Center for High-Performance Computing

NEWS

THE UNIVERSITY OF UTAH FALL 2012

Research

HPC Aids Study of Massive Galaxies

By Janet Ellingson and Adam Bolton, Ph.D.

University of Utah astronomers are using CHPC parallel computing resources to study galaxy evolution and cosmology. Adam Bolton, assistant professor in the department of Physics and Astronomy, and his research group are members of the Baryon Oscillation Spectroscopic Survey (BOSS) project within the international Sloan Digital Sky Survey III (SDSS-III) collaboration. BOSS is currently building the largest ever three-dimensional map of galaxies using a 2.5-meter telescope at Apache Point Observatory (APO) in New Mexico. Researchers are measuring the statistical patterns within this map in order to understand the nature of the mysterious "dark energy" that seems to be accelerating the expansion rate of the universe. Kyle Dawson, another member of the University's Department of Physics and Astronomy, is also heavily involved in the BOSS project.

"Baryon oscillations" refer to the primordial cosmic sound waves that are imprinted on the large-scale distribution of galaxies. "Spectroscopic" refers to the use of astronomical spectroscopy to disperse the light of galaxies into its constituent wavelengths in order to make precise measurements of the cosmological Doppler shifts ("redshifts") that encode the distances of the galaxies from Earth.

Prof. Bolton's work within the BOSS project is to develop, test, and supervise the complex software programs that take the raw digital data from the telescope and translate it into calibrated spectra, classifications, and distance measurements. Since BOSS is observing millions of astronomical objects, this analysis must be accurate, robust, and fully automated. It is routinely run in parallel on CHPC computer systems, which allows for rapid verification of results. [Figure 1]

The research group has taken advantage of massive storage and computing at CHPC to integrate additional numerically intensive scientific analyses with the core BOSS software pipeline. Dr. Joel Brownstein and Dr. Antonio Montero-Dorta, both postdoctoral researchers, have integrated multiple codes to measure the mass in stars contained within each of the BOSS galaxies. This has led to the largest current sample of uniform galaxy "stellar mass" measurements, which are being used throughout the BOSS collaboration to investigate the evolution of galaxies across cosmic time. Yiping Shu, a Ph.D. student in Physics and Astronomy, has developed and implemented a novel statistical method for measuring the characteristic velocities of the stars in the BOSS galaxies, which has uncovered evidence of greater dynamical diversity of galaxies at earlier cosmic times.

(continued on page 2)

