Center for High Performance Computing NEWS

THE UNIVERSITY OF UTAH SPRING 2010

Research

The Study of Proton Transport in Condensed Phase

by Wim R. Cardoen, C. Mark Maupin, and Anita M. Orendt

Dr. Greg Voth, a member of the Department of Chemistry's faculty and head of the U's Center for Biomolecular Simulation (CBMS), conducts research that is extremely computationally intensive and as such is a major user of CHPC resources. In this article, one aspect of the Voth group's research in the field of computational/theoretical chemistry and biophysics is presented, specifically the work in the field of proton solvation and transport in condensed phases. The Voth group has developed the Multi State Empirical Valence Bond (MS-EVB) method to treat the dynamically changing bonding environment that occurs during these processes. This method has been successfully applied to proton solvation and transport in systems such as bulk water, aqueous weak acids, at interfaces such as the water liquid vapor interface, in water alcohol mixtures, in hydrophobic channels, trans-membrane bimolecular channels and enzymes, to name a few.

When dealing with the breaking and formation of chemical bonds, as is the case in the processes mentioned above, one typically needs to use quantum mechanics, specifically the time-dependent Schrödinger equation. However, due to the size of the systems under study, an accurate quantum mechanical treatment is not feasible in a reasonable time. By imposing several constraints the quantum mechanical formalism can be recast into a classical mechanical formalism (molecular dynamics), where the atoms are treated as hard spheres and interactions between atoms are represented by empirically developed force fields. However, these constraints imply that the property to break and form chemical bonds is lost. To overcome this challenging situation Dr. Voth and his co-workers introduced the MS-EVB method. To illustrate the method, we will consider the special case of one hopping (excess) proton (H^+) in bulk water (H_2O). In molecular dynamics the excess proton remains always bound (and therefore localized) to the same one water molecule. The excess proton can only move through bulk water by piggybacking the water molecule to which it is attached. As a consequence, the self-diffusion constant of the excess proton in bulk water (obtained by molecular dynamics) is far below its experimental value.

The MS-EVB formalism overcomes the localization issue in the following way: at every time step the excess proton will be shared by several molecules (delocalization of the excess proton). Before going deeper into the MS-EVB formalism we will introduce two important concepts, namely the

