

Minutes of the HiPCAT Meeting
22 Sep 2005
Authors: P. Teller and J. Yu; edited by A. Kerstens

HiPCAT attendees:

- Borries Demeler (UTHSCSA)
 - Nick Grishin (UTSW)
 - Stuart Johnson (TACC/UTSW)
 - Lennart Johnsson (UH)
 - Kirk Kirksey (UTSW)
 - Chuck Koelbel (Rice)
 - Jeremy Mann (UTHSCSA)
 - Rosalinda Mendez (UH/TLC2)
 - Shan Mishra (UTSW)
 - Suzanne Montague (UTA)
 - Lee Panetta (TAMU)
 - Richard Scheuermann (UTSW)
 - Phil Smith (TTU)
 - Warren Smith (TACC)
 - Pat Teller (UTEP)
 - Jae Yu (UTA)
-
- Nick Grishin hosted the meeting. Nick orchestrated a well-planned and interesting program.
 - The meeting opened with the welcome from the Dean of UTSW, Dr. Alfred Gilman. He said that UTSW will soon have a new imaging facility – it will be possible to image everything including thought processes in the brain. UTSW had modest computational resources 10 years ago; now with computational methods (collaborating with TACC—onsite person migrating applications to TACC systems) the computational resources and computational collaborations are important. Applications include predicting protein structures, evolution of proteins, and modeling of complex biological systems. New endowment: the Green Center of Systems Biology (GCSB) has intense computational needs and recently recruited two people with PhDs in Math who worked 5 yrs with Microsoft followed by some companies focused on the genome. GCSB is looking forward to attracting more people at this time and has strong collaborations with UT-Dallas and UT-Arlington. UTSW has a new bioinformatics program and -building on campus.
 - As the introduction progressed, the discussion smoothly transitioned into cross-chained education and bio-informatics. Pat Teller brought up the effort in education between UTEP and TTU and emphasized the focus on non-CS majors. Lennart Johnsson emphasized the importance of training programs in attracting external funds. Warren Smith mentioned that Jay Boisseau of TACC is pursuing a plan for a broad, bio-science focused program as part of the EPIC program. Lee Panetta raised a question on the definition of bio-informatics. After this discussion, the definition of

bio-informatics still seemed quite broad; including the goals of bio-informatics to link biologists to CS. Phil Smith mentioned that TTU is looking for a bio-informatics person.

- Stuart Johnson, a TACC staff person is working at UT-Southwestern to get ten applications running on TACC's Lonestar cluster (10M processor hours/year). Stuart is working with 10 different labs – see presentation for a description of the applications associated with these labs. These applications range from refinement of structures from NMR data to protein structure comparison to electron microscopy image processing. General application characteristics include the following: large range of data movement (e.g., O(10 TB)/year, 1 TB/experiment per week), at least O(10K) CPU hours (some O(100K) CPU hours), community codes used by numerous research groups, and need for scaling up (or a shared resource needs to be adapted). At UTSW there are workstations and clusters; these are connected to TACC, where larger systems will run the applications. One problem is overloading the file system. Challenges include: clusters are not easy to manage and use, and shifting data around is not easy. Platform Computing's LSF is used for scheduling. People at UTSW know exactly what needs to be done but the hurdle is getting it done on a cluster (lots of nodes and a shared machine/distributed memory model (many tricks here: handle data movement, catch errors, etc). Jae Yu pointed out that this scheme is very similar to HEP's application and the update of the software scheme. The same difficulties as HEP in implementing and upgrading software on a local site is shared by TACC and UTSW. Some work that Stuart has done regarding these apps: changed on-demand code style to batch scheduling; wrote master/slave C/MPI code to manage serial code runs, data movement, local/global storage, logging and RTE copy to /tmp. Scripting languages are used a lot: mpi_ruby to implement parallelism – install, test, code mpi/ruby example. Ruby essentially is a Japanese Python – see Ruby/MPI/mpi_ruby/Fortran/C; added LSF (batch scheduler used by TACC) compatibility to code; installed Matlab compiler – deployed RTE and compiled apps (large scale systems biology – image processing, data reduction); adapted code for the largest part of the problem from another research group (finding modularity in large (100 KN/E) networks of interactions -- 10-20 terabytes of data per year – O(10,000) CPU hours for first end-to-end data analysis.
- A discussion was started concerning MatLab, SyLab, Mathematica (3D) and IDL.
- Prashant Mishra from the Ranganathan lab at UTSW presented his research on the fruit fly. Prashant gave a very interesting talk that demonstrated progression from experimental data to a model to insights that can be used to reduce experimental space and refine model. The eye of a fruit fly is what is being studied – this is a good model for understanding human diseases and genetics – it is powerful in identifying all the protein components, although there is a lot more information to be discovered. How do you model functionality of the cells? Oscillator built into system. Mathematical model of cell: 7 different kinetic equations and 23 parameters. Model explains some of the observed behavior, in particular, non-dampening of response. Must run simulations to find solution to 7 non-linear equations; small numbers of active species (molecules), therefore, stochastic simulations; no spatial dimensions; precise nature of the interactions is often unknown. Gillespie algorithm used – state variable at time t, find out state at another future time. Modeled as a step function.

