after work. The next challenge was to move to the analysis of real communities to demonstrate the usefulness of the metric established in these models. Prof. Farber and his colleagues studied the communities in Davis, Salt Lake, Utah and Weber counties using data from the US Census Bureau Transportation Planning Products.

The challenge of determining a metropolitan area's social interaction potential becomes more complex when the number of variable increases. So far, Prof. Farber's study has restricted the variables to home and work site location and commuting paths. To make the study more useful, many more variables will be introduced, including varying work hours, speed of travel, number of persons meeting together (this study only considered two people meeting), and distribution of social activity locations. Further studies could also include socio-economic data.

Fascinating questions arise from Prof. Farber's study. One of his future research objectives is to consider the consequences of high and low social interaction potential. For example, is there a connection between economic creativity and social interaction? Is Silicon Valley especially conducive to innovation because it has a high social interaction potential? Are economic inequalities related to inequalities in social interaction potential? Are citizens more likely to vote in areas with high social interactions? If such connections exist, knowing how to construct an urban area that facilitates social interactions will be of great benefit to all of us. CHPC computing resources will help find the connections.

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 at data centers 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.

News

CHPC at SC11 in Seattle

by Sam Liston

CHPC will once again participate in the ACM/IEEE Super Computing Conference 2011. This year the conference will be held in Seattle, WA. As part of the participation CHPC will host a booth. In it we will showcase research being facilitated by using CHPC resources.

CHPC will also take advantage of the geographic location of the booth. Being located outside the University of Utah campus firewall we plan to put on a demonstration testing two network paths. The first path will take traffic through the existing campus firewall the second will take us through a bypass route that is not protected by the firewall, but rather is monitored for malicious activity by a device that is not inline with the data path. The goal is to demonstrate that we can achieve better performance through the bypass while maintaining a similar level of security.

For the purpose of the demo we will attempt to recreate a standard research data workflow in order to generate a significant amount of bandwidth. We will transfer files, send and receive high-definition video, and perform remote visualizations. We will make several runs, switching back and forth between the two paths and monitoring the difference in achieved bandwidth.

As an additional method for creating bandwidth during the demo, CHPC has teamed up with HP and the University of Wyoming to set up three instances of HP's x9000 hardware. x9000 is a scale-out network attached storage product that CHPC has deployed as scratch disk space for the clusters, and is working to deploy in other capacities. There will be hardware in place at CHPC, at the University of Wyoming, and in Seattle at the HP booth on the SC11 show floor. We will experiment with some of the feature of the x9000 hardware, such as replication, not only as a method for generating bandwidth for the demo, but also to understand how the hardware works in this particular setup.

SC12 Salt Lake City, Utah

Salt Lake City will host the 2012 SuperComputing Conference November 10 - 16 at the Salt Palace. Attracting thousands of participants from throughout the world, this event is the premier international conference on high performance computing, networking, storage and analysis. CHPC will have a booth at the conference that will highlight the research done with CHPC's compute resources. If you would like your research highlighted, contact Sam Liston at sam.listron@utah.edu.

CHPC's booth at SC10 in New Orleans

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
Jody Smith	Database Manager	647-3042	jody.d.smith@utah.edu	405-12 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	User Services	N/A	N/A	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
Derek Huth	Software Development Support	N/A	N/A	405-8 INSCC
Sam Liston	Systems, Multimedia	232-6932	sam.liston@utah.edu	405-39 INSCC
Kyle McGuire	User Services	N/A	N/A	405-14 INSCC
Jimmy Miklavcic	Multimedia	585-9335	jimmy.miklavcic@utah.edu	296 INSCC
Beth Miklavcic	Multimedia	585-1066	beth.miklavcic@utah.edu	111 INSCC
Michael Palmer	Systems	435-720-3261	michaelj.palmer@utah.edu	405-28 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	N/A	N/A	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

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

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, 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