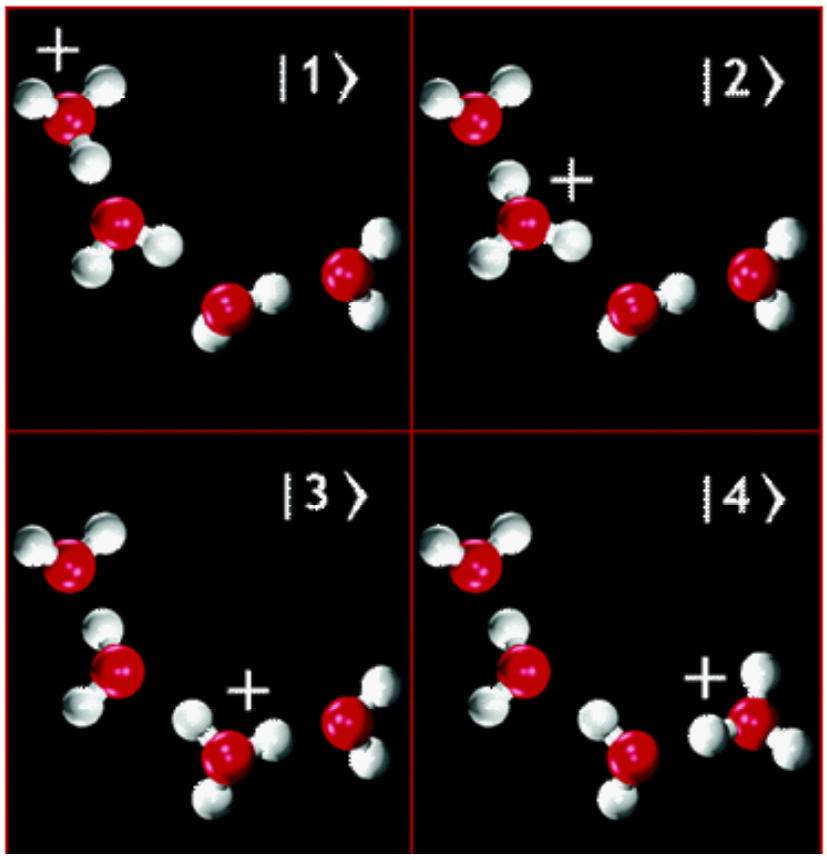

Figure 1: Molecular system containing four H_2O molecules and one excess H^+. Four MS-EVB states are represented.

MS-EVB state and the pivot molecule. An MS-EVB state is defined as the complete binding topology of a molecular system, e.g., a H_2O molecule bound to an excess proton H^+ (forming the H_3O^+ molecule) and a set of surrounding molecules (additional water molecules in the case of bulk water). The pivot molecule (molecule bound to the excess proton H^+) is defined as the reference molecule used to establish the binding topology defining the molecules that will share the excess proton. (continued on page 2)

In the first stage the molecules with which the pivot molecule will share the excess proton must be chosen. In praxi, one chooses the other water molecules within, for example, the first and second solvation shell of the pivot molecule. In the example in this article we are, for simplicity, choosing three additional water molecules. Note that at this point, we have only one MS-EVB state, namely the pivot state, along with the binding topology of the pivot molecule with all the other molecules involved in the simulation. In the second step, we generate additional MS-EVB states by allowing the excess proton to break thebond with the pivot molecule and to bind with each of the other molecules that were selected by the pivot molecule in the first stage. The MS-EVB states of our example are shown in Fig 1. In the third stage a matrix that contains all the interactions between the different MS-EVB states is built. After the diagonalization of this matrix, we find the most stable state (ground state). This ground state is a weight-based linear combination of all the MS-EVB states that share the excess proton. The forces on the atoms are also calculated and the atoms are moved to a set of updated coordinates based on these forces. This procedure is iterated many times, with the MS-EVB state having the highest weight in the ground state becoming the new pivot molecule in the next iteration, resulting in a molecular dynamics trajectory as in traditional molecular dynamics.

