Computing Community Consortium Reverse Site Visit 2/16/2010 Review Panel Report

Review Panelists

Rick Adrion, University of Massachusetts Amherst
Al Aho, Columbia University [chair]
Fran Berman, Rennselaer Polytechnic Institute
Ken Kraemer, University of California, Irvine
Glenn Ricart, National LambdaRail
Burton Smith, Microsoft Corporation
Telle Whitney, Anita Borg Institute for Women and Technology

Introduction

The Computing Community Consortium (CCC) Reverse Site Visit Panel reviewed presentations from CCC Council members Andrew Bernat, Susan Graham, Anita Jones, Ed Lazowska, and Fred Schneider on February 16, 2010 at NSF in Arlington, VA. The purpose of this review was to determine how well the CCC is meeting its stated objectives. This report contains the findings of the panel.

CCC Goals

The original creation of the CCC stemmed from a perceived need for the computer science community to "articulate and coalesce around exciting research visions in computer science – research visions that would galvanize the public, policymakers, researchers and students" (from the NSF solicitation).
In its proposal to the NSF, the CCC more or less adopted the goal of articulating future research needs, but also broadened it as follows (from the CRA proposal and CCC strategic plan):

1) Bring the computing research community together to discuss, prioritize and envision their future research needs and thrusts.
2) Communicate these challenges, needs and thrusts to the broader national community.
3) Create within the computing research community more audacious thinking.
4) See the ideas developed in (1) and (3) turn into funded research programs and/or instruments.
5) Increase the excitement within computing research and use that excitement to attract students of both genders and all ethnic groups into computing research careers.
6) Establish the Computing Community Consortium as a widely accepted catalyst and voice for the computing research community.
7) Inculcate values of leadership and service in the computing research community – by example, by inclusion and by mentoring.

The panel felt these are all important and useful goals. An implication of this broadening, however, is that the CCC may have diluted its original mission. It has shifted somewhat from its original mission of defining grand challenges in computing to helping others to define challenges within various subcommunities of CS. To date this has been a slow process and has not yet resulted in many vision statements or their diffusion to agencies or policymakers. However, there are a few good examples of success, especially in the case of robotics research.

On the positive side, the CCC has expanded its role to responding to requests for guidance and to opportunities that come from federal agencies and the new administration. We perceive the CCC as highly effective in doing this. Moreover, there is evidence that the agencies regard the CCC as useful and important in setting their own agendas and national priorities.

For the future, the CCC needs to pull together the grand challenges that emerge from the various visioning groups and to issue a “blue book” on the grand challenges in computer science. Moreover, these grand challenges should be widely communicated within the CS community and to government agencies and policymakers via multiple means. The CCC should also be proactive in communicating directly with policymakers about CS grand challenges.

CCC Strategies

Proactivity, opportunism, and agility are important to let the CCC respond well to evolving realities. Community visioning is fundamental, should be continued, and will undoubtedly always have mixed outcomes. Workshops should be used for convening constituencies and exposing new ideas as well and for forming prototype agendas. Senior leadership of the field must be recruited since they can play a key role in community visioning and in selling the result. The health and balance of the field should be part of the CCC's mission – arguably, as part of its charge from NSF. CCC needs to build bridges to the research leadership of other fields, both to help define credible interdisciplinary research agendas and to help motivate other funding agencies. Community visibility targets should be ranked roughly as follows:

1) Computer research leadership
2) Research leadership in other disciplines
3) Computer research people
4) Research people in other disciplines
5) Computer research agencies
6) Other research agencies

Implementation Plan

In their implementation plan (dated 8/11/09), the CCC leadership laid out a strategy for connecting their goals and outcomes with measures for their success.

They expect to interact with agencies in two ways: increasing the high-level understanding of the role computing plays in the national priorities (Outcome 1) and supporting the development of new programs (Outcome 3). These two outcomes are connected with almost all of the CCC goals. They expect to measure success based on:

- "The number and quality of the contacts between the CCC and the relevant agencies, the quality of the roadmaps and other reports generated from both our visioning workshops and our targeted efforts, feedback from the agencies as to their agreement in the relevance of computing research their missions."

- "Agency interest and willingness to fund the proposed research."

These are admirable goals and hoped-for outcomes. The CCC should try to recruit "advocates" from other disciplines, perhaps limited to current initiatives, to help promote interdisciplinary agendas. Diversifying funding beyond NSF will also help convince those agencies funding the CCC to implement programs in reaction to CCC recommendations.