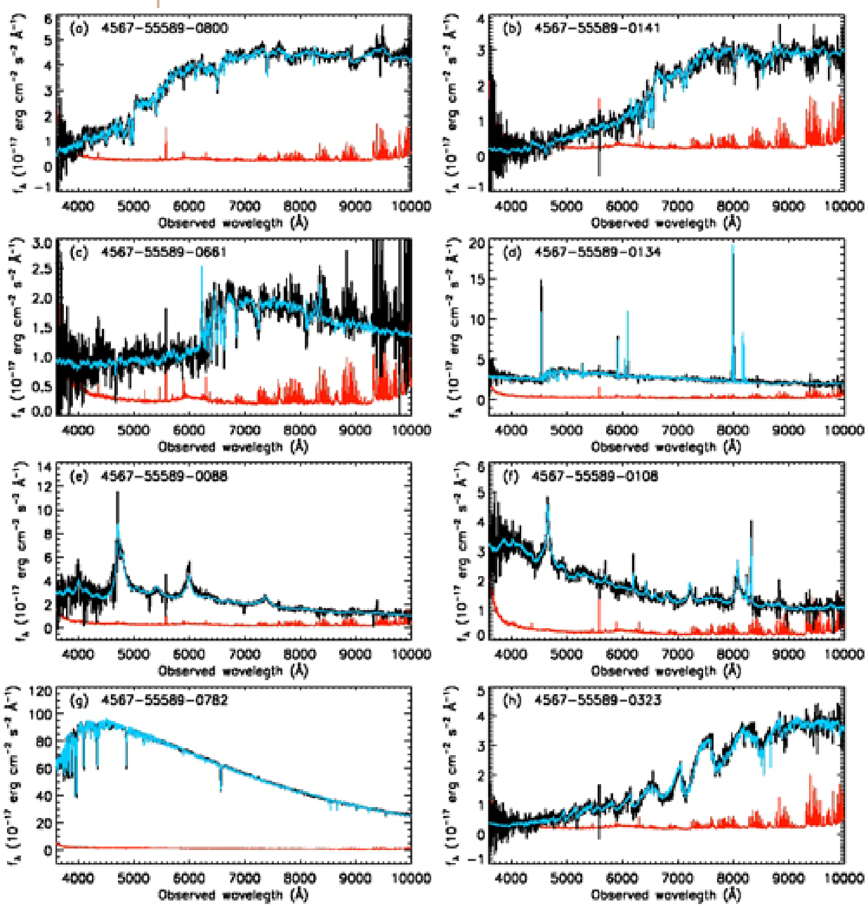

Figure 1. Representative spectra of galaxies (top 2 rows), quasars (third row) and stars (bottom row) from the SDSS-III BOSS survey. Black lines show the actual data for the variation of object brightness with observed wavelength, while blue lines show the best models selected by algorithms for the automated classification and distance measurement of these objects. Red lines indicate the noise level in the data as it varies with wavelength.

(continued from page 1)

Prof. Bolton's research also makes use of the phenomenon of "strong gravitational lensing" to measure the structure and evolution of the most massive galaxies. This phenomenon occurs when two galaxies are aligned along the same line of sight; the gravity of the closer galaxy deflects the light of the galaxy behind it, forming an "Einstein ring" whose size provides a direct measurement of the mass of the nearer galaxy. These rare objects are found through data-mining techniques that search the massive database of BOSS spectra and spectra from earlier SDSS projects. Once located, the gravitational lens is then observed with the Hubble Space Telescope in order to obtain the high resolution imaging necessary for measuring the ring's size and shape. [Figures 2 and 3]

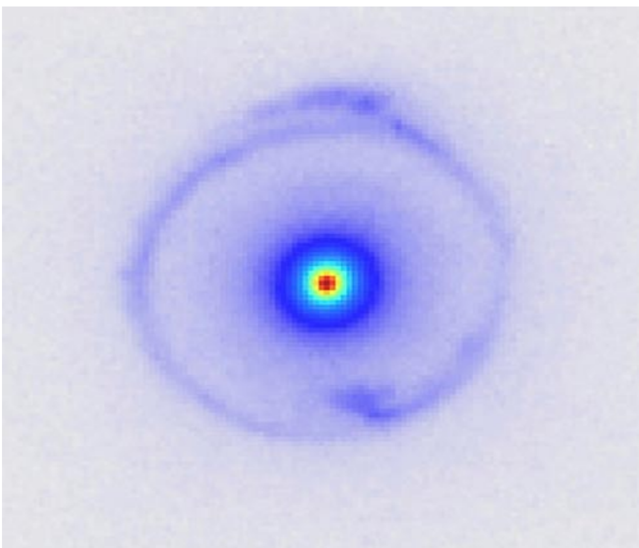

Figure 2: A Hubble telescope image of an "Einstein ring"

Working in collaboration with Prof. Bolton, Dr. Brownstein has implemented a parallel-computing-enabled suite of analysis software that allows for the rapid interactive analysis of the Hubble images of these gravitational lenses and for a quicker code development cycle than would be possible within a non-parallel framework. Their recent results have combined new BOSS-selected lenses with previously known lenses to push the technique across cosmic time and discover that massive galaxies are getting more centrally concentrated as the universe evolves. This provides evidence for the importance of major merging of galaxies with other galaxies, and was described in a recent news release from the U's College of Science: http://unews.utah.edu/news_releases/when-galaxies-eat-galaxies/

CHPC computing resources were also a crucial tool for a research project by Ryan Arneson, an undergraduate student who worked with Prof. Bolton and Dr. Brownstein. Mr. Arneson's work addressed the question of whether gravitational lenses found within BOSS and similar databases are truly representative of the parent population of galaxies from which they are selected. To do this, he assembled the ingredients to compute the observational signatures of millions of simulated gravitational-lens objects, and to select the subset that would be identifiable to existing detection methods. By running these simulations in parallel in order to generate a huge statistical sample, this project demonstrated that the lenses were indeed representative with respect to their central mass concentration. Mr. Arneson is now a doctoral student in the Department of Physics and Astronomy at UC Irvine.

Looking toward the future, Prof. Bolton will serve as Princpal Data Scientist for the "After-SDSS-III" (AS3) project, the next generation of experiments at the Sloan 2.5-meter telescope after the conclusion of BOSS and other SDSS-III projects in 2014. AS3 will run from 2014 to 2020, and will encompass three unique sub-surveys. The extended BOSS (eBOSS) project will extend the cosmological mapping reach of BOSS to previously unmapped epochs in the history of the universe. The project for "Mapping Nearby Galaxies at APO" (MaNGA) will obtain spatially resolved spectroscopy of 10,000 galaxies in the local universe, in order to build maps of their internal structure, dynamics, and star-formation activity. The Apache Point Observatory Galactic Evolution Experiment 2 (APOGEE-2) will continue the work of the current APOGEE project to map the dynamics and chemistry of hundreds of thousands of stars in our own Milky Way Galaxy. Current plans are for the CHPC to host the primary data archive for the AS3 collaboration, further capitalizing on the U's investment in the Center's staff and infrastructure to bring some of the world's most ambitious astronomical science to Utah.