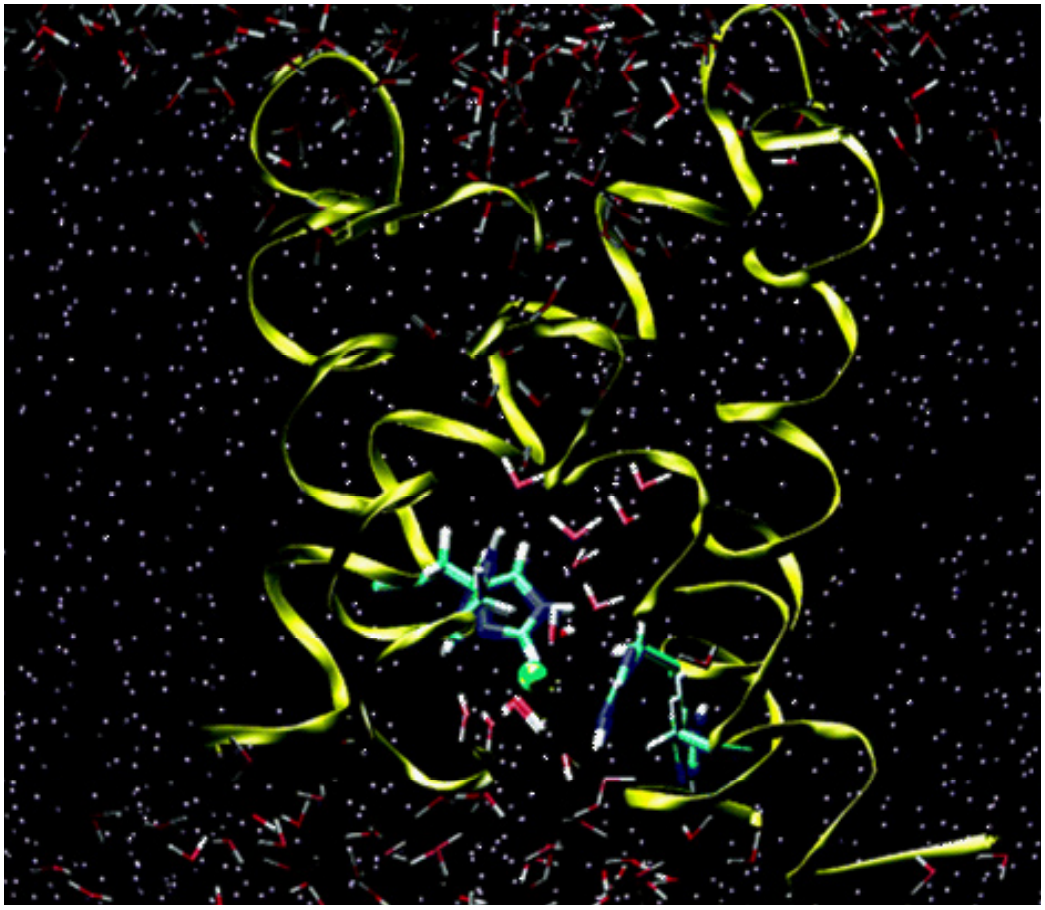

Figure 2: Snapshot from an MS-EVB simulation of the M2 proton channel (influenza A virus - the ribbon) in a DPMC lipid bilayer (blue dots). The excess proton (green sphere) is hopping through a water wire as it moves between the two sides of the lipid bilayer.

As previously mentioned, the MS-EVB method has been successfully applied to an array of interesting problems. One of the first problems that had been studied was the proton solvation and transport in bulk water. The self-diffusion properties for the hydrated proton, the likely solvation structure, and the observed deuterium isotope effects were all reproduced using the MS-EVB model. The MS-EVB model has also been used to study the mechanism of proton translocation through narrow hydrophobic channels (shown in Fig. 2) filled with water. This study revealed a significant increase of the rate of proton transport when only one string of molecules can fit within the channel.

In the example discussed in this article there is only one excess proton. Over time the MS-EVB method has been extended to the shuffling of multiple excess protons (SCI-MS-EVB) and species other than H^+. The majority of the work conducted using the MS-EVB methodology to date has centered on the properties and reactive nature of the solvated excess proton. However, the MS-EVB methodology is general in nature and can be used to simulate a wide range of reactive events. Therefore, the future of the MS-EVB methodology is to expand its ability to model reactive systems in an effort to incorporate a wider range of critically important reactions such as enzyme mediated cellulose hydrolysis for the cost effective production of bio-fuels, various biologically activated processes that use ATP hydrolysis, and chemically activated complexes for CO_2 sequestration in an effort to reduce green houses gases.

Acknowledgement: We express our gratitude to Dr. Voth for his permission to use the graphics in this article.

Article

IPv6 - What Is It and Why Does It Matter?

by David Richardson and Joe Breen, CHPC Network Administration

You've probably seen an IP address. It looks something like 192.168.0.1. Every computer on the internet has one. Today's system of IP addresses is called IPv4 (version 4). But there's a problem: IPv4 can't count past four billion (2^{32}). This may sound like a lot, but there are many devices that need IP addresses. Your computer at work, your laptop at home and your cell phone all need a unique address. There are nearly seven billion people on Earth. Four billion addresses is a lot, but it's not enough.

So far, there have been enough addresses for everyone who needed one, but they are running out. In late 2011 the pool of free addresses will be empty. The short-term work around for the impending scarcity is called NAT (network address translation). It's a way for multiple computers to share one address. You probably have a wireless router at home. If so, you have a NAT. The wireless router gets one address from your ISP and gives fake addresses to your computers and other devices.

NATs works okay for surfing the web or reading email, but it has problems. A computer behind a NAT can start a conversation, but it can't hear a new one. If two computers want to talk, but they're both behind different NATs, they're out of luck. If there are multiple layers of NATs, things get even harder.

