


ICL NEWSLETTER

April 2008

March Issue of *CTWatch Quarterly* Released



Harnessing and managing massive amounts of data in real time for solving real world problems as they are occurring is a challenging task. By exploring the diverse demands that are placed on supercomputing in urgent computing, this issue compellingly demonstrates the critical nature of a dependable cyberinfrastructure where computing failure can have disastrous consequences.

Pete Beckman of Argonne National Laboratory, the guest editor for this issue, has assembled a group of experts whose work in this emerging supercomputing arena is serving as the standard for applying high performance computing to time sensitive problems. In "Urgent Computing: Exploring Supercomputing's New Role," the authors explore the many challenges and crucial applications of on-demand, urgent computing, including space shuttle damage analysis, severe weather forecasting and warning, and real-time medical simulations.

Also included in this issue are some reflections on the past and future of *CTWatch Quarterly* by its Editor-in-Chief, Jack Dongarra. Please visit <http://www.ctwatch.org/quarterly/> at your convenience to get your copy. This issue marks the 14th and final issue of the quarterly.

Recent Conferences

VI-HPS Tuning Workshop - Aachen, Germany
March 5-7
Karl

SIAM Conference on Parallel Processing for Scientific Computing - March 12-14
Dan, Jakub, Wes, and Stan

High Performance Computer Science Week (HPCSW) - Chicago, IL March 31 - April 5
Julie, Keith, Shirley, Tom

Upcoming Conferences

VGrads Workshop April 6-7
Asim, Shaun

2008 PET Technical Overview April 7
David C., Tom

Recent Papers

Fault Tolerance Management for a Hierarchical GridRPC Middleware [pdf]
Aurelien Bouteiller and Frederic Desprez

Detection and Analysis of Iterative Behavior in Parallel Applications [pdf]
Karl Fuerlinger and Shirley Moore

Tip of the Month from Tom Cortese



Machines within machines – the wonders of VMWare

Due to an unforeseen lack of compatibility between Debian Linux (running on just about every cluster to which I have access within UTK) and a software package which I am evaluating, I ended up installing VMWare Server on my Windows laptop in order to create a “virtual machine” running Ubuntu Feisty Linux.

After much moaning and cursing (just ask Barry!), I finally got everything installed and running, and it is just great! Here are some of the “gotcha”s:

Do not use a “beta” version of VMWare Server!

Your Windows laptop must have a non-blank password (and you may need to re-boot a few times after the initial installation), else you won’t be able to log in to the VMWare control panel.

Choose “Network Address Translation” when setting up a new virtual machine to allow it to access the internet via the Windows connection. Transferring files between the host and virtual OS can be accomplished by installing VMWare Tools, but I chose to use a third machine (or a CD-RW) as an intermediary.

Once the virtual machine is created, open “Edit Settings” and change the CD-ROM to point to either a “.iso” image or the physical CD drive containing the virtual machine OS installer.

Once the virtual machine OS is installed, connect Windows to the internet, switch to the virtual machine, and use the “`sudo apt-get install <package>`” command to install all of the packages that one would hope to find (even though Ubuntu Feisty already had a “gcc”, I had to download and install another one since it didn’t work), e.g.,

```
sudo apt-get install gcc
sudo apt-get install g77
sudo apt-get install fort77
sudo apt-get install g++
```

Here’s a good one: Once your cursor is active inside the virtual machine it becomes constrained within that window, hit “Ctrl-Alt” to free it.

Voila! Now you can run Ubuntu Linux within a Windows window, transfer files and/or connect to the internet using the Windows network connection, and easily switch back and forth between them!

Recent Lunch Talks

03-07-2008 **Julien Langou**
The Convergence of Restarted GMRES for Normal Matrices is Sublinear [pdf]

03-14-2008 **Pierre Lemarinier**
Large Scale Self-Stabilization [pdf]

03-28-2008 **Yuanshun (Shaun) Dai**
Uncertainty Analysis in Software Reliability Modeling by Bayesian Analysis and Maximum-Entropy Principle [pdf]

04-04-2008 **Franck Cappello from INRIA-Futur**
Grid'5000: Exposing Distributed Systems to their Limits

Upcoming Lunch Talks

04-11-2008 **Stan Tomov**

04-18-2008 **Peng Du**

Interview with Heike Jagode

Senior Research Associate



Welcome to ICL, Heike!

Where are you from?

I am from the eastern part of Germany, Dresden. I was born and grew up in the former German Democratic Republic (DDR). I was quite young during the German reunification in 1990, an extraordinary event.

Where did you receive your education? Where did you work before joining ICL?

I received my first Master's Degree in Applied Mathematics from the University of Applied Science Mittweida (a university also located in the eastern part of Germany) in 2001. After graduating I started work at the Center for Information Computing and High Performance Computing (ZIH) at the Dresden University of Technology (TUD).