As a basis for the programs, to help energize the community (towards "more audacious thinking"), and to promote understanding of computing, the CCC strategy is primarily to solicit, develop, support and promote "visioning," typically through workshops, that would lead to roadmaps. These have clearly led to some successes and some pending, but likely successes. There are initiatives still at risk and one failure. They need to measure not only participation in these activities and the number that led to actionable or influential outcomes, but also develop measures of the key factors for success and failure that can be used in their process of solicitation and negotiation for program initiatives.

While "visioning" is their primary strategy, questions arose. *How does one set priorities for selecting a set of activities?* They cannot cover everything and while they suggest that they will not overlap agency (e.g., ISAT and CISE AC) agendas and program manager efforts and initiative, they will still need to make choices. Here the leadership of the CCC Council, officers and staff is critical. The Council adequately addressed concerns about the bottom-up appearance of their visioning strategy. *Beyond diversifying and prioritizing the topics, how can they assure that they reach the appropriate communities to carry out visioning?* They need, for interdisciplinary areas, to bring in representatives from all of the communities, not just computer science. Again, advocates from other disciplines

could be used as advisors in broadening interdisciplinary representation. *How are they going to extend their reach to insure broad disciplinary and interdisciplinary participation?* They should contact other communities to solicit input and proposals and to disseminate outcomes: the professional societies, the Anita Borg Institute, the Tapia conferences, the BPC alliances, etc.

Two other goals (attracting "students of both genders and all ethnic groups into *computing research careers,"* and inculcating *"values of leadership and service in the computing research community"*) are perhaps too ambitious. The former is addressed by a number of organizations, while the latter seems more appropriate to CRA and Snowbird.

Overall, the implementation is good, but should be adjusted as evaluation data becomes available.

CCC Organization, Management and Budget

There are two parts of CCC management – the CCC Council and the CCC staff, and both play an important role. What is most important to the success of the CCC is that over time, they work well together.

The representation on the CCC Council at this point in time, seems to include broad representation, and the reviewers applaud the care with which the council is formed. We applaud the breakdown in gender, institution type, and research area of the current council. Over time, it is important to continue to think about diversity in its broadest sense – research, institutions, gender, age, as well as networks. But it is also important to include senior, well respected and connected leaders that currently serve on the Council. The leaders like Ed Lazowska, Anita Jones, and Susan Graham (and others) play a unique role. That means it is critical for the executive leadership to carefully think through the desired characteristics of some of the council members.

We applaud the decision to hire a full-time director. This is key to the success of the future of the CCC. A full-time employee can manage the process and progress of CCC activities and its governing body, ensuring that the CCC is focused on appropriate outcomes.

The responsibilities of the council and the director need to mature over time. At this point, the role of council members seems somewhat ad hoc. We encourage the council to think carefully about where they spend their time and to think what activities will have the most valuable impact on the long-term goals of the field. For example, over time the number of proposals may increase substantially, and it is worth understanding how the council would handle a large influx of proposals. The council has the opportunity to identify people and communities who could process the applications, and not necessarily do the work itself. The

other critical opportunity is to bring groups together who are currently stovepiped. In other words, the council should continue to be an enabler. This is work that the council is already doing, but as the organization scales, it is important to understand how to scale the current processes.

It is also important to recognize the importance of the CCC council and the director in following through after a report is delivered, ensuring that new opportunities are identified and communicating to the visioning leadership possible collaborations and funding.

The success to date is tied to the sterling reputation of many of the founders. It is important to think through the transition to new leaders as the CCC scales its vision.

The budget is appropriate for the current activities.

Outcomes and Assessment

Key outcomes for a community-based CCC:

A. *Build and strengthen the Computer Science Community as a community* (Goals 1, 3, 5, 6, 7)

B. *Advance the Computer Science agenda in the research funding community, and increase the influence of Computer Science in the national science community* (Goals 2, 4, 6)

Feedback on Outcome A (Build and strengthen the Computer Science Community as a community):

The long-term health of the Computer Science Community – its ability to create and forward a compelling scientific agenda, its ability to develop community leaders, and its ability to bring together a diverse and engaged set of participants, is critical for its success and survival. The CCC has done a good job at engaging an initial set of leaders and participants to build greater – and sorely needed – community within computer science.

The review committee encourages the CCC take an even bigger role in community building and leadership at all levels:

- *Junior Level*: Innovative programs such as the Computing Innovation Fellows can do much to seed the research community in computer science and provide early mentorship. This activity should be continued and scaled if possible.