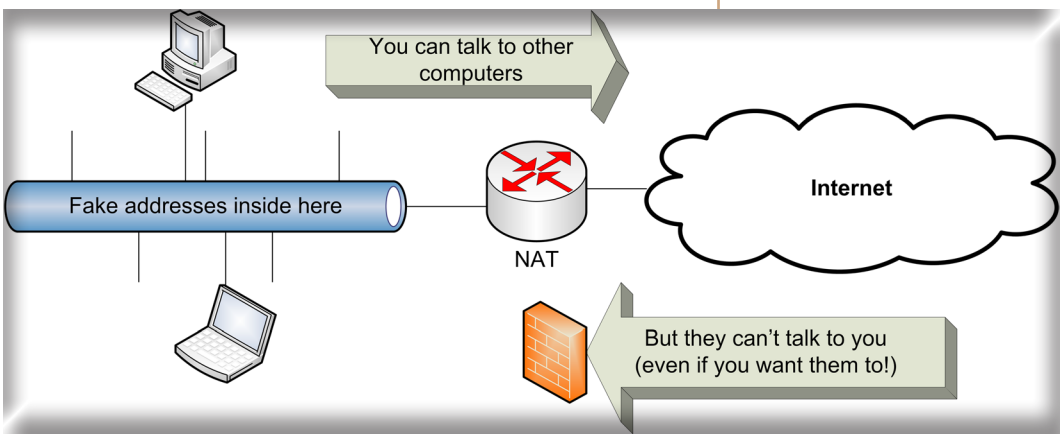

What's the real solution to this problem? IPv6 (version six). While IPv4 only has four billion addresses (2^{32}), IPv6 has 340 trillion trillion trillion (2^{128}). That's enough for every cell in every person on earth to have 3 trillion addresses. Why so many? Changing the address structure of the internet is hard. No wants to do it again later. By having such a large address space, there is room for the growth that we cannot foresee. The new IPv6 address will look like this: 2001:1948:414:3::1.

No one is saying we should drop IPv4. No one intends to suggest that there be a specific date on which IPv4 is replaced by IPv6. Instead, organizations will run both IPv4 and IPv6. ISPs are beginning to provide IPv6 addresses to their subscribers and companies are beginning to make their services available via IPv6.

CHPC has implemented IPv6 on most networks in concert with IPv4. This scenario is called "dual-stack." CHPC has structured its Domain Name System (DNS) to resolve names to IPv6 addresses, something that is very useful with the long IPv6 numbers, and has set up IP tables rules and router Access Control lists to mimic the security rules of those in IPv4. Most of CHPC's hosted webservers now have IPv6 addresses.

CHPC is pressing forward with enabling other server applications for IPv6. Not all applications support IPv6 yet, but CHPC is identifying those that do and is moving forward to support both protocol environments. From the client side, CHPC has different Operating Systems (OS) clients running IPv6 and accessing remote sites on the Internet via IPv6, as well as local services via IPv6. CHPC has also setup its enterprise monitoring and performance monitoring infrastructure to monitor both IPv4 and IPv6 devices, interfaces and applications.

Looking to run IPv6 in your area? Here are some things to consider:

* IPv6 Routing on backbone and within LAN environments
* IPv6 Point of Contact data base (if you have a large number of networks and supported groups)
* IPv6 flows for top talker reports, security visibility, etc.
* IPv6 Monitoring tools
* IPv6 DNS capability and transport
* IPv6 NTP capability
* IPv6 DHCP - necessary for client support if not using built-in StateLess Address Auto-Configuration [RFC2462].
* IPv6 stack enabled on target service boxes
* IPv6 IP tables for Unix hosts
* IPv6 enablement for any proxy services, caching services or firewall services
* IPv6 enablement of service applications
* IPv6 enablement of VPN, mail, filesystems, etc. - future services to prioritize and target

Remember, the credo for IPv6 architecture: if you do it in IPv4, you'll do it in IPv6 -- the address numbers are just bigger.