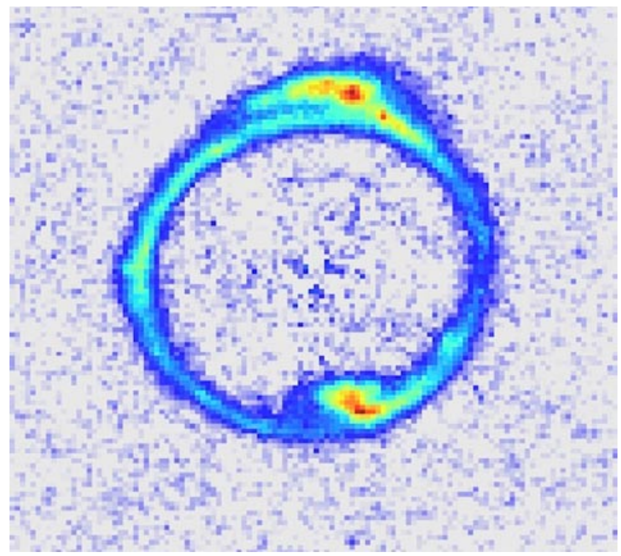

Figure 3: The same image processed to remove the light of the massive elliptical galaxy at the center, in order to reveal the lensed ring more prominently.

Research

Air Quality Modeling at CHPC

by Patrick Barickman, UDAQ

Collaboration between the Utah Division of Air Quality (UDAQ) and the University of Utah's Center for High Performance Computing (CHPC) now gives the air quality modeling group at UDAQ access to advanced computing resources. This cooperative agreement began with a request from the Utah office of the Bureau of Land Management (BLM) for consultation on air quality modeling to support environmental impact analysis needed to process natural gas drilling permits in the Uintah Basin. This collaboration between UDAQ and CHPC is now a critical element in UDAQ's ability to conduct air quality modeling for a wide variety of applications throughout the state, from the urbanized Wasatch Front to energy production areas of the Uintah Basin.

Air quality scientists at UDAQ have been doing regional air quality modeling for over 12 years to support efforts to improve air quality along the Wasatch Front. Since these models are run multiple times to assess the effect of different pollution control strategies, a significant amount of computational power to complete in a reasonable amount of time is required. UDAQ has used its own in-house computer resources to run the models for the State Implementation Plan for PM2.5 in northern Utah. However, with the large increase in computer processing and disc storage space afforded by CHPC, large scale air quality simulations over the state of Utah can be run in about one-fifth the time required in the past. This capability will allow UDAQ to assist and analyze other air quality projects of interest to the state of Utah.

While high performance computer hardware is an important part of this work, the limited number of staff with the knowledge to set up and interpret the information also limits the number of projects that can be undertaken. CHPC provides computer support personnel to solve problems in the setup of the specialized computer programs used by UDAQ. This support, plus the day to day system administration that CHPC provides, allows the UDAQ scientists to spend their time working in the area that they know best.

Currently, outside of the PM2.5 modeling being done along the Wasatch Front and in Cache County, there is need for a better understanding of the causes of episodic increases of wintertime ozone in the oil and gas producing region of the Uintah Basin. An air quality model that simulates winter ozone formation will be used to quantify the effectiveness of mitigation strategies and to tailor an emissions reduction program that is appropriate for the Uintah Basin. This modeling framework will rely on the data collected from studies in the Basin as well as state-of-the-science meteorological model output. This meteorological modeling, focused on the northeastern section of the state, will begin this fall and preliminary air quality modeling will begin early next year. Working with the Regional Modeler for EPA Region-8, UDAQ will complete an initial modeling simulation of winter air pollution in the Uintah Basin. The model results will be compared to field measurements taken during the 2011 Winter Ozone Study as a means of testing the validity of the model to predict ozone concentrations.

There is also multi-agency cooperative agreement among state and federal regulatory agencies to assess the environmental impact of oil and gas development in the region of northeast Utah, northwest Colorado and southwest Wyoming. The addition of the computational capabilities will allow UDAQ to both contribute meaningful technical resources to the project as well as take advantage of the additional research and data that will result from this regional study.

