

CASC Meeting Washington, DC

Tuesday, March 22, 2005

SURA Grid Working Group Reports.

Ed Seidel led this discussion. CASC attended the final half-day of the Southeastern Universities Research Association (SURA) IT Committee meeting. Reports included those on infrastructure (Art Vandenberg, Mary Fran Yafchak); Applications (Ed Seidel, Sara Graves); Awareness (Steve Moore, Mary Fran Yafchak); and Funding (Gary Crane, Sue Fratkin). Charts are on the SURA website.

SURA TeraGrid Participants Update.

Jay Boisseau gave an update on how SURA members are participating in the NSF TeraGrid project. Charts are on the SURA website.

SCOOP Grid Activities.

Sara Graves made a presentation about SURA's SCOOP (SURA Coastal Ocean Observing Programs) project. This is both economically and politically important, because 75% of the damage to U.S. property from tropical storms and hurricanes takes place along the coastline of the SURA states. Charts are on the SURA website.

International Connectivity.

Don Riley gave a presentation showing all the various emerging networks around the globe and how they hope to be tied together. Charts are on the SURA website.

Meeting Summary/Next Steps.

Dave Lambert led a brief discussion summarizing the meeting and praising the group for excellent progress on all fronts.

Dinner Meeting that evening was held at the Old Ebbitt Grill with Dr. Tony Tether, Director, DARPA, as our guest.

Wednesday, March 23, 2005

Business Meeting

Frank Williams, CASC Chair

Frank raised two issues to the group. The first one was a proposal by Craig Stewart that CASC sponsor several Traditionally Underrepresented Groups (TUG) attendees from Minority Serving Institutions (MSIs), Historically Black Colleges and Universities (HBCUs), or tribal colleges, at the annual SC convention.

Many members liked the suggestion of pairing attendees with CASC members and CASC having a more active role at the convention. Suggestions included: extend invitations to underserved groups and at those institutions that serve them; do a pilot program this year and assess its success; and invite faculty and staff to apply as well. A motion was made and passed to implement the proposal for SC05. Craig will take the lead on implementation.

Frank also presented the second issue: that CASC consider taking on sponsorship of conferences such as the Comprehensive Supercomputing and Computational Science User Forum or being the not-for-profit front organization for the HPC Use Forum. This idea was from the High Performance Computing Program Office (HPCMPO) and later clarified by Cray Henry in his presentation. Comments ranged from “it couldn’t hurt for us to be involved” to “do not have enough staff and conference planning is a major effort”. CASC is considering this but did not take action.

Lastly, the group discussed questions that they would like to ask Dr. Marburger of OSTP at their upcoming meeting.

Financial Report. Dick Pritchard, OSC, Secretary/Treasurer CASC, reported on financial accounts as of March 21, 2005, as follows:

Total in the CASC account	\$138,281
---------------------------	-----------

We have one outstanding bill due for \$16,240. One member is past due on membership dues for the 2004-2005 year; they have been contacted and arrangements have been made for payment. The annual revenue will be enough to cover current and projected expenses for the remainder of the year. (For additional details, contact Dr. Pritchard, rhp@osc.edu)

Computer Research Association (CRA) Meeting Report.

Craig Stewart

The two issues discussed were funding for computer research and workforce development for computer science.

Funding rates for computer research have gone from bad to abysmal. Numerous good, meritorious projects have not been funded and those projects that are fortunate enough to be funded are not funded at high enough level. The combination of less funding, funding

constraints, increased faculty looking for funding, and more proposals being submitted, there is no good news. Very little talk at the meeting was dedicated to solutions to the funding challenges.

The second issue, workforce development, also has its challenges. Data shows that there is declining interest in all levels in computer science. Some proposed that is because the dot.com bubble burst. CRA has an action plan: develop an ad campaign to raise the interest in computer science and develop curriculum standards for computer science minors.

Sue talked briefly about visibility. AAAS has begun an ad campaign on science. She will be working with the AAAS and meeting with the National Press Building's science and technology to discuss how to bring attention to the work CASC members are doing. She also now has a contact at NPR who is among their science reporting group.