ARTICLE

Using Infiniband for Long-Distance Bulk Data Transfers

by Tom Ammon, CHPC Network Admin

The Problem

As researchers gain access to more powerful computing resources, they are also having to learn new ways to deal with very large sets of data (multiple-Terabyte sizes). These data sets are powerful sources of information that may contain the answers to many problems that society faces, and researchers are working to find better ways to use this data to solve real problems. But most researchers also use several physical computing systems, often located across the country from each other, resulting in the need to move these large data sets between computing facilities. The networks that connect these sites are unable to handle the large amounts of data. Take, for example, a 10-Terabyte data set. In theory, using a current research network, we should be able to move this data set between Los Angeles and New York City in about 8 hours. This time is calculated based on achieving the maximum 10 Gb/s transfer rate supported by the network hardware. The reality is that it often takes several weeks (based on typically achieved transfer rates of appoximately 1 Gb/s) to move that much data over current networks. It is literally more efficient to copy the data onto a couple of hard drives and ship the hard drives via FedEx to the cross-country location.

The Technologies

Current Situation: 10 GbE with TCP transport protocol

The above transfer rates, both the maximum and typical, are based on 10 Gigabit Ethernet, or 10GbE; this is a high-speed version of the same type of network that connects your desktop computer to the Internet. It is used in most high-performance research networks, such as Internet2, to connect universities and other government entities (such as DOE labs) to each other. These specialized networks were built for the purpose of allowing scientists to move data between sites around the country in a way that is much more efficient than moving data over the open Internet. The problem is that the current transport protocol, TCP, that runs over 10GbE in these research networks wasn't built for high performance - it was built for resiliance. So while the underlying network (10GbE) is fast, the protocol (TCP) is failing to move data fast enough, leading to the reduced transfer rates.

Infiniband

Infiniband, or IB, is a network that was designed from the ground up to connect servers at extremely high performance inside a single datacenter. However, the maximum distance that can be spanned with IB is 500 meters. IB provides much more bandwidth than 10GbE - the latest version provides a whopping 32 Gb/s. But more bandwidth doesn't necessarily mean that the network will work better, as evidenced by the troubles with FTP over 10GbE. IB also includes an alternative transport protocol to TCP and therefore is not constrained by the limitations that hinder TCP. The IB transport protocol can move large amounts of data much more efficiently than TCP running over 10GbE.

A Proposed Solution

If IB is capable of moving large data sets much more efficiently, is there a way to blend the performance characteristics of IB transport protocol with the already-established research network infrastructure? The key is in a technology that encapsulates, or translates, the data that travels over IB into a form that can be transported over the TCP/IP research networks that we currently use. This is the technology that we demonstrated at SC09. The concept is illustrated in the figure below.

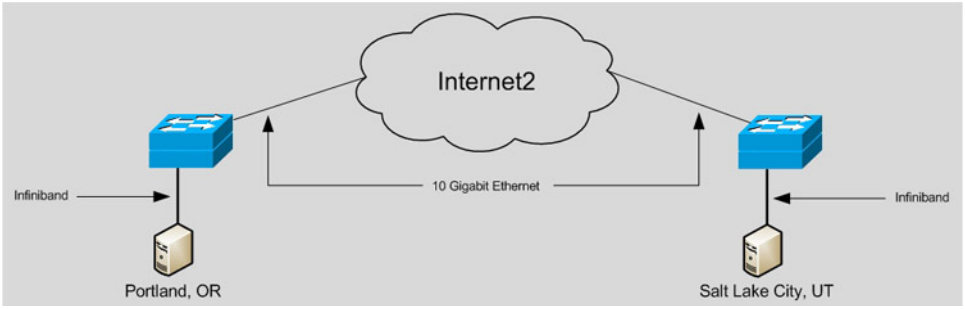

The blue boxes are devices made by Bay Microsystems (www.baymicrosystems.com) for the purpose of connecting multiple infiniband networks over long distances, over traditional TCP/IP networks.