In 2006, I earned my second M.Sc. in High Performance Computing at the University of Edinburgh (Scotland) – Edinburgh Parallel Computing Centre (EPCC) – and afterwards I returned to my position as a research staff member at ZIH, TUD.

How did you hear about the Innovative Computing Laboratory? When did you first meet Jack? What made you want to work for ICL?

I met Shirley Moore at the kick-off meeting of the VI-HPS project in Julich in July 2007. We were talking about ORNL and possible collaboration between UT and the lab. Later, when she came to Dresden for a couple of days, she invited me to give a talk at ICL the following October. I came here for one week to give two talks and to learn a little bit more about ICL and ORNL. That was also the week when I met Jack for the first time.

At that time I wasn't sure whether I could do it: come here alone, for a certain period of time and work in a foreign country ... sometimes it's scary. Anyway, I keep myself motivated to learn as much as I can from experts in the HPC area, and for that reason I would say ICL is a good place to be.

Can you explain your dual position with UTK and ORNL?

I am considered a full-time employee at UT, even though roughly half of my time will be spent at ORNL working in Jeffrey Vetter's group. My current position includes performing basic research in computer science for performance of HPC applications and architectures, focusing primarily on developing methods and tools for automatic performance tuning on HPC applications. Since Jeff's group at ORNL is involved in performance modeling, his initial idea for me was to work at the intersection of modeling and autotuning, which equally includes ICL.

Are you collaborating with anyone outside of UT and ORNL?

I am still in collaboration with EPCC in Edinburgh working on multi-dimensional FFTs and Blue Gene related things. Furthermore, the plan is to keep me involved in the OE project (stands for OTF and EPILOG), working on a common trace format. The latter is part of the VI-HPS (Research Centre Julich) project and involves collaboration with the Dresden University of Technology.

What are your occupations/hobbies outside work?

I have been horseback riding for about 18 years. My horse is a huge red gelding (Trakehnen breed) with big ears and a strong character. He and I were quite busy winning a couple of dressage awards (it is like a "dance contest" where you have to study and perform your own choreography on your horse). I have to say that he was actually the one who taught me in my younger days rather than me teaching him some fancy moves. Hope I can find an animal friend here, too.

Tell us something about yourself that might surprise everyone.

I am afraid of human beings! That makes me keep my pepper spray always unlocked.

People

Visitors



March 18th – 20th
Emmanuel Jeannot
from INRIA in Villers lès Nancy, France



March 2nd – 11th
Marc Baboulin
from CERFACS in Toulouse, France



March 3rd – 11th
Julien Langou
from the University of Colorado at Denver



March 10th
Bo Kagstrom
from the University of Umea in Sweden



March 10th
Petter Bjorstad
from the University of Bergen in Norway

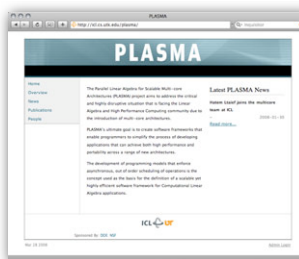
Departure



After nearly 15 years with ICL,
Jan Jones is leaving to join her
husband Reed in New Zealand. Her last
day with the group is April 30th. We'll miss
you Jan!

News

PLASMA Online



The Parallel Linear Algebra for Scalable Multi-core Architectures (**PLASMA**) project now has a Web presence. According to ICLer Jakub Kurzak, "The direct goal of the PLASMA project is to develop a framework for a dense linear algebra library, a successor to LAPACK and ScaLAPACK. The

main motivation for this effort is poor performance of these libraries on multi-core architectures. PLASMA will rely on explicit expression of parallelism through dependencies between tasks. Based on this information, PLASMA will employ dynamic, data-driven, scheduling of tasks and handle the underlying communication appropriately to the actual hardware architecture (SMP, NUMA, MPP, etc.) Different language constructs are under investigation for the algorithm definitions. We are mostly looking at solutions appropriate for massively parallel, PetaFLOP systems (parametrized DAGs), while keeping an eye on recent developments for small scale systems (e.g., Cilk, Intel Threading Building Blocks). The indirect goal of PLASMA is to guide solutions for software development for general-purpose computing on parallel systems (server and desktop multi-core computing)."

New Collaboration

ICL is part of a new NSF-funded collaborative effort with the University of Oregon, NCSA, and the Pittsburgh Supercomputing Center (PSC) called Petascale Productivity from Open, Integrated Tools (POINT). According to ICLer Dan Terpstra, "The POINT project gives ICL an opportunity to work formally with several researchers with whom we've already developed strong informal relationships, including Allen Maloney and Sameer Shende at University of Oregon, as well as Rick Kufrin at NCSA. This is a great opportunity to develop a robust performance infrastructure in which the whole is greater than the sum of its parts."

[More information from [HPCwire](#) about the project]