Speakers

Dr. Dan Hitchcock, Senior Technical Advisor, DOE, Office of Science
ASC Program

Where is DOE going with HPC? DOE cares about the science, not necessarily about the hardware. It is interested in solving the same problem in less time and more complicated problems in the same time.

Oak Ridge has two new machines. Not more than 10 teams per year can use the machines. DOE is looking to buy the next set of machines; must do evaluations now to see what is out there. Hitchcock talked about a few projects: applied mathematics and computational chemistry. He briefly talked about the Insight program and its focus on computations that are too big to do anywhere.

Discussion focused on the partnerships between DARPA, NASA, DOD for HPC. Hitchcock pointed out that more processors don't necessarily mean better science. The challenge seems to be in software and user codes and of course grid computing. There was some discussion of "end stations." Hitchcock mentioned programs with HBCU to prepare the next generation of scientists and the Computational Science Graduate Fellowship. The group raised the issue of foreign nationals and cybersecurity. Hitchcock explained the group the new regulations that the use of supercomputers is now considered an export and usage is now subject to export controls. [This is very problematic] Dan talked about a government-wide effort to issue personal identity cards in order to access a government machine. Basically, homeland security/cybersecurity trumps everything.
Presentations are on the CASC website.

PITAC ASC Report

Dr. Miner gave a talk on the future of computational science from the NCO perspective. Reed, after being congratulated for approval of UNC as a new CASC member, gave a summary of the latest PITAC activities. *Presentations are on the CASC website.*

William “Buff” Miner, PhD
Technical Coordinator, NCO

NITRD focuses on interagency issues. By leveraging resources, agencies can take advantage of other programs. Coordination groups keep abreast of what all the agencies are doing; discuss what the next roadblocks are; and what new programs should be initiated. Twelve agencies are involved and have a total budget of \$2B. Currently, a draft definition of computational science is being developed. It is very broad and inclusive; includes the social sciences and humanities.

Buff then spoke about PITAC. PITAC has had a historic role in computer science. Since computer science is all pervasive, we now need new tools and software to go forward. – We have enjoyed progress but are we progressing at the right rate. Buff suggested that perhaps CS is being held back. He cited the LAX Report from 1982 that recommended increased access to good networks and improved computation resources. He too mentioned the educational pipeline is slowing down.

High End Computing Revitalization Task Force believes that peak performance is no longer as important as time to solution. We need better code, algorithms, and models. Storage and bandwidth are huge problems. The private sector will fund HPC research. We need better ways to develop software, so there is a proposal for creating software centers. Buff also spoke about “leadership systems” – systems large enough for cutting edge research but are too expensive for one agency to pay for. A competitive proposal process would be used to bid for access.

Industry needs to understand how HPC is essential for their business. Often, CEOs don’t understand what the possibilities are or how HPC can greatly increase their company’s bottom line.

For more resources, go to: www.nitrd.gov

Daniel Reed, University of North Carolina, Chapel Hill
PITAC, Subcommittee on Computational Science
Dan_reed@unc.edu

PITAC’s Subcommittee on Computational Science is almost at the end of the process. At SC ’04, the team gathered more input and a formal call for information from agencies were made. The roll-out of the draft report on computational science will be on April 14. [You can watch via the web.]

The definition for computational science recognizes the needs of social sciences and the humanities. It captures how broad CS has become. We discovered how CS seems to be everyone's second priority, but very few made it their first.

A few highlights:

- The ecosystem is unbalanced.
- Software needs are unmet.
- We need continuity in implementing a long term plan.
- No roadmap for future operational guidelines exists.
- To create a long term plan, we need to prioritize funding and develop a Federal plan.
- Sustaining infrastructure is important.
- Large scale data management is a challenge.
- DARPA is ending their CS program and no agency has yet picked up the slack.
- We need to focus on building a supercomputer for next generation applications.
- We should create software sustainability centers.
- Training is not integrated and is inadequate. No interdisciplinary degrees exist – need to change university policies and curriculum.
- Greater interagency coordination is needed. Use a real roadmap that drives agency budgets for coordination. (HR 28)
- There is a disconnect between commercial needs and science. Investment is not coming from industry.
- Barriers: software, access, difficulty

Q&A:

DARPA:

HPCS does exist in DARPA thanks to Bob Graybill, but should such a program be dependent on a single individual? Challenge is attracting CS people to DC and working at DARPA. We need to improve incentives to come to DC.