During our demonstration at SC09, we transferred data from a server at the University of Utah to a server which was located in Portland, OR, at a rate of 4Gb/s using this technology. The only reason we stopped at 4Gb/s was because the people operating the networks connecting the University of Utah to Portland wanted to make sure that we didn't crowd out other users of the network. We could easily have filled the 10GB network pipes with our single data set!

In other words, a solution to the problem of data transfer efficiency over long-distance networks has been demonstrated. This means that a researcher can now move his

data sets between locations around the country, and the data will make it to its destination faster than the FedEx truck will. This will also enable much larger-scale collaboration between scientists, which will, in turn, lead to better science.

Conclusion

As computational resources become more powerful, they also generate larger and larger sets of data. The ability to efficiently move or copy these data sets between distant locations is key to enabling collaboration between scientists and to enable scientists to make the most of the computational resources that they have access to, regardless of their physical location. It is clear that the current approaches to transferring large data sets are struggling to keep data moving across the network with adequate speed and efficiency. Novel ways of moving data, including the hybridization of long-distance and datacenter-distance protocols, as we showed in our SC09 demonstration, may provide the solutions that researchers need.

CHPC Expands Windows Support Staff

CHPC administers desktop and laptop computers for many of the users in the departments of Atmospheric Sciences, Physics and Astronomy, Chemical Engineering, Chemistry, Geophysics and Geology. With the recent addition of Biomedical Informatics to our responsibilities, we added another member to the Windows support team. Eric Hughes brings 15 years experience to our staff. He was the IT director at the Jack Johnson Company in Park City. Before that he was a systems administrator at the U's Department of Pediatrics. As well as playing nice with Irv and Steve, Eric golfs on the weekends and cheers for the St. Louis Cardinals, his hometown team.

We also have support teams for Linux and Mac machines. A note of interest: Linux leads as the preferred OS among CHPC supported desktops and laptops with 166. Windows follows with 147 and Mac tags along with 91.

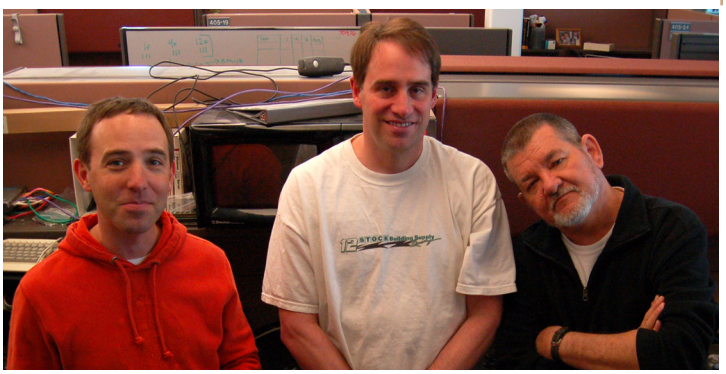

Eric Hughes, Irv Allen and Steve Smith, CHPC's Windows Support Team

For CHPC Users:

Batch Jobs with Virtual Wallclock Time Scaling

by Brian Haymore, HPC Team Lead

Since the deployment of the Updraft cluster we have had requests from users for help in handling jobs that could run just before one of the system reservations, weekly DATs, or large bigrun jobs. As we approach one of these events, it is difficult for users to submit work targeting this decaying window of time. One approach is for users to watch the system closely and change their walltime request in their job submission script. This approach proves to be taxing and awkward as many of these lost time windows happen late at night and early in the morning, leading to the underutilization of the time leading up to one of these events.

Over the past couple weeks we have been testing a new option to assist users in utilizing the time before these events. The new option is called **Virtual Wallclock Time Scaling**. An important requirement for using this option is that the job must be able to handle being killed before the full requested wallclock in the batch system script. In order to make use of this option, users need to specify both a walltime and a time value for minwclimit – the minimum time that the job must be able to run. This is done by the following line in the batch script:

```
#PBS -l walltime=24:00:00,minwclimit=2:00:00.
```

In this example the system knows that while you want to run for 24 hours you would take as little as 2 hours, or any amount of time in between. The system will find a place to start your job and then once it runs up to the minwclimit value, it will be evaluated to see if your job can be extended. If we have reached the start time of the scheduled event, your job would not be extended. If the system can extend your job it will do so in small increments until it either runs into a conflict or reaches the walltime limit specified for the job. This submit time option allows users to run jobs more easily in situations where the exact available walltime is not certain or is in a constant state of change. If you would like to explore this option please feel free to contact us at CHPC, issues@chpc.utah.edu, for more information.

