

ICL Newsletter February 2007

ICL Winter Banquet

The annual Winter Banquet was held on Thursday, February 1st from 5:30-7:30 at the Riverside Tavern. Forty-five or so ICLers and guests enjoyed the evening to celebrate the many outstanding accomplishments of 2006. Time went fast, however, as did the food. Chicken fingers, stuffed mushrooms, quesadillas, and vegetable spring rolls were the main buffet items along with fresh fruit, vegetables, shrimp cocktail, and other appetizers. A good time was had by all as the pictures below can attest.

ICL Winter Banquet 2007



People



We had a visit from one of the first ICLers, **Adam Beguelin**, VP of AOL video, who gave a talk kicking off our new lunch Seminar Series. This was followed by a talk by **Marc Garbey**, the department chair at the University of Houston.

The first week of February, **Julie** is visiting from Denver the ICL East laboratory!



February also marks the end of two appointments in ICL. **MyungHo Kim**, our visitor from Soongsil University in Korea, left after spending 1.5 years here. This was his second lengthy visit to ICL.

And the longest term ICLer (other than Jack), **Keith Moore**, left ICL this month to pursue something a little different.



January was a very busy month and we are happy to have a calm month to catch up on things before we get into spring.

Call for Papers

Power Architecture developer conference 2007, deadline: April 20, 2007 [[more info](#)]

Copper Mountain Multigrid Methods Conference, deadline: February 10, 2007 [[more info](#)]

2007 World Congress in Computer Science, Computer Engineering, and Applied Computing, deadline: February 20, 2007 [[more info](#)]

Did you know?

TOP 5 longest term ICLers

- 1 - Jack
- 2 - Tracy R.
- 3 - Shirley
- 4 - Tracy L.
- 5 - Scott

Friday Lunch Presentations

2007-01-05 David Cronk
Prototype Visualization Tools for Multi-Experiment Multiprocessor Performance Analysis of DoD Applications. [[PDF](#) 1.2M]

2007-01-12 Jakub
Mixed-Precision Iterative Refinement Using Cholesky Factorization on the CELL Processor [[PDF](#) 1.0M]

2007-01-19 Adam Beguelin of AOL, ICL alumni, *AOL Video* [[PDF](#) 2.5M]

2007-01-26 Marc Garbey from the University of Houston, *Innovative Computational Medicine* [[PDF](#) 5.1M]

2007-02-02 Thara, *Binomial Graph: A Scalable and Fault-Tolerant Logical Network Topology* [[PDF](#) 1.0M]

Jack introduced in Germany as the “Pope of Supercomputers”



“Last month I went to Germany for the [Jülich Research Center's](#) celebration of a Parliamentary Evening in Berlin on January 22, 2007. The major topic was supercomputing and computational science. I was speaking on the future of supercomputing, the role of Europe and the international standing of Jülich in the area of High Performance Computing.” [[see press release](#)]



Friday Lunch

2007-02-09 Julie.

What's up in Colorado?
Snow/LAPACK/Snow/LAPACK/
Snow/...

2007-02-16 Dan
PAPI


2007-02-23 TBA

Talks

2007-01-25 Jack, Oak Ridge, TN
Exploiting the Performance of 32-bit Floating-Point Arithmetic in Obtaining 64-bit Accuracy (Computing on Games) [[PDF](#) 0.9M]

Plasma Project

A hot project has emerged recently at ICL. PLASMA, **Parallel Linear Algebra for Scalable Multi-Core Systems**, (a bow to Piotr for the acronym) is going to be an investigation into future directions for development of numerical software, and linear algebra in particular, on emerging, next generation architectures. The ultimate goal would be a solid numerical library delivering a large portion of the functionality of LAPACK, which is still a long way off.



In the meantime, answers are still needed for how to address the obstacles LAPACK and ScaLAPACK face to achieve the best performance on the new machines. One of the powerful tools in this pursuit will be the utilization of data-flow and data dependency analysis techniques similar to those employed by super-scalar processors, like Scoreboarding and the Tomasulo algorithm. Many new directions have been inspired by developments related to the CELL processor, which has proven to be an excellent testing ground for novel ideas.