Cybersecurity:

Cybersecurity has a huge impact on CS. Cybersecurity is a cross-agency initiative. PITAC has a technical focus – looking at data issues.

Funding:

Real issues are political and societal. Need to examine how agencies work together; need to determine who their core constituency is. A systemic change is process. The roadmap is the actual budget guide that is used by OMB. We need interagency cooperation – and traction at those individual agencies.

PITAC Report:

The report shows that as access to large scale data sets increases, the need for high speed networking becomes apparent. Antidotes are the most important/have the biggest impact when making an argument. We all need to describe where HPC has made the greatest advances and enhanced the US economic advantage.

Eric Jakobssen, NIH Biomedical Information Science and Technology Initiative
Extramural Computational Biology at NIH

Computation and IT are needed for bio-research; however, many challenges exist such as lack of interoperability and difficulty in using high-end machines. There is very little support for computational biology – too few personnel; no significant advances in education. We need better supply of trained people.

Computation is at the leading edge of biomedical science. Future progress is in cataloguing and analyzing individual genome-based variations to permit customized diagnosis and drug therapy. Infectious disease modeling will need computational advances. New tools are needed for biomedical research. Right now, there is a lack of software and algorithms.

In 2004, the first set of National Centers for Biomedical Computing established. In 2005, a second set of centers will be established. Small science will work hand-in-hand with big science. The next stage will be a coordinating group of NIH staff to monitor computational activities. There is no strategic plan or vision for biomedical computation. NIH wants to work with computer scientists and help identify areas of computer science that are important for biomed research and should be supported by NIH. Input is welcomed.

NIH is a mission driven agency but NIH is increasingly supporting CS and IT.
Resources: CRISP database; BISTI homepage; CSR (study sector information).

NIH uses HPC but doesn't necessarily pay for them. The big limitation is in the software, not hardware. NIH rather invest in software so researchers can better use equipment. Infectious disease modeling on HPC is coming soon.

Success is shown that computational science is now critical in biomedical research; however, now it is critical for CS to continue in order for future biomedical advances.

Cybersecurity and Supercomputing Centers

Charts are on the CASC website.

Ginger Caldwell, NCAR
Cybersecurity Conference Report

The Cybersecurity Conference was in response to the “spring incident” of last year – what is now referred to as “FBI Case 216”. A variety of individuals attended the meeting, with university system administrators being the majority. The main conclusion was that universities and research facilities need to develop or enhance communications with each other. Trust relations need to be established, alternatives for email need to be chosen, and secure networking must be developed. A point person at each facility needs to be designated for such incidents.

Breakout groups identified the challenges and drafted suggested solutions:

- Incident Response
 - Communication between centers is poor.
 - Discovered that many centers did not apply necessary software patches. Security starts at home! Centers should assess their own vulnerability.
 - Create a response plan and set up security procedures.
 - Procedures for peer group notification and alerting users must be developed.
- Training and Education
 - Passwords are weak.
 - Understanding the risks and protections is poor. [This includes system administrators!]
 - Develop a security plan that includes education, enforcement, and collaboration.
 - Encourage NSF to provide security training for system administrators and users.
- Security Planning
 - Security cannot be reactive.
 - Suggest NSF impose security requirements for grants
 - Examine better password alternatives

George Strawn, NSF CIO

NSF & IT Security

Strawn spoke on cybersecurity program at NSF.

NSF does not want to be the boss of awardees and users – they don't use the federal rules [FISMA], nor do they want to.

NSF wants to help with security. NSF is focusing on large scale facility IT security. 10% of the budget at NSF is used for internal security. A security working group has been established at the NSF and is tasked with developing plan of action; review & assessment tools, and training curriculum. Facility security is now a NSF security working group. NSF now uses standards for connectivity; firewall architecture; patch management; and intrusion detection. NSF has two networks: external and internal.