Crystal Woods, a grad student in Environmental Engineering at USU, prepares an instrument launch to measure ozone levels in the Uintah Basin

Updraft HPC Cluster - photo by Sam Liston

Spring 2013 CHPC Presentation Schedule
Tuesdays and Thursdays at 1:00 PM - 2:00 PM
INSCC Auditorium (Room 110)
(north of the Park Building)

DATE	SUBJECT
2/05	Overview of CHPC (by Wim Cardoen)
2/07	Introduction to BRISC (Biomedical Research Informatics Service Core) Services (by Bernie LaSalle)
2/12	Introductory Linux for HPC - Part 1 (by Martin Cuma)
2/14	Introductory Linux for HPC - Part 2 (by Martin Cuma)
2/19	Introduction to Parallel Computing (by Martin Cuma)
2/19	HIPAA Environment, AI and NLP Services at CHPC (by Sean Igo) *Time: 2:00 – 3:00 PM, Location: HSEB 2908
2/21	Introduction to MPI (by Martin Cuma)
2/26	Using Gaussian09 and Gaussview (by Anita Orendt)
2/28	Introduction to GPU Programming (by Wim Cardoen)
3/5	Intro to Programming with OpenMP (by Martin Cuma)
3/7	Using Python for Scientific Computing (by Wim Cardoen)
3/21	Intro to I/O in the HPC Environment (by Brian Haymore and Sam Liston)
3/26	Hybrid MPI-OpenMP Programming (by Martin Cuma)

For more detailed description of the presentations go to: http://chpc.utah.edu/docs/presentations/

Everyone is Welcome

Student Cluster Competition at SC 12

by Janet Ellingson

SC12 will be in Salt Lake City this year. In addition to sponsoring the largest booth CHPC has had at the annual event, CHPC staff are also mentoring the University's team in the student cluster competition. The team is made up of six computer science students: Leif Andersen, Bruce Bolick, Ian King, Tom Robertson, Kathryn Rodgers, and Tyler Sorenson. Professor Mary Hall, School of Computing, is the team's faculty advisor. The two-day competition requires that each team build a small cluster from component parts they will receive on the first day of the competition. This cluster, known as the LittleFe, is a complete 6-node Beowulf style cluster. Each node has two 64bit Intel CPU cores (4 hyper-threaded cores) and a cuda capable nvidia geforce embeded GPU. Brian Haymore, a CHPC hardware expert, has guided the students as they practice building their cluster with parts that replicate what they will get at the competition and as they learn how best to exploit the hardware's capabilities.

Once their machine is built, the students have the remainder of the first day to install the operating system and open-source software of their choice, test, tune and run HPCC benchmarks. Brian, along with Martin Cuma and Wim Cardoen, CHPC software specialists, are helping the students as they choose software and learn how to optimize and compile code for best results.

At the end of the first day, the teams will receive the data sets. This year the challenge is the computationally complex "Traveling Salesman Problem" (TSP). The data include three sets of cities and coordinates. The students must construct an optimal algorithm to find the shortest route that includes a visit to each city once and a return to the city of origin for each of the data sets. The teams will be judged on how frequently they have the current best score and the visualization of their results. The judges will also interview the teams to access their knowledge of HPC, parallelization and the TSP.

The teams will be working at their tasks in full view of conference attendees. If you are attending SC12, find the University's team on the convention floor and cheer them on.

Examples of Recent Research Using CHPC Resouces

Baron, R, Vellore NA (2012). "Reversible Opening/Closing Dynamics of LSD1/CoREST: Discovery of a Nanoscale Clamp for Chromatin and Protein Binding." Biochemistry 51: 3151-3153.

Beveridge, DL, Cheatham, TE, Jr., Mezei, M (2012). "The ABCs of molecular dynamics simulations on B-DNA, circa 2012." J. Bioscience 37: 379-397.

Crosman, E., Horel, JD (2012). The impact of the Great Salt Lake on Salt Lake Valley persistent cold air pools. 15th Conference on Mountain Meteorology. Steamboat Springs, CO.

Davis, K, Staes, C, Duncan, J, Igo, S, Facelli, JC (2012). "Identification of pneumonia and influenza deaths using the death certificate pipeline." BMC Medical Informatics and Decision Making 12(37).

Kim, J. (2012). Investigating Stratospheric Influences on Tropospheric Predictability. Doctoral Dissertation, University of Utah.

Lin, Z, Flores, M, Fortez, I, Henriksen, NM, Concepcion, GP, Rosenber, G, Haygood, MG, Olivera, BM, Light, AR, Cheatham, TE, Jr., Schmidt, EW (2012). "Totopotensamides, polyketide-cyclic peptide hybrids from a Mollusk-associated bacterium Strepomyces sp." J. Natural Prod 75: 644-649.