A likely collaboration is emerging with [Barcelona Supercomputer Center](#), where very interesting developments are taking place under the project name [CELLSs](#) (CELL Super-Scalar). Presently, CELLSs offers the ultimate automation, basically allowing the programmer to code an algorithm like Cholesky factorization with as little effort as writing three nested loops. It has actually been done for the Cholesky factorization and, quite expectedly, delivers less than desirable performance. On the other hand, the work on hand-written Cholesky for the CELL has just been finished at ICL, resulting in an algorithm that delivers the ultimate performance at the cost of mind-blowing complexity. The goal of the collaboration will be to close the gap between the automation and performance with the ICL team putting the emphasis on the latter, while trying to deliver as much functionality as possible utilizing the former. Although the CELL processor is the point where the interests of the two groups originally intersected, the ultimate goal is to apply the technology to non-CELL processors, as well as distributed memory systems.

Conference

2007-02-08/07 DoD
benchmarking and portfolios
meetings La Jolla and San Jose,
CA david c. and Shirley

2007-02-19/23 [2007 SIAM Conference on Computational Science and Engineering](#) Piotr and Jakub

2007-02-26/27 [First Annual Cray Technical Workshop](#) Nashville, TN David C. and Dan

2007-02-21/22 GARPA-2
Workshop (Geoscience
Application Requirements for
Petascale Architectures) San
Diego, CA Karl

SiCortex by Philip Mucci



[Philip Mucci](#) is a computer scientist and consultant. He is currently working as a consultant for [SiCortex](#) in addition to his permanent positions as a Research Scientist for the [Center for Parallel Computers](#) (PDC) at the [Royal Institute of Technology](#) in Stockholm, Sweden and as a Research Consultant for the ICL. His research is focused on the performance and optimization of high-end technical compute platforms.

By now, many of you have heard of [SiCortex](#), the supercomputer startup based in **Maynard, Mass.** We've received loads of press since our announcement at [SC06](#), including The Registers' **TOP FLOP Award** and **HPCWire's top 5** companies to watch. In addition, we recently announced that our previous chairman, **Bob Metcalfe** (inventor of Ethernet if you failed CS History 101) stepped down to make room for **John Rollwagen**, former CEO and Chairman of Cray for numerous years.

Our flagship product is the [SC5832](#), a **six Teraflop** massively parallel machine that consists of **972 chips each with six cores running at 500Mhz**. Each chip is a SOC or System on a Chip design, with two memory controllers, pci express six modified MIPS64 cores, L1 and L2 caches, Coherency engine, Fabric Switch and micro programmed DMA engine. If you had the opportunity to see the 27 chip board, the integration is impressive, a pizza size board with just chips and DRAMS.

The system is diskless, with massive external I/O bandwidth, and up to 108 pciE slots can be populated with FC, IB or Gig/10G Ethernet. The system is an ultra-low power design with each CHIP (a six GFLOP unit) consuming around 10 Watts. The system has been designed for performance and low power from the outset, with the ultimate goal being maximum return on investment.



My dad being recognized for HPCWire's top 5 companies to watch.



SC5832

Just how cost effective is it? Well, internally **we suspect that our 6T system will outperform commodity clusters 3-4 times it's size on applications and draw far less power**. One fully populated 6T machine draws 17KW. The smaller version of the product (.8TF), draws less than 2KW of power and can be plugged into a wall outlet and cooled with a commodity air conditioner.

The system has a very special topology called a [Kautz graph](#), which essentially allows us to build arbitrary size machines with a fixed degree yet near logarithmic diameter. You can't do that with a hypercube. At the network level, it has hardware RDMA semantics...with very tight processor/memory integration. The DMA engine deposits data into L2 for maximum performance. Just how fast is it? We are predicting near 1us latency for ping-pong MPI between two nodes across the machine. Each chip has three independent paths to every other chip, each path supporting 2GB/s. Interestingly enough, the DMA engine is microcoded and that code will be made open-source for those that are interested in doing other forms of research on communication API's. There is hardware assisted broadcast and barrier semantics for maximum scaling.

The cores do not have to be involved, which is a huge win on larger processor counts.

Software wise, we are an **open source company**. The system runs a 64 bit 2.6 Linux in little endian mode. The file system, including root, is Lustre, which we have done lots of work on for robustness and scalability. Different from other clusters, we have provided the whole software stack from day one. **No need to download and compile various libraries** in order to get your app up and running. Almost everything is there, tested and supported. Various components like FFT, BLAS and the math library have been extensively **hand tuned** to provide **maximum performance** out of the cores. We also are providing the **most complete set of performance and analysis tools** of any Linux based HPC vendor out there.

Where are we now? We just received our first batch of chips over Christmas and as of today we have MPI applications running internally on a modest number of nodes. The first tape out of the chip has been a great success with no show-stoppers. The chip, while very modest compared to your billion transistor Intel heating unit, went through massive verification cycles.