Strawn believes that facilities are not using the tools currently available; of course we do need next generation security. Yet, a culture change is necessary – all levels need to learn IT security. The key in IT security is keeping intruders out. Facilities must develop proactive security measures and evaluate risks and vulnerabilities. Centers need to create a web of trust within the community and users. User authentication and communication are important.

Rodney Peterson, Task Force Coordinator, Educause

Educause/Internet 2 Security Project

The project evaluates higher education institutions and cybersecurity. Institutions should focus on training students in cybersecurity in effort to develop future security leaders. More research in cybersecurity and effective tools are needed. Colleges and universities must secure their networks. Strong scientific computation is our goal. In order to attain that goal, intellectual property policies should be developed for academics and integrity of data must be preserved. Continual security breaches may cause projects to be shut down.

Strategic goals for project include:

- Education awareness & training -- Students need to prepare to become IT professionals. Training is critical – perhaps at an annual security professional conference.
- Policies & procedures – legal compliance by universities and students
- Security Architecture & tools – Develop a best practices guide for universities. Most institutions do not use the current tools available.
- Organization & information sharing – New security programs should be offered at conferences. Many opportunities to address security issues.

For more information: rpeter@educause.edu or 202-331-5368

Congressional Updates

Sam Rankin talked about funding for NSF. Kate Bannan talked about funding for DOE Office of Science. Peter Rooney, Elizabeth Grossman and Jim Wilson gave an update on the House Science Committee and the HPC Revitalization Act.

Sam Rankin, Chair, Coalition for National Science Funding

CNSF mission is to raise the visibility of the science budget on Capitol Hill. Due to the diverse membership, we work at the highest level to secure funding for the NSF budget as a whole.

In the mid-1990's, NSF had a spurt of increased funding until about 2003. In 2002, Congress passed a law to double the NSF budget over the next 5 years. Despite that law, NSF funding has continued to decline.

Kate Bannan, Chair, Energy Science Coalition

The Energy Science Coalition has about 150 members from universities, research centers, and national labs. Our goal is to secure funding for science at DOE.

Last year, DOE Office of Science had a 3% increase; yet in FY06, it has been cut by 4%. Currently, we are asking Congress for an additional \$250M (7%) increase. You can help – call your home Senator and ask them to sign the “Dear Colleague” letter. Last year, 55 members signed on to the letter. We especially need Republicans to sign on. In the

House, Rep Judy Biggert will be starting a Dear Colleague as well. 97 House members signed last year.

Get to know your home congressional staff members and your district staffers. Invite them to tour your center.

Kay Bailey Hutchinson is now chairing the Senate Science subcommittee. She and the other committee members need to know what supercomputers are and do. Many new faces in Congress and they need to learn what terrific projects you are doing.

Peter Rooney, House Science Committee

The HPC bill passed last week out of committee. We expect it to pass the House and hope to get it to the Senate floor before Fall. NASA poses a huge challenge this year. The President has a new vision for NASA and the Senate subcommittee will be working with him on that.

NSF: No one really thought NSF would be cut so drastically. The Science budget is NOT healthy this year. There is competition for funds due to the new vision for NASA. There are still many issues, other than funding, at NSF. There is a big hole in supercomputing software and hardware. The Committee wants to influence them in the right way, want them to go forward in research. We are concerned with the refusal rate for proposals.

Jim Wilson, Minority Staff Director, House Science Committee

We are trying to pass the HPC bill again this year. We need to focus on the interagency process and develop a roadmap for the long term. The committee is disappointed with NSF's planning and implementation for HECRTF. NSF has now established an internal committee that includes all the other directorates as well as CISE to produce a plan. Additionally, a higher level committee on the Director level will implement this plan. A third committee – the users committee – will also have a role in the plan. Make sure CASC is involved. The committee has a new subcommittee ranking member – Darlene Hooley from Oregon. You should get to know her.

Elizabeth Grossman, House Science Committee

The Research Subcommittee has a new chair – Rep Bob Inglis of South Carolina. He is very interested in NSF and understands that investment in science is important.

The Appropriations Committee has been reorganized. A new Science subcommittee has been established – Rep Wolf from Virginia is now the subcommittee chair and Rep Mollohan of West Virginia is the ranking member. Now competition will be between NASA and NSF instead of NSF and Veterans Affairs.