- *Senior Level*: The participation of dedicated and outstanding senior leaders has done much to jumpstart CCC. Engagement of a broader set of community leaders would benefit the CCC, spread the workload at the senior level, provide greater mentorship to emerging community leaders, and contribute to a broader set of outreach activities. The CCC should seek to incorporate a broader set of computer science leaders including National Academies' members in CS, leaders of Professional Societies, Senior Computer Scientists serving on Federal and influential Advisory Committees, Awards winners, etc.
- *Rank and File*: There should be more engagement of the broader community of computer scientists who are members of ACM, IEEE, SIGs, etc., and who congregate at professional conferences in their areas. These computer scientists and venues are critical to creating a broad-based computer science vision and agenda, as well as creating openness and opportunities for creating "audacious" ideas. Integrating the "visioning" activities in mainstream community activities will help build both the community and CCC.

Metrics of Success for A:

The CCC needs to develop useful criteria for measuring their success in building the community. Such metrics should be measurable ("clear actionable roadmaps for visionary research"), rather than non-measurable ("societal understanding of the foundational impact of computing research"). Measures such as those used by NSF's BPC program could be adapted to measure the success of CCC community building efforts.

It is also important for CCC metrics to track inclusiveness (has the fraction of the community that is involved with, and benefits from, CCC increased?), as was tracked for the Computing Innovation Fellows Program.

Metrics of Success for B:

Although impact and influence is difficult to quantify, it is important for the CCC to demonstrate return on investment of NSF funding in terms of increased impact and influence of the community on the national agenda. Has the representation of computer scientists in positions of leadership at the national level increased? Have the budgets that federal agencies devote to supporting computer science increased? Developing a set of measurable criteria for success for CCC's impact and influence is important to demonstrate return on investment.

Conclusion

Intellectual Merit and Broader Impacts

The CCC provides vital national functions. It successfully helps policy-makers understand the role of computing research in progressing important societal issues. It helps develop new leaders in the computing research community. It accelerates the pace of the computing and information sciences by convening appropriate internal communities and encouraging them to set appropriately ambitious goals. The reviewers note the success of the CCC white papers, the huge interest in the Computing Innovation Fellows program from both faculty members and applicants, and the CCC's close and useful connections with the National Science Foundation.

The addition of a full-time Executive Director in March 2010 is likely to give the CCC a significant boost. We recognize the efforts by the CCC to be inclusive by gender, by institutional size, and by research area, and encourage them to continue to spend extra effort on inclusivity.

Recommendations

1. Focus on connections. Please continue to take every opportunity to connect the dots between our internal computing research community and external parties who need to better understand the role of computing research in solving their problems. This includes a broad range of science and engineering leaders in other disciplines, funding bodies, and policy-driven organizations. Consider how CCC members can facilitate the placement of thoughtful people with computing research backgrounds on the advisory groups of science and engineering activities of other disciplines.

2. Consider pulling together the grand challenges that emerge from the various visioning groups and issue a "blue book" on the grand challenges in computer science. These grand challenges should be widely communicated within the computer science community and to government agencies and policy makers via multiple means. Some student-oriented challenges may also be appropriate, working with the ACM and IEEE Computer Society.

3. Begin now to seek cross-agency funding for CCC. In addition to diversifying funding, the process of obtaining funding may raise the profile of CCC within those agencies.

4. Take advantage of the SRI evaluation to further energize and drive mindshare for CCC and its mission as well as better understand how the community views CCC and its accomplishments. Focus on evaluating performance

against aspirational goals more than tactical goals. Request separate reports for the CCC and the Computing Innovation Fellows.

5. Now that CCC has sponsored some successful visioning, additional effort should be given by the CCC to hone crisp and understandable research agendas based on those visions and matching those agendas to people and agencies who can advance those agendas in pursuit of their own missions.
6. Don't forget to bring along the needed basic research. In connecting the dots between societal issues and the needed computing research, there will likely be a closer association with the applied side of computing research. However, underlying that applied research are more basic research efforts that also need to be encouraged.
7. Groom the next set of CCC leadership. The CCC officers and council members are all respected members of the community and are reasonably diverse. It's never too early to recruit and groom new leadership candidates and give them appropriate recognition.

The unanimous consensus of the panel is that the CCC is an excellent project that has huge potential payoff not only to the computing research community but to all of science. It recommends the NSF should continue to fund the project at the requested level through year 4. Done well, the benefits of the CCC are likely to extend into the foreseeable future.