Congratulations to Dr. Anita Orendt

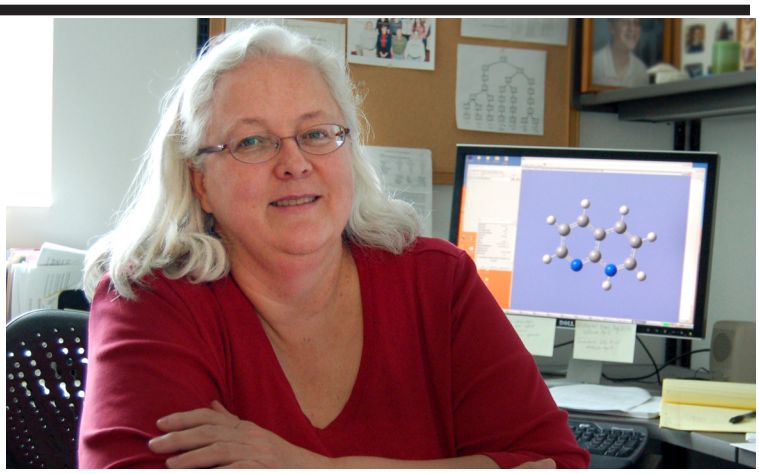

Anita Orendt, Ph.D., CHPC's molecular science specialist, received the W.W. Epstein Outstanding Educator Award from the University's Department of Chemistry. Graduate students in chemistry determine the recepient of this award and they chose to honor Dr. Orendt this year. The award is a well deserved recognition for the help she gives to all researchers who need assistance with the computational component of their research.

FYI

Published Research Using CHPC Resources

Examples of recently published research that used CHPC resources:

Bankiewicz, B. H., L. K.; Ratkiewicz, A.; Truong, T. N. (2009). "Kinetics of 1,4 Hydrogen Migration in Alkyl Radical Reaction Class." *J Physical Chem A* 113: 1564-1573.

Demille, R., Molinero, V. (2009). "Coarse-grained Ions Without Charges: Reproducing the Solvation Structure of NaCl in Water Using Short-Ranged Potentials." *J. Chem. Phys.* 131(3): 034107.

Du, J., Wright, G., Fogelson, A. (2009). "A Parallel Computational Method for Simulating Two-Phase Gel Dynamics." *International Journal for Numerical Methods in Fluids* 60: 633-649.

Lu, J., C. Deser, and T. Reichler (2009). "Cause of The Widening of the Tropical Belt Since 1958." Geophys. Res. Lett 36: L03803, doi:10.1029/2008GL036076.

CHPC maintains on its web site a complete list of publications and presentations that acknowledge the use of CHPC's resources. You can find the current listing at the following address:

http://www.chpc.utah.edu/docs/research/CHPCBibliography.pdf

If you utilize CHPC resources in your research, please include an acknowledgement in your publications and presentations. Also, please send us a copy for our records.

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

Technical Support Staff	Group	Phone	Email	Location
Ty Adams	User Services	N/A	ty.adams@utah.edu	405-28 ISNCC
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
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
Jimmy Miklavcic	Multimedia	585-9335	jimmy.miklavcic@utah.edu	296 INSCC
Beth Miklavcic	Multimedia	585-1067	beth.miklavcic@utah.edu	111 INSCC
Victor Morales	User Services	N/A	N/A	405-14 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
Neal Todd	Systems	201-1761	neal.todd@utah.edu	405-30 INSCC
Paul Vandersteen	User Services	N/A	paul.vandersteen@utah.edu	405-19 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, 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