PITAC is scheduled to be renewed this year; there is a 2 year sunset provision that can be used continually. You should raise PITAC with Dr Marburger.

The number of proposals in CISE increased dramatically. Cybersecurity activity has increased over the past 3 years. DARPA has gotten out of funding unclassified research so the burden shifts to NSF. More proposals, same amount of funds results in the lowest success rate. Supercomputers are a thorn to NSF right now. They don't like long term commitments. There is a view that computer science community doesn't directly benefit from supercomputers. The tension will never go away, but NSF must continue to support it. In regards to ITR, the bottom line remains the same, but those programs have moved from CISE to other directorates.

The NSF Director Committee will decide how money is divided up. They may revisit how this is done in the future but right now looking at interdisciplinary research for cybersecurity.

Grossman spoke a bit about a "leadership machine" versus 43 centers. The Committee view is that NSF needs to plan balanced research and maintain current investment in supercomputers. NSF has signaled that they want to back out of high-end computing. They are no longer thinking of upgrading 2-3 years. The supercomputer community should help decide what technology should be invested and activities to undertake. There is a big shift in users – most users of supercomputers don't have experience in computational research. Efforts need to be made to train these SC users. A long term plan should be developed – something other than funding 2 ½ machines. The problem with one leadership class systems (such as Oak Ridge) is that NSF users will be involved in DOE politics. For example, NERSC is not as widely used as hoped. NSF is the only agency with a broad scientific agenda.

Sue added a few comments as well. All the forthcoming reports buttress what we are all saying. There seems to be hardship at NSF/Cyberinfrastructure – fewer staff, low morale, director departing this year, difficulty in placing IPAs... NSF will have a hard time when these reports come out. But NSF finally has plans underway for cyberinfrastructure. There is an internal advisory group where the directorates are soliciting opinions and comments from their users. CASC has not been involved in this process.

Dinner Meeting that evening was held at the Tuscana West. Dr. Arden Bement, NSF Director, was our guest. He brought Dr. Sangtae Kim, Director, SCI, Dr. Guy Almes, IPA

Dr. Ardent Bement, Director, NSF

Dr. Bement gave some remarks at the CASC dinner. He spoke about the three advisory committees that have been set up to coordinate efforts in cyberinfrastructure. The first advisory group is the external committee – the cyberinfrastructure user advisory group. It will be a committee of grantee institutions involved in cyberinfrastructure. They will get input from the computational community. The second advisory group is the Directorate Working Group. Each Directorate will get input from their staff on cyberinfrastructure

activities. The last advisory group will be the Director' Advisory Committee on Cyberinfrastructure. Bement heads this committee and it involves all the Directorate level directors.

Bement reminded the group that NSF is vastly different from DOE – NSF investment is not geared to building labs, but funding research.

CASC Meeting Washington, DC

Thursday, March 24, 2005

Frank Williams, CASC Chair
Business Meeting

Frank began this morning's meeting with finalizing questions for the CASC meeting with Dr. Marburger of OSTP that afternoon. Issues suggested: funding; vision for high-end computing; investment in software and new tools; cross agency coordination.

Bob Graybill, DARPA High Productivity Computing Systems
Update on HPC

Graybill gave an overview of where the HPCS program is going, including an announcement of HECURA (High-End Computing University Research Activity), an item of interest to the CASC membership. *Preentations are on the website.*

HPCS program looks at the entire stack – including program language. Currently, HPCS is midway Phase 2 with awards given to IBM, Sun and Cray. Phase 1 cost \$18.1M and Phase 2 cost \$229 M. This includes vendor investment (about \$50M), not just research. Goals for the program are to sustain 1 petaflop/sec. Challenges are the framework, bandwidth, and interoperability. A notable accomplishment of Phase 2 is that all government partners were very involved in the reviews.

DARPA has a large role for universities -- \$10-\$12M per year is spent on these centers. They are examining operating systems and language. This is a joint effort between DOE, NSF and DARPA. These university projects have a direct impact on industry. They are working with IBM and Cray as a team.