SiCortex has now been in existence about five years and fully funded for the last three. It was started by two Alpha architects who had the idea, and then they needed a business guy. That guy turned out to be father John, an ex Dec and Thinking Machines exec. I was brought in very early to work on the software stack for the machine, as well as what kind of hardware and performance the machine needed to have to be competitive. It has been an incredible learning experience for me, not only to work so closely with my father, but also to have a hand in everything from working with the design teams on the performance hardware and the DMA engine (including my first patent) to working on the Linux kernel to fleshing out a software plan for the entire machine.

My role is currently a consultant and that is probably how it will stay, being the CEO's son can be politically sensitive as you might imagine. Plus I like my ties to academia. The company now has around **55 people**. With the power of agile design teams and proper hands-off management, we have been able to deliver a world class product, coupling massive savings in cost of ownership with similar gains in delivered productivity. Don't believe me? We'll it is very likely that many of you will have access to early systems by mid-summer of 2007. The proof is in the pudding as they say. I will end this note to say that we are actively hiring software folks. Should you be interested, please drop me a line. - Phil

Interview: Karl Fuerlinger (ICL's newest Post-doc)



Where are you from? Where did you graduate from?

I am from Salzburg, Austria, studied Mathematics and Computer Science there and then went to Munich, Germany to get a PhD. Salzburg is best known to Americans probably by that "Sound of Music" film. Interestingly the movie is not at all well known in Austria (it has probably been shown once or twice on Austrian TV only), yet you can get on several Sound of Music Tours in Salzburg, it is a place that lives from tourists, after all...

Munich on the other hand is home to the famous Oktoberfest that attracts some 6 million visitors each year, a large (and loud) fraction of them coming from overseas.

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

I do know Felix for quite some time. I was working on a performance analysis project related to Kojak called Periscope for my PhD. And we had some collaboration on an European-level project on automated performance analysis called APART. Felix told me about the position at ICL in summer and since I always wanted to go to the US this was the perfect opportunity. I only met Jack personally for the job interview which was in Germany on occasion of the EuroPVM/MPI conference in Bonn.

Is this your first work experience in the US? What is the most striking difference between US and Germany?

Yes, it is my first work experience in the US.

On which project are you working at ICL? What are your objectives? Who are your external ICL collaborators?

I'm working on Kojak/Scalasca I'm also hoping to continue some of the work I did in Munich. We plan include support for UPC: Unify Parallel C and other PGAS (partitioned global address space) languages in Kojak/Scalasca. And I'm also interested in many other things like ScaLAPACK, multi-core architectures....

Was your integration in the group easy?

I had the opportunity to attend an ICL party even before I was a member (at Supercomputing in Florida), so that was easy.

Is US and especially Tennessee like you expected?

I didn't really have any expectations. There are obviously many differences and some things I like better and some I like less.

I manage to have quickly my driving license, to buy a car, to get a bank account, and so on.

I already have one of TN specialties: tipsy cake: a cake soaked in Jack Daniel's Whiskey....heavy food ...

What would you be doing if you were not a computer research scientist?

I would probably be a forester because my father is one.

What are your occupations/hobbies outside work?

I like to hiking and some mountaineering. I like to be in the nature.

I already drove around and it was quiet a foggy day that I did not see that much. I guess this is why they called it "Smoky Mountain".

What is you favorite beer?

Augustiner Helles. Augustiner is one of the last smaller and mostly privately owned breweries in Munich.

Tell us something about yourself that might surprise everyone.

I do like ice-climbing (climbing frozen waterfalls).

New hardware at ICL

PS3 Cluster: ashe.cs.utk.edu

Hardware: Sony PS3 60 GB, 1 Gigabit Ethernet switch and a HD-TV

Software: Fedora Core 6, Yellow Dog Linux 5.0, OpenMPI

Project: OpenMPI, iterative-refinement.

[[more info](#)] [[HowTo: Huge TLB pages on PS3 Linux](#)]



Intel Cloverton machine now available : brutus.cs.utk.edu.

Hardware: 2x4 (dual processor, quad core) machine with 16G of RAM

Software: Fedora Core 6.

ICL reminders

Wednesday, February 14th 2007 – Valentine's Day

The newsletter is available online at: <http://icl.cs.utk.edu/newsletter/2007/icl-news-2007-02.pdf>

Comments and contributions can be sent to: icl-news@cs.utk.edu