Determine time steps based on rates – rates high => small time step; rates low => larger time step. 5000 time steps needed to see a quantum bump – 100 simulations on desktop – 1000 simulations needed. Shows that it models biology then start analyzing mutations. Models provide interesting intuitions. Allows hypotheses re: molecular reasons for observations/results. The modeling – experiments are rate limiting – can the simulations identify parameters of interest / dependency among parameters – can reduce time to discovery. Direct experiments in an intelligent way.

- Richard Scheuermann from the department of Pathology at UTSW described two large, very comprehensive Bioinformatics and Computational Biology projects. Bioinformatics Resource Center (BRC) for Biodefense and Emerging/Re-emerging Infectious Disease (\$17M) NIH and Bioinformatics Integration Support Contract (BISC) (\$28M). BRC focuses on giardia lamblia and influenza virus (also Francisella tularensis, mycobacterium tuberculosis, microsporidia, and ricinus communis). BISC involves the development of an immunology database (for diagnostics, therapeutics, and drugs) and analysis portal. In addition, Richard and his colleagues are starting in the area of data analysis and algorithm development – he also discussed their work in this area, i.e., parallel MoNet, which decomposes biological network systems into modules of various complexity – MoNet utilizes topological analysis and graph theory to distinguish edges between and within modules – the findings were validated with biological data in a gene ontology (frequency of clustering indicated functionality) – a completely independent database (e.g., a yeast protein interaction network was experimentally derived and then validated). The error rate for this methodology, although not known, is much higher than that associated with genome sequencing. There are lots of false positives and some false negatives. The time complexity of MoNet analysis is large – e.g., analysis of B lymphocyte genetic network; it is 50 days for about 65,000 edges. This complexity is being reduced by parallelizing the calculation of all shortest paths in G-N algorithm to obtain edge betweenness. The computation problem in a nutshell is that experiments can now be run in a high throughput fashion => lots of data. Efficient and effective methods are needed to analyze that data. Collaborations with researchers at UT-Southwestern are possible -- bring expertise to bear on medical problems. For example, SMU statistical faculty/students are working on some of their problems.

- K. Kirksey, the VP for Information Resources at UTSW Kirk presented an overview of computational, networking, and storage facilities, including related issues, e.g., redundancy, security, and customization to provide for various different needs.

UTSW:

- is a member of Internet2 and NLR,
- is now using a commercial network provider but plan to join LEARN,
- is developing a voice over IP network,
- just completed a facility to allow patients to electronically access medical records,
- Information Resources has 291 employees and a \$35M budget – note that UT-Southwestern includes several hospitals plus an academic medical center, and
- provides access grid in four rooms.

With respect to security, Kirk provided UT System audit findings, i.e., issues that had to be addressed: backups not required, firewall/virus policies have to be enforced, lab computers must be protected, patch management must be improved, password

policies must be enforced, security logs must be reviewed, users must be made aware, perimeter controls are needed, and wireless transmission must be understood. UT-Southwestern has an IR Contract Program (disseminates relevant information), Systems Administrator Program (in collaboration with UTSA), and a Disaster Recovery Program, and purchased Critical Watch Vulnerability Management Program (for internal security violation detection). With respect to disaster recovery, Sunguard is used (\$150,000 just to subscribe). The UT system is investigating a system that takes advantage of collaborations among data centers, i.e., data center consolidation.