Graybill talked about the vendor innovations. Vertical integration is important. Sun is working on a flat memory system. All vendors are working on separate programming languages. HPC market is somewhat splintered – hopefully this project will be able to bridge those challenges and create product that has broader appeal and is easier to program. Smaller companies need help to improve software – they have little HPC talent.

Major efforts in Phase 2: productivity (software cost; time to solution); benchmarking (8 HPC challenge benchmarks); and porting codes. Graybill talked about a project at UCSB that used a work flow model that students used to evaluate their projects. They spent the most time debugging the program. This workflow model shows a way to improve all over productivity by examining where you spend the most time and address those challenges.

Graybill then spoke about Phase 3. A new area will be the development of high productivity language system. Common threads between the Sun, IBM and Cray approaches will be used to create a universal language. The goal is to extend languages

beyond these 3 vendors. DARPA doesn't just want to create software HPC but software that can migrate platforms. We hope this is a major impact on companies and how they are organized. All mission partners are at the table with vendors.

DARPA will fund industry milestones and is working with mission partners. The entire HPC effort is approaching \$1B. The tradeoff in machines will be cost versus performance. There is a price point for each feature, so you will need to know what you need. Other related programs: new architectures, information processing, and hybrid cognitive techniques.

**Note: Bob leaves his post at DARPA in June.

Cray Henry, DOD HPCMP

High Performance Computing Modernization Program

DOD's mission is to develop and deploy capabilities to help the war fighter today. 561 difference projects that are in 10 technology areas. HPCMP is directly under the Secretary of Defense – integrated and united across DOD/ across all service branches. Some examples are: aerodynamics software, predicting dust storms, terrorist attack modeling, and unmanned aerial vehicles. Program areas: computing center, networking capacity, software support, and resource allocation. *Charts are on the website.*

Resource management is important. DOD has a database from 10 years of projects. Allocations are determined by capability and capacity. There are challenge projects across DOD that get premier treatment. Then there are individual allocations that the service branches determine themselves. Capability allocation project works with an already existing project. For example, reallocation of time at centers for projects is common. DOD also gives out distributive allocations to dedicated projects. These are often smaller projects and receive a one-time investment. Most centers are retired after 4 years.

\$40+ million in funding is given to large centers for capital investments. Security is very important. DOD uses filter devices, firewalls, traffic monitoring, and validation centers for software protection. Software application support is critical – must invest in human capital as well in order to have trained people. Five HPC software application institutes exist that are tasked with developing applications that can solve broader problems.

Access rules. DOD is changing some of its rules. Previously, graduate students had to go through time consuming background checks. Now, open research facilities exist. (Such as the two in Alaska).

Henry closed his presentation with discussion on the Comprehensive Supercomputing and Computational Science User Forum. DOD determined that the forum must be part of a nonprofit entity and is looking for a new home for the forum. The suggestion is that CASC takes it over. This forum puts on 2 or 3 meeting on a year for 150 or so people. There is a mix of participants – industry, academic and government.

Suzy Tichenor, Council on Competitiveness
High Performance Computing Initiative

Suzy gave the latest on the CoC HPC project. *Presentations are on the web.*

What are the drivers of innovation? How do you plan for innovation -- Funding?
Resources? Education?

The Council is developing an innovation strategy. They held a summit in December to examine what is the HPC community doing today. CASC was invited to join the President's Council. The report can be found at: www.compete.org

Tichenor is working on an applications timeline – a history of application breakthroughs in HPC. She is also looking for thoughts of what comes next – what new breakthroughs are on the horizon in HPC.

The group discussed in depth about drivers of innovation and HPC breakthroughs. Many feel that industry is not investing as it should and often CEO's don't understand the value of HPC. Also, industry seems to be hesitant to share its breakthroughs/successes due to proprietary issues.

Meeting with OSTP Director, Dr. Jack Marburger, White House Conference Center. CASC had a cordial meeting with Dr. Marburger and asked him a series of questions that primarily expressed our concerns about the levels of investment in HPC research and the distribution of those investments within the federal FY06 budget.

Next Meeting(s): June 1-2, 2005, in Fairbanks, Alaska with ARSC/University of Alaska as the host institution.