Olson, JK, Boldyrev, AI (2012). "Electronic transmutation: Boron acquiring and extra electron becomes 'carbon'." Chem Phys Lett 523: 83 - 86.

For a full bibliography go to: http://www.chpc.utah.edu/docs/research/CHPCBibliography.pdf

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 the 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.

Farewell to Dee Raynor and Byron Davis

At the end of June, DeeAnn Raynor and Byron Davis retired after over two decades of service to the University of Utah and CHPC.

DeeAnn has been with CHPC from the very beginning in 1988 when it was called the Utah Supercomputing Institute. The USI office was in the Merrill Enginering Building and its first supercomputer, the IBM 3090/600S, was in the Student Services Building. Users accessed the IBM, purchased with an NSF grant, by using telnet, a network protocol that allowed connection to the machine's operating system. For documentation, researchers used a text-based tool called "gopher," the precursor to the World Wide Web. Dee handled all the administrative duties for CHPC, including completing the purchase of the Ember cluster last year, a supercomputer considerably more powerful than the old IBM. After sitting at a desk most of her work life, Dee now plans to spend many years travelling the world with her sisters and her husband.

DeeAnn Raynor

Byron joined the CHPC staff in 1992 as the consultant for statistics and research methodology. In 1993 CHPC acquired its first dedicated stats server, an IBM RS/6000 with 256 MB of memory and 6 GB of disk space. Byron cleverly named it DV8. Over the next two decades he helped many faculty and grad students with their statistical analyses and taught an introduction to statistics to undergraduate student as an adjunct in the Department of Family and Consumer Studies. In retirement he is following his life-long passion for table tennis. In addition to teaching the sport, he and his wife compete in national competitions and frequently take home top medals.

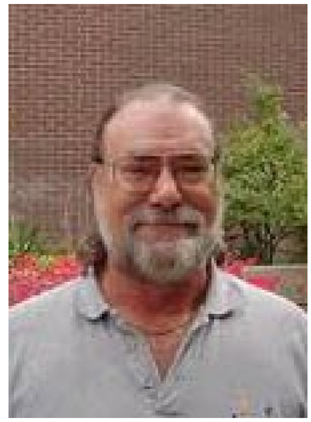

Byron Davis

Both Byron and Dee experienced firsthand the remarkable development in computational resources over the past two decades. They met those changes with good humor and intelligence to ensure that CHPC's users always had the best information and the best support as our resources evolved along with the technology.

CHPC Staff in front of the INSCC building, June 2012

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	585-1869	julia.harrison@utah.edu	430 INSCC
Guy Adams	Assistant Director, Systems	554-0125	guy.adams@utah.edu	424 INSCC
Joe Breen	Advanced Network Initiatives	550-9172	joe.breen@utah.edu	426 INSCC
Jimmy Miklavcic	Manager, Media Services	585-9335	jimmy.miklavcic@utah.edu	296 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	587-7770	martin.cuma@utah.edu	418 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 Eng/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 Admin	231-3194	irvin.allen@utah.edu	405-40 INSCC
Robert Bolton	Systems Admin	528-8233	robert.bolton@utah.edu	405-24 INSCC
Wayne Bradford	Systems Admin	243-8655	wayne.bradford@utah.edu	405-41 INSCC
Erik Brown	Systems Admin	824-4996	erik.brown@utah.edu	405-29 INSCC
Brandon Day	Tech Assistant	N/A	N/A	405-20 INSCC
Nathaniel Ellingson	Tech Assistant	N/A	N/A	405 INSCC
Jake Evans	Network Engineer	718-1526	jake.evans@utah.edu	405-22 INSCC
Steve Harper	Systems Admin	541-3514	s.harper@utah.edu	405-31 INSCC
Brian Haymore	HPC Systems	558-1150	brian.haymore@utah.edu	428 INSCC
Eric Hughes	Systems Admin	879-8449	eric.hughes@utah.edu	405-18 INSCC
Samuel T. Liston	Systems Admin	232-6932	sam.liston@utah.edu	405-39 INSCC
Beth Miklavcic	Multimedia	585-1066	beth.miklavcic@utah.edu	111 INSCC
Michael Palmer	Systems Admin	435-720-3261	michaelj.palmer@utah.edu	405-28 INSCC
David Richardson	Network Engineer	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 Admin	581-7552	steve.smith@utah.edu	405-25 INSCC
Neal Todd	Systems Admin	201-1761	neal.todd@utah.edu	405-30 INSCC
Alan Wisniewski	Network Support	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

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