- There was some discussion on TIGRE concerning the use of DOE CA and the use of Globus as the common framework. It has been agreed that the steering committee will make a decision concerning these matters in 2 weeks time scale.
- A possible date for the next in-person meeting at UTHSA was discussed (in January 2006). Pros and cons about the meeting at UTSW (meeting date should be chosen so that preparation/communication is not during a time when host is on travel; dinner the night before should likely not be planned).
- The education thrust was discussed informally to determine interest. There seems to be interest in remotely taught courses (need to work out the details/issues) and course modules that could be used to construct one-semester long courses and/or mini-courses and workshops. There may be funding available for such an effort, e.g., from NIH or NSF (CCLI program). Leadership is needed; the membership needs to vote on leadership; Pat Teller indicated that she would be investigating institutional support if she were to lead the thrust. University of Houston may have a person who is interested in taking the lead.
- There was some discussion on shared courses for students' credit across the university boundaries using distance education.
- An action item identified at the meeting was for Phil Smith to contact Josten Ma to allow HiPCAT people to be adequately authorized to post documents on the HiPCAT web page.
- Cluster information/questions/thoughts: LSF (best scheduler even with its problems) -
- LSF is a better scheduler than PBS but it has a price tag – LSF can be used on top of ROCKS – PBS pro is free for academic use (OpenPBS is part of ROCKS). Problems with scaling – LSF assumes all nodes are working – scripts needed – memory bandwidth issues – file access issues. Power, UPS and cooling needs of clusters are significant.

Institutional updates

- UT
 - TACC's computer center expansion is done including enhanced air conditioning. The new building is 18 months away.
 - TACC is upgrading their 512-node Lonestar cluster.
 - Received IBM second year funding for UTGrid.
 - 2 grid technology people from IBM are working onsite at TACC and will be working on TIGRE, Teragrid and campus cyber-infrastructure.
- Rice

- Brought up a 128-processor Cray cluster with FPGA's that can be used to control network or to act as co-processors.
- UTHSCSA
 - Inherited an Anthalon cluster (that was built by UTHSCSA). The cluster will be used by researchers in crystallography and NMR for structure modeling.
 - Purchased storage facility (8 TB for backup and storage – coupled with another backup and storage facility).
 - Will be connected to LEARN shortly.
 - Working with TACC on TIGRE use.
 - VPN bridge between clusters at UTSA and UTHSCSA (MPICH 2.5 – known problems with latency) tested. Tests show that latency across communication link is negligible.
 - Establishment for Bioinformatics core facilities – steering committee members needed for Bioinformatics thrust -- \$25M budget.
- UTEP
 - 64-node, 128-processor Opteron cluster ordered by Andre Kerstens.
 - A large Dell cluster to be used in the GEON project will soon be ordered by Steve Riter (VP of IT).
 - 24-processor IBM p590 is up and running and users are being migrated to this machine from the p690.
 - 4-processor IBM p550 is up and running and used for virtualization research.
 - The IBM p690 has been extended with a fast Fiber Channel external disk array.
 - The 36 dual-core Opteron processor Cray XD1 cluster is up and running and currently being benchmarked.
 - Pat Teller's research group has two papers at OSIHPA.
 - A \$1.2M grant with the Scripps Research Institute & UC Berkeley (PI Michela Taufer; Co-PI's Pat Teller, Martine Ceberio) has been awarded by NSF. The objective is to build a protein-ligand docking system using desktop machines on the Internet. The project name is DAPLDS.
 - Computational studies of Giardia Lamblia are being initiated.
 - Andre Kerstens, Michela Taufer and Pat Teller are preparing the Supercomputing 2005 Research Exhibit for UTEP.
 - Pat Teller has been promoted to full professor (congrats!).
- TTU
 - Bringing up a 256-processor Dell system with 4 TB storage (1/2 price of other vendors like HP) with Rocks / Infiniband roll – problems with switches GigE and Myrinet networks; setting up cluster so that researchers can buy 32-processor units to add on (with GigE); really parallel jobs will use Infiniband partition.
- Given the time limitations and the imminent landfall of the Hurricane Rita, many participants from Houston area had to depart early. The tour of UTSW's HPC Center was cancelled and made the meeting shorter than usual. Phil Smith adjourned the meeting.